Research Article

Slow Light with Photonic Crystals for On-Chip Optical Interconnects

Sean P. Anderson,[1] Ashutosh R. Shroff,[1] and Philippe M. Fauchet[1, 2]

[1] *The Institute of Optics, University of Rochester, Rochester, NY 14627, USA*
[2] *Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627, USA*

Correspondence should be addressed to Sean P. Anderson, sanderso@optics.rochester.edu

Transistor scaling alone can no longer be relied upon to yield the exponential speed increases we have come to expect from the microprocessor industry. The principle reason for this is the *interconnect bottleneck*, where the electrical connections between and within microprocessors are becoming, and in some cases have already become, the limiting factor in overall microprocessor performance. Optical interconnects have the potential to address this shortcoming directly, by providing an inter- and intrachip communication infrastructure that has both greater bandwidth and lower latency than electrical interconnects, while remaining safely within size and power constraints. In this paper, we review the requirements that a successful optical interconnect must meet, as well as some of the recent work in our group in the area of slow-light photonic crystal devices for on-chip optical interconnects. We show that slow-light interferometric optical modulators in photonic crystal can have not only high bandwidth, but also extremely compact size. We also introduce the first example of a multichannel slow light platform, upon which a new class of ultracompact optical devices can be built.

1. INTRODUCTION

Transistor scaling has been the crux of the rapid growth in microprocessor performance over the past forty years [1]. More recently, however, the performance of the electrical interconnects, which are responsible for transporting data within the microprocessor and between the microprocessor and memory, has been unable to keep pace. This is true because as the interconnect is scaled down along with the transistors, resistance and capacitance grow, limiting performance. This not only decreases the bandwidth of the interconnect, but also increases both its latency and power consumption. In fact, in modern microprocessors, over half of the dissipated power is dissipated by the interconnects [2, 3]. These issues will continue to worsen as chip technology continues to scale [3–8]. Optical interconnects can directly address these problems at the system level by replacing electrical interconnects [6–10]. To do so, they must meet the performance requirements of modern and future microprocessors while achieving both compact size and low power consumption.

Implementing optical interconnects in silicon has the advantage of maintaining maximum compatibility with the existing CMOS fabrication infrastructure. Additionally, silicon is transparent in the range of telecom wavelengths ($\lambda \sim 1.5\,\mu$m). However, it is not an ideal choice as an active optical material due to its relatively weak electro-optic response [11]. In general, this results either in large device size or decreased operational bandwidth. The former is true especially for phase-shifting approaches, while the latter applies mostly to resonant approaches. To work around this limitation, it is possible to use some other, extrinsic material as the active medium, such as liquid crystals [12–14], quantum dots [15], or electro-optic polymers [16]. These have the advantage of potentially smaller device size, but at the expense of bandwidth. This is because their response times can be orders of magnitude slower than that of silicon. To date, devices based on such extrinsic materials are unable to meet the bandwidth requirements for optical interconnects. Such hybrid approaches also have the drawback of reduced compatibility with CMOS processing techniques. An all-silicon approach would thus be favorable,

although devices that employ silicon as the active medium have so far been unable to meet the size requirements.

Using slow light, however, it is possible to create a device that meets both of these seemingly contradictory sets of requirements. The advantage of using slow light is that, because the group velocity of light is decreased by two or more orders of magnitude as compared to that in bulk silicon, the effective photon-material interaction length inside an active device is increased [17]. This allows greater use to be made of the small refractive index changes available in silicon, and thus can enable silicon-based photonics to meet the requirements set out by optical interconnects.

In this paper, we first review the rationale behind the push toward optical interconnects, as well as the bandwidth, latency, power, and footprint requirements that an optical interconnect must satisfy in order to be competitive. We then explain how slow light can be used as the basis for an optical interconnect technology that satisfies these requirements. Specifically, we describe how interferometric optical modulators based on slow light in photonic crystals can exhibit not only high bandwidth but also extremely compact size. We also report on recent results regarding a new optical structure that merges slow light with wavelength-division multiplexing to form the first multichannel slow-light platform for silicon photonics. Finally, we outline some of the challenges surrounding the implementation of these novel structures and devices and, more broadly, how a slow-light approach may fit into what has been called "the interconnect era" [18].

2. THE NEED FOR OPTICAL INTERCONNECTS

The electrical interconnects inside modern microprocessors consist of copper wires surrounded by a low-k dielectric [19]. This represents a transition from previous generations of electrical interconnects, which used aluminum wires surrounded by a dielectric. Both the switch from aluminum to copper, and from dielectric to specifically a low-k dielectric, were made to reduce the RC time constant of the interconnect (the former determining the R part, and the latter the C part). Since the RC time constant is essentially a measure of the amount of energy required to operate the interconnect, reducing it allows the interconnect to operate more efficiently. A lower time constant also allows for lower transmission latency and for reduced crosstalk between adjacent wires. Regardless of the materials used, however, the performance of an electrical interconnect, weather measured by bandwidth, latency, power consumption or crosstalk, worsens as its dimensions are scaled down. This is due to the fact that the resistance of a metal wire grows as its cross-sectional area is reduced. The result is a hyperbolic increase in the RC time constant per unit length of interconnect as the chip feature size is scaled down [6]. Optical interconnects, on the other hand, do not suffer from this constraint because they are not subject to an RC time constant.

For this reason, optical interconnects (OIs) offer a number of advantages over electrical interconnects (EIs). The first is a significantly lower signal propagation delay. Figure 1(a) shows the relationship between propagation delay and interconnect length both for EIs based on copper, and for OIs based on silicon waveguides. Silicon waveguides possess an intrinsic advantage over copper wires because of their higher signal propagation speed, which is in turn due to the absence of RC impedances [10]. Because there is an encoding/decoding penalty associated with optical interconnects (the time spent converting the signal from the electrical domain to the optical, and vice versa), OIs based on silicon waveguides may be best suited for longer length interconnects (i.e., global, as opposed to local, interconnects) [10, 20].

A second metric by which OIs must be compared to EIs is bandwidth density. Although the bandwidth available from a single wire in an EI decreases as the chip feature size is scaled down, the cross-sectional area it occupies also decreases. The net result is that the bandwidth density, as measured by the number of bits that can be transmitted per second per unit lateral width, increases with further chip scaling. Additionally, repeaters can be used to further enhance bandwidth, though at the price of increased size and power consumption [21]. This creates a moving target for OIs to beat. Figure 1(b) compares the bandwidth density of EIs to that of OIs, based on reasonable estimates of the size of and spacing between the silicon waveguides. The bandwidth density of OIs is assumed to scale linearly with the clock speed of the chip, while the bandwidth density of EIs scales at a slightly higher rate. EIs, however, have the disadvantage that they can support only a single data channel at a time. OIs, on the other hand, can employ techniques such as wavelength-division multiplexing (WDM) to support a large number of channels simultaneously. Even at the 22 nm node, only three WDM channels are needed to match the performance of EIs [4]. Although the additional channels come at the price of increased power consumption and footprint, these parameters scale much more weakly for OIs than for EIs. This is due to the fact that the extra space and power are required only by additional modulator/demodulator pairs at the ends of the interconnect. The size of the waveguide itself does not grow (and it consumes no power because, unlike with EIs, it contains no repeaters). This gives WDM OIs a significant advantage over EIs, especially for longer-length interconnects.

An additional metric, the power-delay product, is often used to evaluate interconnect performance because it is a measure of both the power consumed, and the delay introduced, by the interconnect. The power-delay product (PDP) of an OI is dependent on both its length and the technique used to modulate the optical signal. Modulation techniques fall generally into two categories: resonant and interferometric. Figure 2 compares the PDP for these two types of modulators to that of EIs as a function of the achievable index shift Δn. For interferometric modulators, Δn determines the length required to achieve a required phase shift; the lower Δn, the greater the length, and thus the greater the power consumption and delay. A resonator-based modulator has the advantage that its active area can be made much smaller than in an interferometric modulator, resulting in both lower delay and lower power consumption. In fact, the PDP for a resonant modulator is nearly two orders
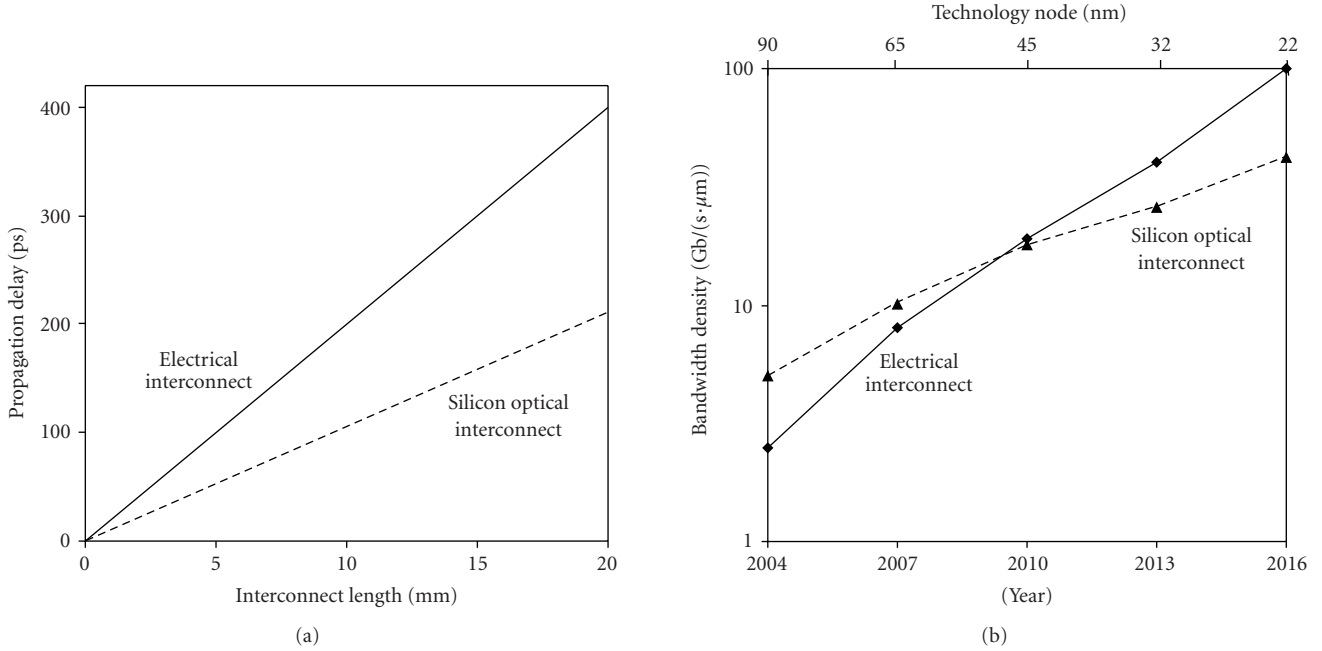
FIGURE 1: (a) Propagation delay of silicon optical waveguides as compared to copper electrical wires (cf. [4]). (b) Comparison of bandwidth density of EIs to that of a single-channel OI as a function of ITRS year and technology node. The use of a multichannel OI can increase bandwidth density over that of EIs (cf. [4]).

of magnitude less than for an interferometric one, assuming each has sufficient bandwidth for a single optical channel. This shows that a resonant approach is generally preferable over one based on interference [4]. Resonant modulators also have the advantage of a smaller footprint, and thus a smaller size penalty for each additional WDM channel. In principle, these advantages come at the price of decreased bandwidth. However, the bandwidth of any interconnect channel is limited by the clock speed of the chip, and that speed is in the range of tens of gigahertz. Therefore, a favorable PDP can still be achieved without using Q-factors so high that they compromise bandwidth.

Optical interconnects thus have advantages in terms of signal propagation delay, power consumption, and bandwidth density. These advantages are especially compelling when the length of the interconnect in question is long. We can then estimate the *critical length* over which OIs are preferable to EIs. This length is plotted in Figure 3 for several technology nodes of the ITRS [22]. The comparison is made separately for three different criteria: signal propagation delay, power consumption, and bandwidth density scaled by delay. In all of these comparisons, OIs are favorable for interconnect lengths over a few millimeters, assuming the use of WDM and resonator-based modulators, with each channel operating at a bit-rate equal to the clock rate of the chip [7, 10, 20].

Therefore, the development of a successful optical interconnect technology must include the development of a compact and low-power modulator, upon which a WDM communication system can be based. However, the size advantage of resonant modulators is not fundamental. In


FIGURE 2: Power-delay product of EI and OIs as a function of Δn for the 90 nm technology node, assuming a length of 10 mm. OIs based on resonator-based modulators offer a significant advantage over both EIs and OIs based on interferometric approaches (cf. [4]).

fact, it is possible to produce interferometric modulators with superior power and size characteristics if the group velocity inside the devices can be sufficiently reduced. This is the principle behind *slow-light* devices. We next review recent work in our group in the area of slow-light interferometric devices in silicon.
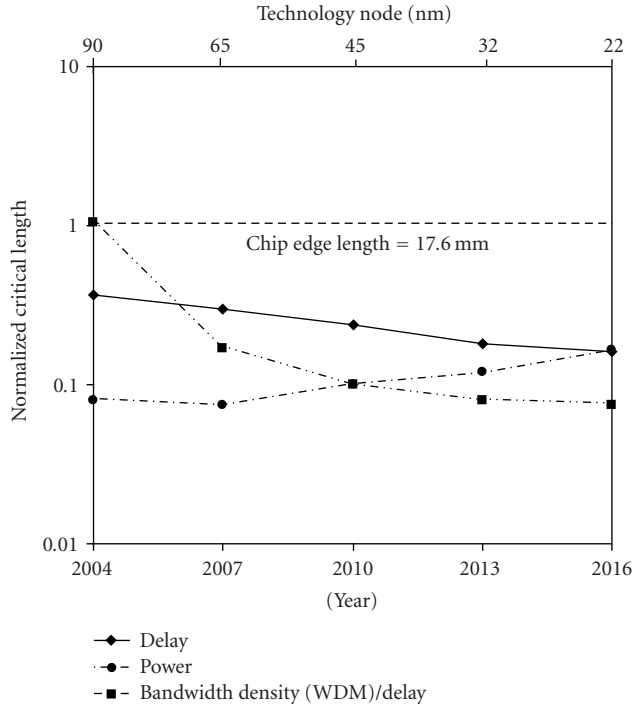
FIGURE 3: Critical length, normalized to the chip edge length, above which OIs (using WDM) are advantageous over EIs in terms of delay, bandwidth density (scaled by delay), and power consumption (cf. [10]).

3. SLOW-LIGHT MACH-ZEHNDER INTERFEROMETERS

Integrated Mach-Zehnder interferometer (MZI) devices are used extensively in optical modulators and switches. Liu et al. demonstrated a silicon high-speed optical modulator with an operational bandwidth of 40 GHz [23]. While the device proves the feasibility of silicon for optoelectronic applications, it suffers from a significant disadvantage: its size is on the order of a millimeter. This is consistent with an earlier analysis presented by Giguere et al. [24]. The large device footprint is a result of the small value of Δn available in silicon. This makes the distance, L_π, required to produce a π phase shift in one arm of the MZI to be very long.

However, Soljačić et al. proposed that by increasing the group index (decreasing the group velocity) of light propagating in the arms of an MZI, the sensitivity to small changes in the material refractive index of the arms could be amplified [17]. More recently, Shi et al. showed experimentally that the sensitivity of an interferometer is dependent on the group index rather than the material index [25]. One way to understand this is that inducing a change δn in the material refractive index in one arm of the interferometer causes the electromagnetic bands to shift in frequency by an amount $d\omega$. Because frequency is kept constant by the choice of the operating wavelength, the propagating wave experiences a change in wavevector magnitude by an amount dk. Therefore, since $n_g = c/v_g = c/(d\omega/dk)$, the larger the group index, the larger the change in wavevector, and thus the shorter the interferometer arms

can be made, because $L_\pi = \pi/dk$. Reducing the device length not only allows for space savings, but also decreases device power consumption. Additionally, the electrodes can be made smaller, thus reducing parasitics and further increasing operating bandwidth.

The photonic crystal platform [26–28] is ideal for implementing such a device because it provides the design flexibility required to increase the group index while allowing the tunability required of an active device. In a photonic crystal coupled-cavity waveguide (PC-CCW), the group velocity can be controlled by changing the spacing between adjacent cavities. This is because the spacing controls the spatial overlap between the optical fields in adjacent cavities, which in turn determines the width (and slope) of the optical miniband and thus its group velocity [29–32]. An MZI based on such a design was proposed and analyzed by Soljačić et al. [17], although that analysis did not include the possibility of optical jitter. The concept of optical jitter in devices is important because it can cause pulse distortion and thus reduced bandwidth. In fact, in MZIs based on slow-light CCWs, deterministic optical jitter is a significant source of pulse distortion at high bit-rates. Further, when multiple devices are cascaded together, uncertainty due to jitter compounds, which can result in asynchronous operation. In slow-light MZIs, deterministic optical jitter grows as the device length is reduced, resulting in a tradeoff between bandwidth and device size reduction. Optical jitter can, however, be minimized by carefully choosing the operating wavelength of the device. This effectively removes the tradeoff between bandwidth and size reduction, but replaces it with a tradeoff between size reduction, and sensitivity of the device to material and fabrication variance. This new tradeoff, however, is a very favorable one because the semiconductor industry excels at minimizing variation.

3.1. Increased sensitivity in slow-light MZIs

Figure 4(a) shows the concept of an MZI in which each arm consists of a PC-CCW. We modeled five different PC-CCWs, denoted by the separation Δ between adjacent cavities, using the MIT Photonic Bands software package [33, 34]. We calculated the TE bands for the structures assuming air holes ($n = 1.0$) of radius $0.3a$ in a silicon slab ($n = 3.4$), where the lattice constant a is chosen to be 400 nm for use at $\lambda = 1.5\,\mu m$. Each cavity consists simply of a missing hole. The Brillouin zone in each case is chosen to include the entire repeat unit cell of length Δ, which determines the distance between adjacent cavities. For ease of comparison, all results are normalized to the lattice constant a. Each PC-CCW has a significant photonic bandgap with a single defect mode (Figure 4(b)). As the separation between adjacent cavities is increased, both the bandwidth and the group velocity of the defect band are reduced because the cavity lifetime of a photon within the defect band is increased. The greater the group index, the shorter the interferometer can be made. For example, if we are able to inject a free carrier concentration of $\Delta N = 10^{18}$ cm^{-3} into the silicon, corresponding to an index change of approximately $\delta n = 0.001$ [11], then the length of the MZI can be reduced to approximately $56\,\mu m$
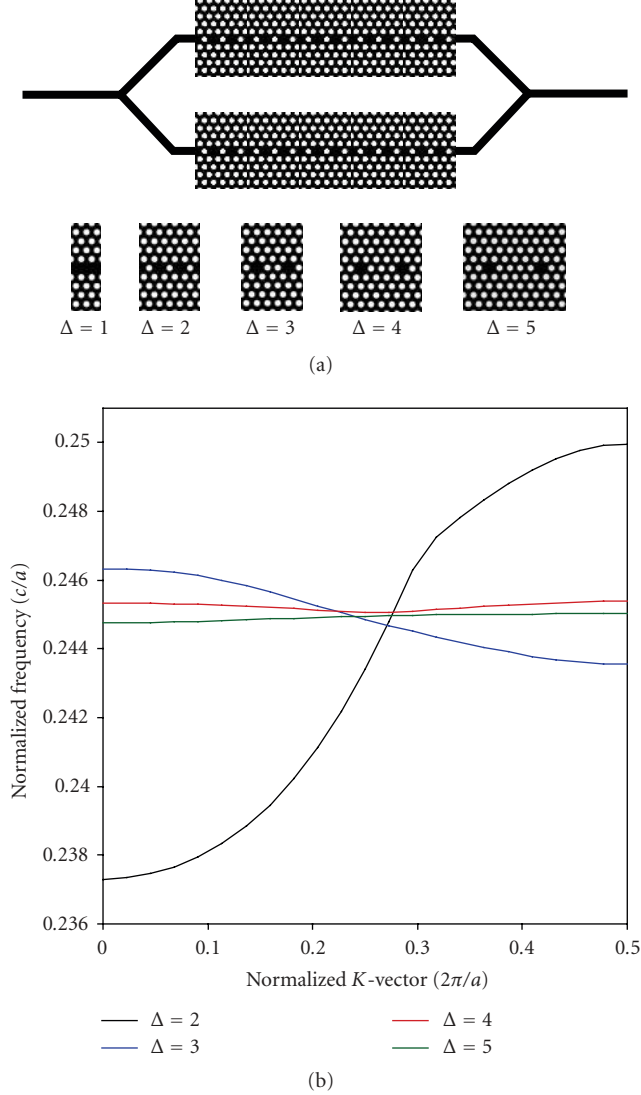
$\Delta = 1$ $\Delta = 2$ $\Delta = 3$ $\Delta = 4$ $\Delta = 5$

(a)

(b)

FIGURE 4: (a) Concept of PC-CCWs used to design ultracompact MZIs. Here Δ is the degree of freedom for varying the device structure. (b) The dispersion curves for the CCW defect bands. As the separation Δ between the cavities is increased, the spatial overlap between the fields localized in each cavity is reduced, thus flattening the dispersion curve and reducing group velocity.

(for the $\Delta = 3$ case), as compared to over 600 μm for an MZI based on a simple photonic crystal line-defect waveguide (the $\Delta = 1$ case). In fact, the arm length could be reduced even further by using PC-CCWs with greater separations Δ between adjacent cavities.

3.2. Optical jitter

When modeling these PC-CCW structures with a small increase δn in the material refractive index induced in the entire silicon slab, we find, as expected, that the bands shift downward in frequency, causing a change Δk in the magnitude of the corresponding wavevector. Associated with this change in wavevector, however, is a change in group


FIGURE 5: When an index shift is induced in one arm of the interferometer, the group velocity in that arm changes. The pulses traveling in the two arms of the interferometer therefore arrive at the output separated in time by Δt, resulting in pulse distortion (broadening). In the worst case, where the pulses interfere destructively, the output pulse can even become double-humped (inset).

velocity. This occurs because the slope of the defect band is not exactly linear at its center, so that when the magnitude of the wavevector changes by Δk, the corresponding point on the dispersion curve has a slightly different slope, and thus a different group velocity. Because the index change is induced in only one arm of the interferometer, the result is that the pulses in the two arms of the MZI do not propagate at the same speed, and thus arrive at the output shifted slightly in time from one another (Figure 5). We have termed this difference in arrival time deterministic optical jitter, Δt. The effect of this type of jitter on the output pulse depends upon both the magnitude of the optical jitter in relation to the pulse width, as well as the bias of the MZI itself. When the optical jitter is comparable to or smaller than the pulse width, the result is pulse broadening in the case of constructive interference, or gross pulse distortion in the case of destructive interference. When the optical jitter is large in comparison to the pulse width, the pulses fail to interfere at all, resulting in two distinct output pulses regardless of the interferometer's bias.

In previous work [17] on slow-light MZI modulators, it was implicitly assumed that $\Delta t = 0$, that is, upon producing a small refractive index shift, δn, in the material, the defect band moves to a higher or lower frequency without changing its slope. While this assumption is true in the case of a single defect in an infinite PC slab, it does not generally hold in the case of PC-CCWs. Because the group velocity is sensitive to changes in the material refractive index, the potential for pulse distortion due to deterministic optical jitter must be taken into account when designing devices based on PC-CCWs. While pulse distortion is insignificant at moderate bit-rates, it becomes more important at higher bit-rates, where Δt becomes significant compared to the FWHM of individual pulses. Depending upon the parameters of the device, jitter can be quite large; over ten picoseconds for the $\Delta = 5$ PC-CCW device, for example. In fact, the amount of optical jitter introduced by the MZI is a function of the separation between adjacent cavities in the PC-CCW. There is thus a tradeoff between the extent to which the arm length can be reduced and the amount of optical jitter introduced (Figure 6). We note also that this behavior is not restricted to the specific geometry used. We have observed the same effect

FIGURE 6: Achievable arm length L_π and optical jitter Δt as functions of the separation Δ between the cavities of the PC-CCWs, assuming $\delta n = 0.001$.

in CCWs based on other cavity geometries, including that proposed by Akahane et al. [35, 36].

To characterize the importance of deterministic optical jitter in PC-CCW MZIs, we compared the pulse distortion it causes to that caused by waveguide dispersion. We assumed a 100 Gbits/s Gaussian pulse train, where the pulses have a FWHM of 3.33 picoseconds, and the 10-picosecond bit-slot thus contains 99% of the pulse energy. Because the pulse bandwidth for a 100 Gbits/s signal is significantly smaller than the channel bandwidths of the PC-CCW defect bands, waveguide dispersion should be minimal. To verify this, we calculated the temporal envelope of a single pulse in the 100 Gbits/s pulse train after propagating through an MZI with arm length L_π and optical jitter Δt. We then compared the resulting output pulse to the input pulse using the $(1 - R^2)$ metric, where $0 \leq R^2 \leq 100\%$. This metric is commonly used in regression analysis to calculate the goodness of fit, and is analogous to the mean-squared error metric [37]. In each case, the amount of pulse distortion due to waveguide dispersion was less than 1%. In comparison, approximately the same level of pulse distortion is reached when Δt is only 0.2 picosecond. (This assumes that the pulses interfere constructively at the output; the distortion would be worse for destructive interference.) Because pulse distortion increases rapidly with Δt, deterministic optical jitter has the potential to be the dominant source of pulse distortion in PC-CCW MZIs.

3.3. Design considerations

In previous work, it was assumed that a small change in the material index would result only in a change in the magnitude of the propagating wavevector, and not in the slope of the dispersion curves. Thus, by operating a PC-CCW MZI in the center of the band ($k = 0.25$, where k is the normalized wavevector) where the dispersion curve is approximately linear, no change in group velocity would be observed. In practice, however, the operating frequency is fixed by the choice of wavelength of the input signal. Therefore, as the dispersion curve shifts up or down, the fixed operating frequency forces a change in the magnitude of the wavevector, and that change may be large enough to shift outside the linear region of the dispersion curve, causing a change in group velocity. This effect is more pronounced in PC-CCWs with larger cavity separations Δ because of the larger changes in k that occur in them. The problem is compounded by the fact that the linear region of the dispersion curve does not necessarily lie exactly at the band center in k-space.

Figure 7 shows the dispersion curves for the $\Delta = 5$ PC-CCW, with and without δn applied. The first operating frequency, ω_1, was chosen because the dispersion curves crossed the band center at approximately that frequency. While this frequency allows for very large changes in the magnitude of the wavevector (k ranges from 0.23 for the $\delta n = 0$ case to 0.29 for the $\delta n = 0.0005$ case, thus resulting in a very short L_π of only 16 μm), the slope of the band changes significantly when the index change is induced, causing the group velocity to decrease from $0.0033c$ to $0.0020c$, a reduction of nearly 40%. (Although the example shown in Figure 7 is an extreme one, and the operating frequency ω_1 would not normally be chosen because of its unfavorable dispersion characteristics, it is illustrative of the issue in point.) By shifting the operating frequency down to ω_2 so that it is closer to the linear regions of the dispersion curves, a much smaller change in group velocity can be obtained, thus greatly reducing optical jitter. Figure 8 shows the arm lengths and values of optical jitter that are achievable when these guidelines are applied to MZIs based on other PC-CCWs. While the arm lengths are slightly increased, the amount of optical jitter is reduced by as much as an order of magnitude by optimizing the operating frequency.

Optical jitter can even be eliminated altogether by choosing the appropriate value of δn. For the $\Delta = 5$ PC-CCW, in fact, choosing $\delta n \approx 0.00055$ results in $\Delta t = 0$ picosecond and $L_\pi < 30 \mu$m (Figure 9). By slightly varying the operating frequency, the "zero-jitter" point can be shifted to other values of δn. For example, choosing $\omega_2 = 0.24486$ as the operating frequency allows for $\Delta t = 0$ at $\delta n \approx 0.0003$, with $L_\pi < 50 \mu$m. The arm lengths given here can be further reduced by a factor of two by employing a "push-pull" design for the MZI, where each arm induces a $\pm \pi/2$ phase shift. Figure 9 also shows, however, that Δt is extremely sensitive to the value of δn being used. Thus, a minute variance in the value of δn could increase Δt to be on the order of tenths of a picosecond, severely limiting the operational bandwidth of the resulting device. This sensitivity is more pronounced for shorter device lengths (greater separation Δ between adjacent cavities, and thus lower group velocity), hence the tradeoff between the size of a slow-light MZI and its tolerance to fabrication variances. This new tradeoff is a very favorable one, however, because the semiconductor industry excels at minimizing fabrication variances, and improves at it with each successive generation of technology.

As long as the refractive index change can be tightly controlled, slow-light MZIs can be made very short while

FIGURE 7: When the operating frequency is chosen to correspond to a normalized frequency of $\omega_1 = 0.24494$, the slope of the dispersion curve changes significantly when δn is applied. When the operating frequency is changed to $\omega_2 = 0.24485$, the slopes become more similar, thus reducing Δt.


FIGURE 8: Achievable arm length L_π and optical jitter Δt as functions of the separation Δ between the cavities when the optical jitter is minimized by optimizing the operating frequencies for each MZI configuration, assuming $\delta n = 0.001$.

still maintaining the bandwidth necessary for on-chip applications. With an arm length of only $30 \, \mu$m and a width under $10 \, \mu$m, slow-light MZIs are competitive with resonant approaches, such as microrings, in terms of on-chip footprint [38]. Furthermore, power consumption scales down along with arm length, so that a $30 \, \mu$m long MZI uses two orders of magnitude less power than, for example, a 3-millimeter long MZI. This brings the power-delay product of slow-light MZIs in line with that of resonator-based modulators, yet with greater bandwidth.

Although the above analysis is specific to slow-light MZIs based on CCWs, the same considerations apply to slow-light MZI approaches in general, including those based on, for example, line-defect waveguides [39, 40]. Just as the slopes of the dispersion curves must be matched in CCWs to minimize optical jitter, so must the slopes be matched in line-defect waveguides or other slow-light media.

4. INTERLACED COUPLED-CAVITY WAVEGUIDE

Even with the use of these ultracompact, high-speed, slow-light MZIs, it is still desirable to use WDM to increase the net bandwidth density of OIs since, in general, the bandwidth of any device is still limited by the clock speed of the chip. To this end, we have also been investigating a novel type of slow-light structure that may be an ideal platform for WDM. It is the first example of which we are aware, of a *slow-light WDM platform* for silicon photonics.

The interlaced coupled-cavity waveguide (ICCW), as we have named it, is a multiresonant structure. Its operation is analogous to that of a normal CCW, where electromagnetic energy couples from one defect cavity to the next, except that there are now multiple cavities, each of which has a different resonant frequency (Figure 10). Thus, light couples from one cavity to the next of the same size, skipping over the intervening cavities. An ICCW therefore exhibits multiple slow-light bands, each corresponding to cavities of a particular radius.

Simulations of a representative ICCW design, using parameters similar to those used for the slow-light MZI above (slab index 3.4, hole index 1.0, lattice constant 400 nm, $r/a = 0.35$), reveal a significant photonic bandgap. In this particular structure, the radii of the defects increase in steps of $0.075a$ from $r/a = 0.0$ to $r/a = 0.225$. There are 10 TE bands in the bandgap, which includes several coupled-cavity modes as well as several waveguide modes. Each of the coupled-cavity modes is localized to a cavity of a particular radius. Figure 11(a) plots the bandstructure of the ICCW, and Figure 11(b) shows the electric field distributions for three of the localized modes. The exact positions in frequency space of these localized bands, and thus the spacing between them, can be adjusted by fine-tuning the radii of the defect cavities. For example, in Figure 11(a), band 4 can be moved such that it lies exactly between bands 3 and 5 by slightly increasing the radius of the second defect, in which band 4 is localized.

Assuming that only one sixteenth of the bandwidth of each band can be used, centered around the zero-GVD points, this ICCW structure exhibits an aggregate bandwidth above 400 Gbits/s at group velocities below $0.004c$. Several other bands also exhibit good dispersion properties, in addition to low group velocities. However, either their bandwidths are smaller due to band-edge effects or their lowest group velocities are only one to two orders of magnitude lower than c. The use of these additional bands can, however, provide a further boost to the total available bandwidth. Pulse propagation studies, performed using the propagation constant β calculated to the second order from the dispersion curves, show that for short distances the device can support high band-widths. Figure 12 shows one such study for the fourth

FIGURE 9: Achievable arm length L_π and jitter Δt as functions of δn for the $\Delta = 5$ PC-CCW for the optimized operating frequency of $\omega = 0.24485$. For a material index change of $\delta n \approx 0.00055$, $\Delta t = 0$ picosecond, and $L_\pi < 30\,\mu$m.


FIGURE 10: Illustration of the concept of a photonic crystal interlaced microcavity. It is a combination of several photonic crystal microcavities into a single waveguide. Repeating the photonic crystal interlaced microcavity yields the ICCW structure.

band in the gap, having an estimated bandwidth of 99 Gbits/s.

4.1. Tuning the slow-light properties of the ICCW

In order to determine the effect of changing the refractive index and radii of the defect cavities, we simulated a second ICCW structure similar to the first but with cavity radii $0.05a$, $0.125a$, $0.20a$, and $0.275a$, corresponding to a $0.05a$ increase over the previous structure. Additionally, the refractive index of the material inside these cavities was set to 1.5 in order to be representative of silicon dioxide, for example, or various active materials such as liquid crystals or electro-optic polymers. This structure exhibits 8 TE bands in its bandgap, which is fewer than the previous structure due to the reduced index contrast. The electric field distributions of the four coupled-cavity modes are shown in Figure 13, and each is localized around a cavity of a particular radius.

We next determined the effect of changing the refractive index of the material in the cavities upon the value of the eigenfrequency of each mode. To do this, we varied the refractive index of the medium in each cavity from 1.5 to 1.501 in steps of 0.0002. Each band exhibits a linear relationship between its eigenfrequency and the refractive index change. This is a very promising feature because it could be exploited to actively tune ICCW-based devices.

The most powerful property of the ICCW, however, is that the cavities can be tuned individually. Figure 14 shows the result of introducing a small index shift (+0.0005) into any cavity. In each case, the mode localized to the cavity in which the index change is introduced experiences a change in frequency on the order of approximately 0.5 GHz (corresponding to a tuning sensitivity df/dn above 1.2 THz/RIU, assuming 1.5 μm wavelength), while the other modes experience only a much smaller change. The fact that the multiple slow-light bands of the ICCW structure are each separately accessible may lead to the development of a variety of novel devices including multichannel modulators and switches, multichannel tunable amplifiers and lasers, or multichannel biosensors. Because all of the channels are contained in a single waveguide, these devices could be made no larger than their single-channel counterparts. An optical modulator based on the ICCW platform would have the advantage of being a resonant modulator, but would be able to handle multiple high-bandwidth optical channels simultaneously. Addressing individual cavities will be challenging, but patterning electrical contacts with size and spacing on the order of one lattice period is possible using UV lithography or direct-write methods, for example. Also, work on single quantum dots has shown that it is possible to manipulate the electrical properties of ultrasmall regions of semiconductor, on the order of hundreds of nanometers in size [41, 42]. Although this work has not yet been extended to photonic crystal devices, doing so should not be beyond the capabilities of state-of-the-art nanofabrication techniques.

It is possible to further optimize the performance of the device by changing the number or radii of the cavities, or the refractive indices of the cavities or the surrounding photonic crystal slab. Simulations of an ICCW with three and five different cavity sizes, for example, have also been carried out with similar results to those above, but with three and five localized modes, respectively.

5. IMPLEMENTATION CHALLENGES

Although photonic crystal slow-light devices have considerable potential in the area of optical interconnects, there are still several issues that must be addressed before they can be practically implemented. One challenge applicable to optical interconnects in general is the availability of an effective light source. While it is possible to develop a highly efficient, off-chip source, coupling its output to the chip can add great cost and complexity to chip package design, in addition to counteracting some of the efficiency advantages. A wafer-bonding approach can help to address some of the complexity issues, but still requires an extra
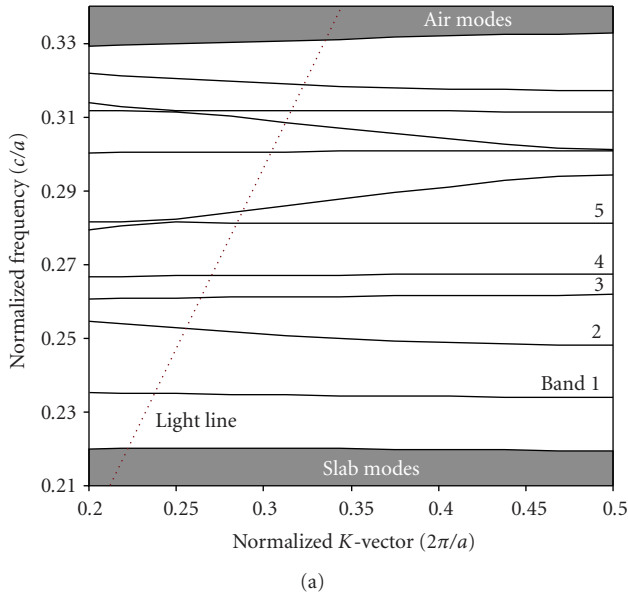
(a)

(b)

Figure 11: (a) Dispersion curves for the photonic crystal modes of the ICCW structure. (b) Electric field distributions for three of the cavity-localized modes.
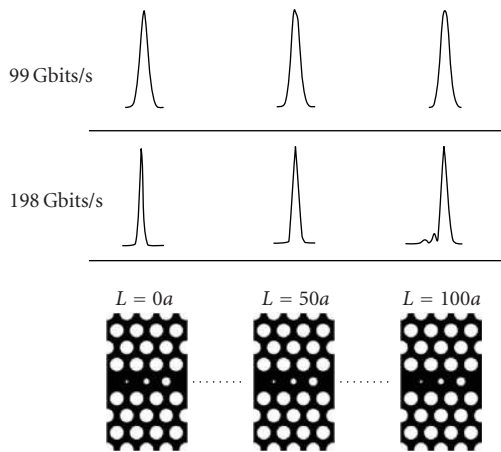

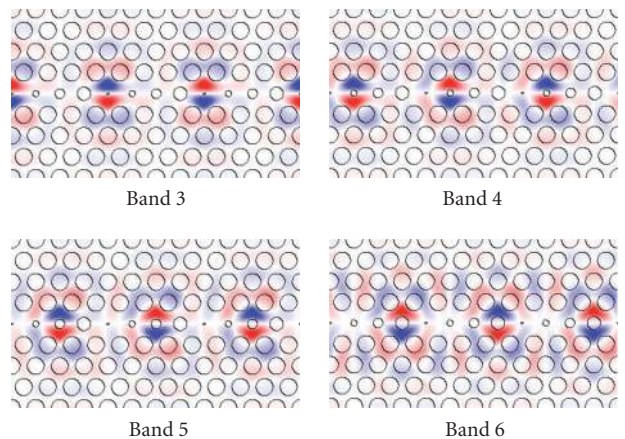Figure 12: Pulse shape evolution for the fourth band in the gap.

Figure 13: Electric field distributions for the four coupled-cavity modes in the ICCW.

layer of fabrication, as well as the inclusion of materials that are not a part of the standard CMOS library [43]. The development of an on-chip light source, however, would eliminate any packaging and output coupling issues, by allowing the light source to be integrated directly on the microprocessor itself [44]. Unfortunately, the constraints of working with silicon (which has an indirect bandgap) in an on-chip platform, where power dissipation, device footprint, and other limitations must be taken into account, complicate the design process. To date, despite evidence for optical gain, there is still no silicon-based laser based on carrier injection [44–46]. Although an optically-pumped silicon laser based on the Raman effect has already been demonstrated [47], an electrically-pumped laser is much more desirable because it would eliminate any need for an off-chip light source. For the

immediate future, a wafer-bonding approach would seem to be the most convenient option, although the inclusion of an electrically-pumped, silicon light source will help to realize the full advantages of optical interconnects.

Several issues specific to the implementation of slow-light devices also exist. The first is that of coupling. In order to maintain the delay and bandwidth density advantages of optical interconnects based on slow-light devices, it is desirable to use silicon wire waveguides as the transmission medium and to use slow-light devices only as the active elements. This is because photonic crystal waveguides, including the ICCW, have a larger cross-section than silicon wire waveguides, the latter of which can have a width under $1\,\mu m$. Additionally, the low group velocities inside coupled-cavity waveguides would increase data propagation delays.
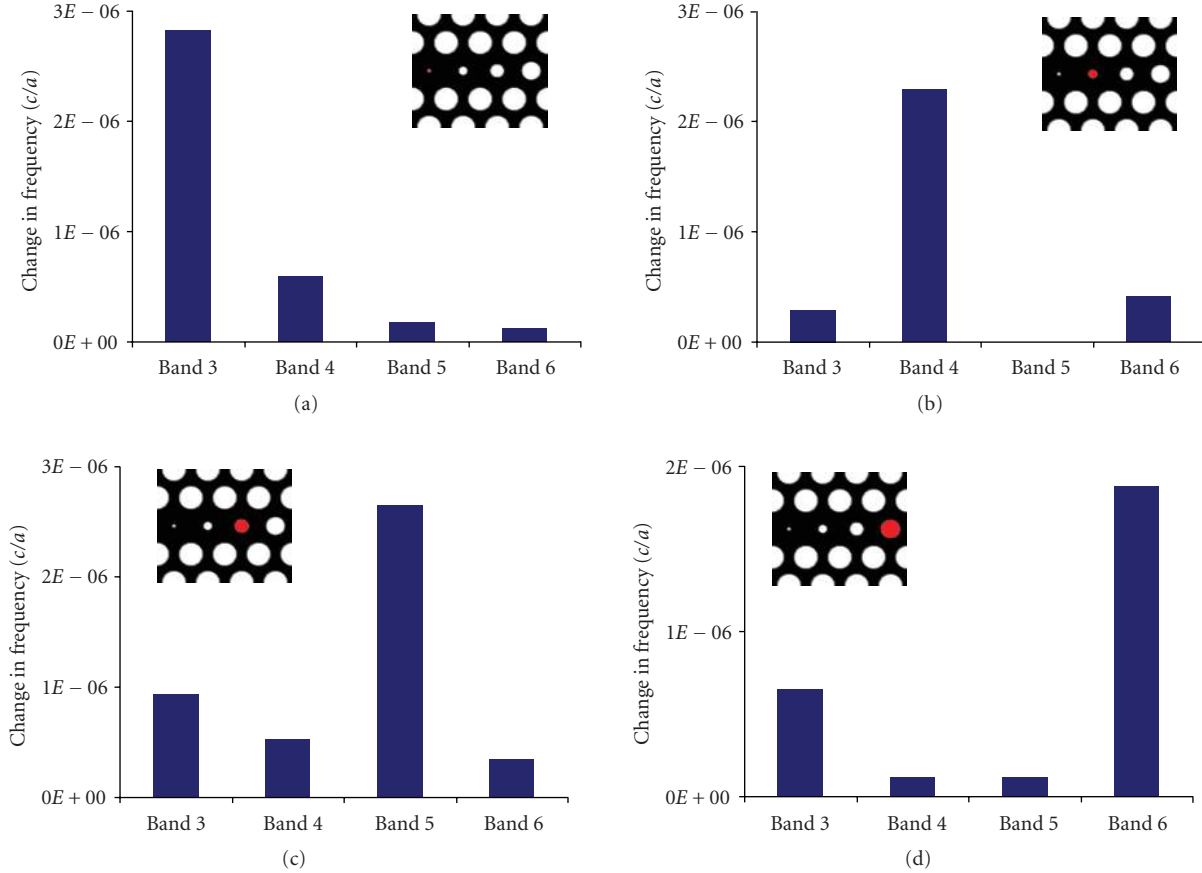
FIGURE 14: Tuning of individual cavities in the ICCW. The cavity tuned in each case is indicated in the schematic.

Therefore, it is necessary to be able to couple the active slow-light devices to traditional silicon wire waveguides. The large mismatch between the group velocities in these two media, however, makes efficient coupling very challenging. An abrupt waveguide-to-CCW interface, for example, causes strong reflection and thus signal loss. Some success has been achieved in coupling silicon ridge waveguides to CCWs [48], although because of the large group velocity difference any approach based on a taper is unlikely to be compact enough for integration [49]. Other methods, such as those based on structural optimization [50, 51] have shown some promise in reducing coupling loss. More recently, another approach based on the use of a photonic crystal waveguide having an intermediate group velocity was demonstrated to have very high coupling efficiency into slow-light modes [49, 52]. While these techniques have so far only been used for coupling into band-edge slow-light modes, they may also prove useful as the basis for coupling into CCW slow-light modes.

A second issue facing slow-light devices is that of insertion loss. Losses in slow-light photonic crystal devices can be divided into two categories: extrinsic loss and intrinsic loss. Extrinsic loss is due to imperfections in the fabricated structure, that is, disorder induced during fabrication. A recent study of slow-light photonic crystal waveguides revealed that extrinsic losses have only a sublinear dependence on group

velocity (proportional to $1/v_g^{1/2}$) [53], which is promising for slow-light devices. However, at very low group velocities (below $0.01c$), extrinsic losses may scale much more strongly and thus may have the potential to become the dominant loss mechanism [53, 54]. Further refinement of fabrication processes can help to control extrinsic losses, though, by reducing the structural disorder that causes it.

Intrinsic loss, on the other hand, exists regardless of fabrication variances. This takes the form of radiation loss, where optical energy leaks out of the structure due to lack of confinement in the vertical direction. In general, these losses grow as the group velocity is decreased. Radiation loss can be controlled, however, by optimizing the quality factor Q of the individual cavities that make up the CCW [55, 56].Much progress has already been made in the area of ultra-high-Q photonic crystal microcavities [35, 36], and it has been shown that the radiation loss of a CCW can actually be made far lower than that of a single microcavity [55]. This can be accomplished in part by optimizing the spatial distribution of the optical modes in the cavities, such as by using a modified cavity geometry. To that end, we have observed that the analysis presented in Section 3 also holds for other cavity geometries, such as that proposed by Akahane et al. [35, 36]. Finally, because of the extremely short lengths of slow-light devices, even with total losses as high as 20 dB/mm, for example, a $50 \, \mu$m long device would

exhibit only 1 dB of net loss, which is well within acceptable levels for optical interconnects.

6. CONCLUSION

The need for an improved on-chip interconnect technology has brought the possibility of optical interconnects to the fore. Before optical interconnects can replace electrical interconnects, however, they must be able to compete in terms of a variety of parameters, including signal propagation delay, bandwidth density, footprint, and power consumption. In terms of signal propagation delay, optical interconnects have a natural advantage over electrical interconnects because of the absence of RC impedances. Likewise, they have an advantage in terms of bandwidth density if wavelength-division multiplexing is used. The use of photonic crystal slow-light devices has the potential to extend this list of advantages to footprint and power consumption as well.

As we have outlined above, the properties of slow light can be used to vastly shrink the size and power consumption of interferometric optical modulators while maintaining very high bandwidth. Furthermore, the ICCW platform, which is the first example of a slow-light WDM platform, has the potential to be the basis for a variety of ultracompact, low-power, high-bandwidth multichannel devices. With further improvements in coupling and the reduction of intrinsic and extrinsic losses, one can envision an optical interconnect built on slow-light photonic crystal active devices. Because of its unique advantages in terms of low latency and high bandwidth even over large distances, such an interconnect may enable the emergence of a new generation of microprocessors that are no longer interconnect-constrained. These chips would be able to leverage the availability of much longer interconnects than are possible using electrical interconnects. Even before that occurs, optical interconnects based on slow-light photonic crystal active devices will be able to function as a drop-in replacement for electrical interconnects in future microprocessors.

ACKNOWLEDGMENTS

REFERENCES

[1] G. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, no. 8, pp. 114–117, 1965.

[2] F. J. Pollack, "New microarchitecture challenges in the coming generations of CMOS process technologies," in *Proceedings of the 32nd Annual ACM/IEEE International Symposium on Microarchitecture*, p. 2, Haifa, Israel, November 1999.

[3] N. Magen, A. Kolodny, U. Weiser, and N. Shamir, "Interconnect-power dissipation in a microprocessor," in *Proceedings of the International Workshop on System Level Interconnect Prediction (SLIP '04)*, pp. 7–13, Paris, France, February 2004.

[4] M. Haurylau, G. Chen, H. Chen, et al., "On-chip optical interconnect roadmap: challenges and critical directions," *IEEE Journal on Selected Topics in Quantum Electronics*, vol. 12, no. 6, part 2, pp. 1699–1705, 2006.

[5] N. Nelson, G. Briggs, M. Haurylau, et al., "Alleviating thermal constraints while maintaining performance via silicon-based on-chip optical interconnects," in *Proceedings of the Workshop on Unique Chips and Systems*, pp. 45–52, Austin, Tex, USA, March 2005.

[6] R. G. Beausoleil, P. J. Kuekes, G. S. Snider, S.-H. Wang, and R. S. Williams, "Nanoelectronic and nanophotonic interconnect," *Proceedings of the IEEE*, vol. 96, no. 2, pp. 230–247, 2008.

[7] M. J. Kobrinsky, B. A. Block, J.-F. Zheng, et al., "On-chip optical interconnects," *Intel Technology Journal*, vol. 8, no. 2, pp. 129–141, 2004.

[8] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proceedings of the IEEE*, vol. 88, no. 6, pp. 728–749, 2000.

[9] J. W. Goodman, F. J. Leonberger, S.-Y. Kung, and R. A. Athale, "Optical interconnections for VLSI systems," *Proceedings of the IEEE*, vol. 72, no. 7, pp. 850–866, 1984.

[10] G. Chen, H. Chen, M. Haurylau, et al., "Predictions of CMOS compatible on-chip optical interconnect," *Integration, the VLSI Journal*, vol. 40, no. 4, pp. 434–446, 2007.

[11] R. A. Soref and B. R. Bennett, "Electrooptical effects in silicon," *IEEE Journal of Quantum Electronics*, vol. 23, no. 1, pp. 123–129, 1987.

[12] S. M. Weiss, M. Haurylau, and P. M. Fauchet, "Tunable photonic bandgap structures for optical interconnects," *Optical Materials*, vol. 27, no. 5, pp. 740–744, 2005.

[13] S. M. Weiss, H. Ouyang, J. Zhang, and P. M. Fauchet, "Electrical and thermal modulation of silicon photonic bandgap microcavities containing liquid crystals," *Optics Express*, vol. 13, no. 4, pp. 1090–1097, 2005.

[14] M. Haurylau, S. P. Anderson, K. L. Marshall, and P. M. Fauchet, "Electrically tunable silicon 2-D photonic bandgap structures," *IEEE Journal on Selected Topics in Quantum Electronics*, vol. 12, no. 6, pp. 1527–1532, 2006.

[15] M. Haurylau, J. Zhang, S. M. Weiss, et al., "Nonlinear optical response of photonic bandgap structures containing PbSe quantum dots," *Journal of Photochemistry and Photobiology A*, vol. 183, no. 3, pp. 329–333, 2006.

[16] R. M. de Ridder, A. Driessen, E. Rikkers, P. V. Lambeck, and M. B. J. Diemeer, "Design and fabrication of electro-optic polymer modulators and switches," *Optical Materials*, vol. 12, no. 2-3, pp. 205–214, 1999.

[17] M. Soljačić, S. G. Johnson, S. Fan, M. Ibanescu, E. Ippen, and J. D. Joannopoulos, "Photonic-crystal slow-light enhancement of nonlinear phase sensitivity," *Journal of the Optical Society of America B*, vol. 19, no. 9, pp. 2052–2059, 2002.

[18] J. D. Meindl, "Beyond Moore's law: the interconnect era," *Computing in Science and Engineering*, vol. 5, no. 1, pp. 20–24, 2003.

[19] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," *Proceedings of the IEEE*, vol. 89, no. 4, pp. 490–504, 2001.

[20] P. Kapur and K. C. Saraswat, "Optical interconnects for future high performance integrated circuits," *Physica E*, vol. 16, no. 3-4, pp. 620–627, 2003.

[21] V. Adler and E. G. Friedman, "Repeater design to reduce delay and power in resistive interconnect," *IEEE Transactions on Circuits and Systems II*, vol. 45, no. 5, pp. 607–616, 1998.

[22] http://public.itrs.net.

[23] A. Liu, L. Liao, D. Rubin, et al., "High-speed optical modulation based on carrier depletion in a silicon waveguide," *Optics Express*, vol. 15, no. 2, pp. 660–668, 2007.

[24] S. R. Giguere, L. Friedman, R. A. Soref, and J. P. Lorenzo, "Simulation studies of silicon electro-optic waveguide devices," *Journal of Applied Physics*, vol. 68, no. 10, pp. 4964–4970, 1990.

[25] Z. Shi, R. W. Boyd, D. J. Gauthier, and C. C. Dudley, "Enhancing the spectral sensitivity of interferometers using slow-light media," *Optics Letters*, vol. 32, no. 8, pp. 915–917, 2007.

[26] J. D. Joannopoulos, R. D. Meade, and J. N. Winn, *Photonic Crystals: Molding the Flow of Light*, Princeton University Press, Princeton, NJ, USA, 1995.

[27] S. John, "Strong localization of photons in certain disordered dielectric superlattices," *Physical Review Letters*, vol. 58, no. 23, pp. 2486–2489, 1987.

[28] E. Yablonovitch, "Inhibited spontaneous emission in solid-state physics and electronics," *Physical Review Letters*, vol. 58, no. 20, pp. 2059–2062, 1987.

[29] A. Yariv, Y. Xu, R. K. Lee, and A. Scherer, "Coupled-resonator optical waveguide: a proposal and analysis," *Optics Letters*, vol. 24, no. 11, pp. 711–713, 1999.

[30] S. Mookherjea and A. Yariv, "Coupled resonator optical waveguides," *IEEE Journal on Selected Topics in Quantum Electronics*, vol. 8, no. 3, pp. 448–456, 2002.

[31] A. Martínez, A. García, P. Sanchis, and J. Martí, "Group velocity and dispersion model of coupled-cavity waveguides in photonic crystals," *Journal of the Optical Society of America A*, vol. 20, no. 1, pp. 147–150, 2003.

[32] T. F. Krauss, "Slow light in photonic crystal waveguides," *Journal of Physics D*, vol. 40, no. 9, pp. 2666–2670, 2007.

[33] S. G. Johnson and J. D. Joannopoulos, "Block-iterative frequency-domain methods for Maxwell's equations in a planewave basis," *Optics Express*, vol. 8, no. 3, pp. 173–190, 2001.

[34] http://ab-initio.mit.edu/mpb.

[35] Y. Akahane, T. Asano, B.-S. Song, and S. Noda, "High-Q photonic nanocavity in a two-dimensional photonic crystal," *Nature*, vol. 425, no. 6961, pp. 944–947, 2003.

[36] Y. Akahane, T. Asano, B.-S. Song, and S. Noda, "Fine-tuned high-Q photonic-crystal nanocavity," *Optics Express*, vol. 13, no. 4, pp. 1202–1214, 2005.

[37] N. R. Draper and H. Smith, *Applied Regression Analysis*, Wiley Series in Probability and Statistics, Wiley-Interscience, New York, NY, USA, 1998.

[38] M. Lipson, "Compact electro-optic modulators on a silicon chip," *IEEE Journal on Selected Topics in Quantum Electronics*, vol. 12, no. 6, pp. 1520–1526, 2006.

[39] M. D. Settle, R. J. P. Engelen, M. Salib, A. Michaeli, L. Kuipers, and T. F. Krauss, "Flatband slow light in photonic crystals featuring spatial pulse compression and terahertz bandwidth," *Optics Express*, vol. 15, no. 1, pp. 219–226, 2007.

[40] L. H. Frandsen, A. V. Lavrinenko, J. Fage-Pedersen, and P. I. Borel, "Photonic crystal waveguides with semi-slow light and tailored dispersion properties," *Optics Express*, vol. 14, no. 20, pp. 9444–9450, 2006.

[41] S. K. Jung, S. W. Hwang, D. Ahn, J. H. Park, Y. Kim, and E. K. Kim, "Fabrication of quantum dot transistors incorporating a single self-assembled quantum dot," *Physica E*, vol. 7, no. 3, pp. 430–434, 2000.

[42] S. Sasaki, S. De Franceschl, J. M. Elzerman, et al., "Kondo effect in an integer-spin quantum dot," *Nature*, vol. 405, no. 6788, pp. 764–767, 2000.

[43] A. W. Fang, H. Park, O. Cohen, R. Jones, M. J. Paniccia, and J. E. Bowers, "Electrically pumped hybrid AlGaInAs-silicon evanescent laser," *Optics Express*, vol. 14, no. 20, pp. 9203–9210, 2006.

[44] P. M. Fauchet, "Monolithic silicon light sources," in *Silicon Photonics*, L. Pavesi and D. J. Lockwood, Eds., vol. 94 of *Topics in Applied Physics*, pp. 177–198, Springer, Berlin, Germany, month 2004.

[45] P. M. Fauchet, J. Ruan, H. Chen, et al., "Optical gain in different silicon nanocrystal systems," *Optical Materials*, vol. 27, no. 5, pp. 745–749, 2005.

[46] L. Pavesi, "Routes toward silicon-based lasers," *Materials Today*, vol. 8, no. 1, pp. 18–25, 2005.

[47] H. Rong, R. Jones, A. Liu, et al., "A continuous-wave Raman silicon laser," *Nature*, vol. 433, no. 7027, pp. 725–728, 2005.

[48] D. O'Brien, M. D. Settle, T. Karle, A. Michaeli, M. Salib, and T. F. Krauss, "Coupled photonic crystal heterostructure nanocavities," *Optics Express*, vol. 15, no. 3, pp. 1228–1233, 2007.

[49] P. Velha, J. P. Hugonin, and P. Lalanne, "Compact and efficient injection of light into band-edge slow-modes," *Optics Express*, vol. 15, no. 10, pp. 6102–6112, 2007.

[50] E. Miyai and S. Noda, "Structural dependance of coupling between a two-dimensional photonic crystal waveguide and a wire waveguide," *Journal of the Optical Society of America B*, vol. 21, no. 1, pp. 67–72, 2004.

[51] Y. A. Vlasov and S. J. McNab, "Coupling into the slow light mode in slab-type photonic crystal waveguides," *Optics Letters*, vol. 31, no. 1, pp. 50–52, 2006.

[52] J. P. Hugonin, P. Lalanne, T. P. White, and T. F. Krauss, "Coupling into slow-mode photonic crystal waveguides," *Optics Letters*, vol. 32, no. 18, pp. 2638–2640, 2007.

[53] L. O'Faolain, T. P. White, D. O'Brien, X. Yuan, M. D. Settle, and T. F. Krauss, "Dependence of extrinsic loss on group velocity in photonic crystal waveguides," *Optics Express*, vol. 15, no. 20, pp. 13129–13138, 2007.

[54] S. Hughes, L. Ramunno, J. F. Young, and J. E. Sipe, "Extrinsic optical scattering loss in photonic crystal waveguides: role of fabrication disorder and photon group velocity," *Physical Review Letters*, vol. 94, no. 3, Article ID 033903, 4 pages, 2005.

[55] M. L. Povinelli and S. Fan, "Radiation loss of coupled-resonator waveguides in photonic-crystal slabs," *Applied Physics Letters*, vol. 89, no. 19, Article ID 191114, 3 pages, 2006.

[56] D. P. Fussell and M. M. Dignam, "Engineering the quality factors of coupled-cavity modes in photonic crystal slabs," *Applied Physics Letters*, vol. 90, no. 18, Article ID 183121, 3 pages, 2007.

Submit your manuscripts at
http://www.hindawi.com

Hindawi