An Energy-Efficient Silicon Microring Resonator-Based Photonic Transmitter

Cheng Li, Chin-Hui Chen, Marco Fiorentino, and Raymond G. Beausoleil
HP Laboratories

Binhao Wang and Samuel Palermo
Texas A&M University

Editors' notes:
The design of two silicon microring resonator-based photonic transmitters, which were fabricated in the 130-nm SOI process and wire bonded with drivers in the 65-nm CMOS technology, is presented. The data rate of the proposed photonic transmitter can reach up to 9 Gb/s.
—Sudeep Pasricha, Colorado State University and Yi Xu, Macau University of Science and Technology

Optical channels are potential candidates to replace conventional electrical channels for efficient interchip and intrachip interconnects due to their attractive properties: flat channel loss over a wide frequency range and strong immunity to crosstalk and electromagnetic noise. An important feature of optical interconnects is the ability to combine multiple data channels on a single waveguide via wavelength division multiplexing (WDM) to greatly improve bandwidth density and amortize connector costs over high aggregate bandwidth.

To take full advantage of these benefits, silicon photonic platforms are being developed that enable tightly integrated optical interconnects and novel photonic network architectures. One promising photonic device is the silicon microring resonator [1], [2] which can be configured either as an optical modulator or a WDM drop filter. Silicon ring resonator modulators and filters offer the advantages of small size relative to Mach–Zehnder modulators (MZMs), and increased filter functionality relative to electro-absorption modulators.

This paper presents silicon ring resonator-based photonic transmitter prototypes that address the limited intrinsic bandwidth of carrier-injection ring modulators, achieving energy-efficient high-speed optical modulation in a compact silicon area suitable for on-chip WDM interconnects.

Digital Object Identifier 10.1109/MDAT.2014.2347932
Date of publication: 28 August 2014; date of current version: 07 October 2014.

Silicon microring resonator-based photonic WDM link

Silicon microring resonator-based photonic links provide a unique opportunity to deliver distance-independent connectivity whose pin bandwidth scales with the degree of WDM. As Figure 1 shows, multiple wavelengths generated by an off-chip continuous-wave (CW) laser are coupled into a silicon waveguide via a grating coupler. This off-chip laser can be either a distributed feedback (DFB) laser bank (which consists of an array of DFB laser diodes) or a comb laser [3] (which can generate multiple wavelengths simultaneously). Implementing a DFB laser bank for dense WDM (DWDM) photonic interconnects (for example, using 64 wavelengths) is quite challenging due to area and power budget constraints.

A possible alternative is a single broad-spectrum comb laser source, such as an indium arsenide

Copublished by the IEEE CEDA, IEEE CASS, IEEE SSCS, and TTTC

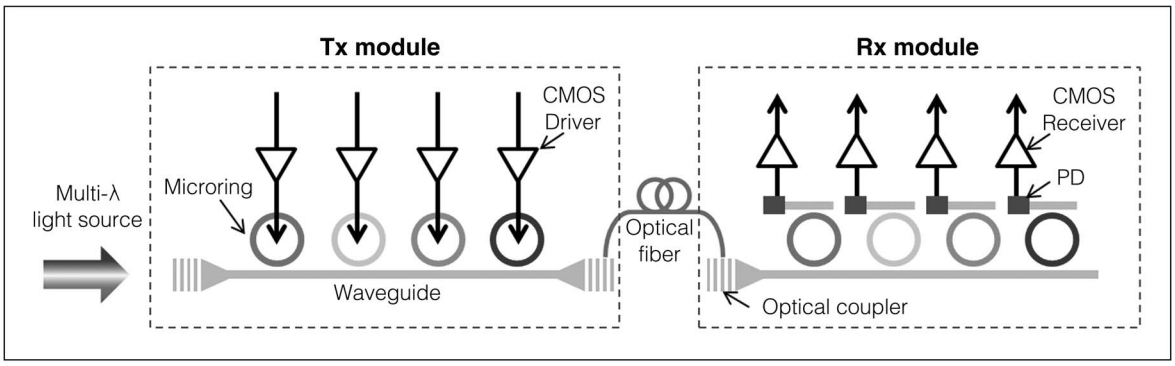

Figure 1. Silicon ring resonator-based WDM link (Rx: receiver; Tx: transmitter).

(InAs) or gallium arsenide (GaAs) quantum dot comb laser that can generate a large number of (for example, 16–64) wavelengths in the 1100–1320-nm spectral range, with typical channel spacing of 50–100 GHz and optical power of 0.2–1 mW per channel [3]. A system operating near a wavelength of 1310 nm (O-band) incurs slightly higher optical loss in fiber compared to a 1550-nm (C-band) system. However, this has negligible impact in short-reach interconnect applications. At the transmitter side, each ring modulator inserts data onto a specific wavelength through electro–optical modulation. The modulated optical signals propagate through the link optical waveguides and arrive at the receiver side. There, ring filters drop the modulated optical signals of a specific wavelength at a photodetector (PD), which converts the signals back to the electrical domain.

Silicon ring modulator modeling

A basic silicon ring modulator contains a straight waveguide coupled with a circular waveguide that has a diameter in the tens of micrometers, as Figure 2a shows. At the resonance wavelength, most of the input light is coupled into the circular waveguide, and only a small amount of light can be observed at the through port. As a result, the through-port spectrum displays a notch-shaped characteristic (Figure 2c). It is possible to shift this resonance by changing the waveguide's effective refractive index through the free-carrier plasma dispersion effect to implement the optical modulation [4]. For example, the ring modulator exhibits a low optical power level at the through port when the resonance is aligned well with the laser wavelength,

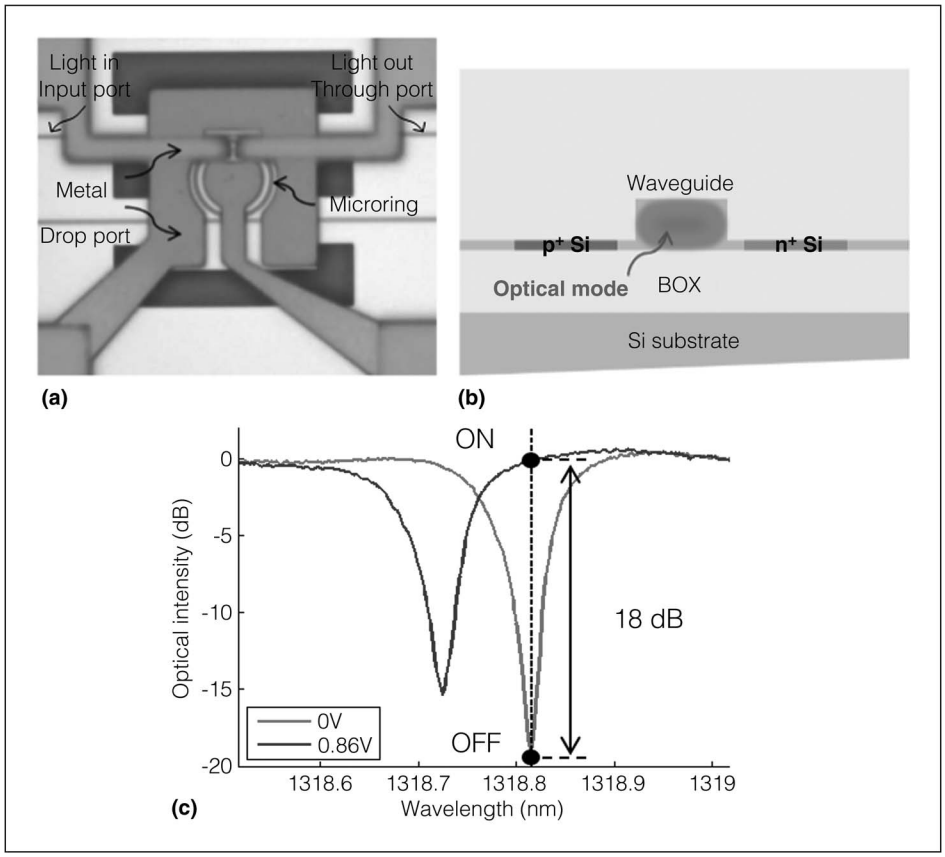

Figure 2. Carrier-injection silicon ring resonator modulator: (a) top view, (b) cross section, and (c) optical spectrum at through port.

whereas a high optical power level is displayed when the resonance blue shifts because of an increase in the waveguide's carrier density that lowers the waveguide's effective refractive index.

Two common implementations of silicon ring resonator modulators include:

- carrier-injection devices [2] containing an embedded p–i–n junction, side-coupled with the circular waveguide, operating primarily in forward bias;
- carrier-depletion devices [1] with only a p–n junction, side-coupled, operating primarily in reverse bias.

A depletion ring generally achieves higher modulation speeds relative to a carrier-injection ring, thanks to the former's ability to rapidly sweep the carriers out of the junction. However, a depletion ring's modulation depth is limited due to the relatively low doping concentration in the waveguide to avoid excessive optical loss. In contrast, carrier-injection ring modulators can provide large refractive-index changes and high modulation depths, but are limited by the relatively slow carrier dynamics of the forward-biased p–i–n junction.

As Figure 2c shows, when a forward-bias voltage is applied over the p–i–n junction of the carrier-injection ring, the resonance shifts toward the shorter wavelength (called a blue shift) because of the accumulated carriers changing the waveguide refractive index. A reverse-bias voltage extracts the carriers accumulated in the junction during the forward-bias modulation and restores the waveguide refractive index. The ring modulator bandwidth is limited by the slow carrier-injection operation due to the relatively slow carrier dynamics of the forward-biased p–i–n junction [2]. Increasing the optical rising transition by simply applying a high modulation swing leads to a slow optical falling transition and causes intersymbol interference (ISI) because the over-injected carriers need more time to be swept out of the ring waveguide.

Although a pre-emphasis modulation scheme has been proposed to break the tradeoff between the optical rising and falling transitions [2], a major ring modulator driver design challenge is the lack of accurate models to predict the high-speed optical modulation signal quality under different pre-emphasis duration and voltage levels. We have developed an accurate Spice model of a silicon ring modulator to enable an efficient cosimulation with electrical driver circuits. A large-signal Spice p–i–n model predicts the carrier distribution in the intrinsic region of the p–i–n junction switched under the modulation signals, and a ring macromodel catches the dynamic electro-optic effects between the junction carrier density and the silicon ring waveguide refractive index at a wavelength close to 1300 nm, as described by

$$\Delta(n) = (6.2 \times 10^{-22} \times \Delta N + (5.9 \times 10^{-18} \times \Delta P^{0.8}) \quad (1)$$

where ΔN (cm^{-3}) and ΔP (cm^{-3}) are the silicon waveguide refractive-index changes due to the electron and hole concentration changes, respectively. The $\Delta(n)$ of 1.5×10^3 was found at a wavelength of 1300 nm, with the injection of 10^{18} carriers/cm^3 [4].

Figure 3 shows the device model simulation results of the first-generation carrier-injection silicon ring modulators [5], with positive and negative 200-ps pulse responses overlaid. The 5-μm diameter ring device exhibits a quality factor of about 9000. A simple $2V_{pp}$ non-return-to-zero (NRZ) modulation produces the excessively long optical rise time shown in Figure 3a, mainly because the 1-V forward-bias voltage is not high enough to overcome the slow carrier dynamics in the p–i–n junction. Increasing the modulation swing to $4V_{pp}$ dramatically improves the optical rise time, at the expense of high-level ringing and a high steady-state charge value. Unfortunately, this large amount of charge results in a slow optical fall time because of the modulator's series resistance (about 2 KΩ), limiting the drift current to extract the excess carriers in the junction. As a result, a deteriorated extinction ratio (Figure 3b) is observed relative to the $2V_{pp}$ NRZ modulation case.

We use a pre-emphasis modulation technique [2] to address the conflicting requirements for fast rising and falling transitions. During a rising-edge transition, the positive voltage overshoots (2 V) for a fraction of a bit period to allow for a high initial charge before settling to a lower voltage (1 V), corresponding to a reduced steady-state charge. A similarly shaped waveform is used for the falling-edge transition in order to increase the drift current to extract the carriers. Because the rising- and falling-edge time constants differ, a nonlinear modulation waveform is applied. We adjust the amount of overshoot/undershoot time of the pre-emphasis waveform for a

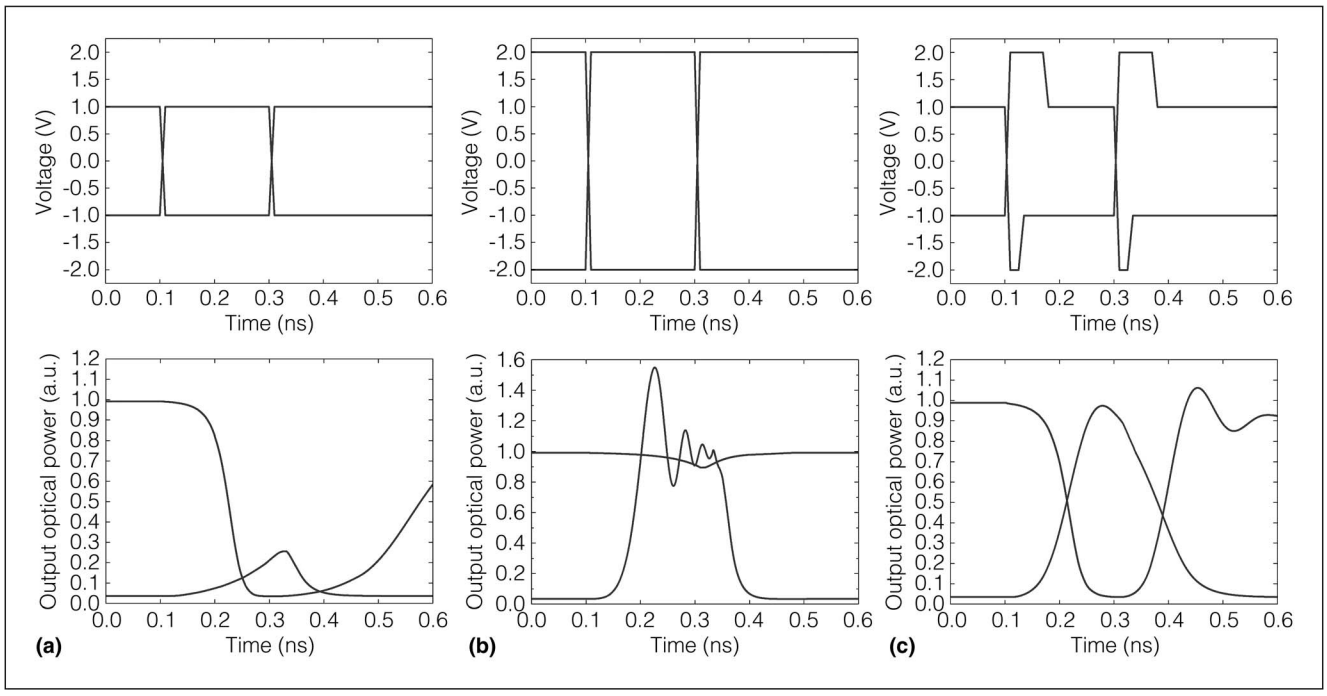

Figure 3. Simulated first-generation ring resonator modulator response (bottom) to 200-ps data pulses (top) with (a) $2V_{pp}$ simple modulation, (b) $4V_{pp}$ simple modulation, and (c) $4V_{pp}$ modulation with pre-emphasis.

specific modulator, with the rising-edge pre-emphasis pulse typically wider than the falling edge. Adjusting the pre-emphasis time rather than using different voltage levels lets us decouple the optimization of the transient response from the steady-state extinction ratio value. Figure 3c shows a fast optical rising and falling with the pre-emphasis modulation.

The major issue of the first-generation ring modulator is that the large series contact resistance (about 2 KΩ) requires a high modulation swing ($4V_{pp}$) to compensate for the voltage overhead on the large 2-KΩ series contact resistance, which also limits the drift current to extract the excess carriers in the junction and leads to a slow optical falling transition. The second-generation ring modulator reduces the series contact resistance down to about 200 Ω [6], providing the potential for high-speed, energy-efficient optical modulation. Unlike the $4V_{pp}$ driver, which outputs a differential voltage swing with an average bias level of approximately 0 V on the p–i–n diode, the $2V_{pp}$ single-ended driver provides a $2V_{pp}$ output swing on the modulator cathode and uses a nonlinear voltage digital-to-analog converter (DAC) on the anode, with adjustable direct current (dc)-bias levels for an optimized eye opening.

The 9-b segmented bias DAC consists of a coarse 3-b, nonlinear R-string DAC to match the p–i–n *I*–*V* characteristics, and a fine 6-b linear *R*–2*R* DAC to achieve linear voltage steps on each nonlinear voltage segment [7]. To overcome the relatively slow carrier dynamics in forward bias, the anode is biased at a voltage level close to the p–i–n junction

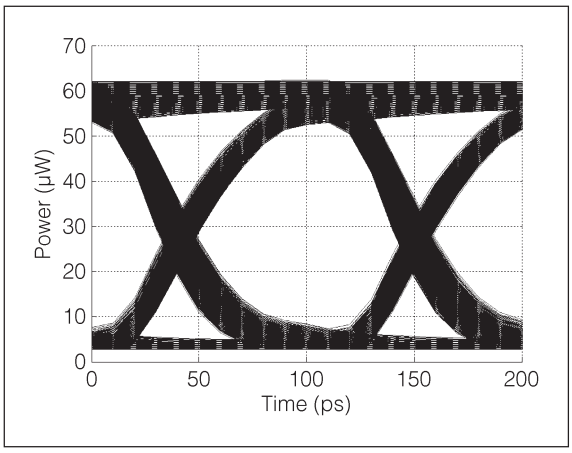

Figure 4. Simulated second-generation ring modulator 9-Gb/s optical eye diagram driven by the $2V_{pp}$ CMOS driver.

threshold voltage. The resonance wavelength blue shifts to shorter wavelengths because of the accumulation of free carriers in the ring waveguide. Therefore, when the resonator p–i–n diode anode voltage is increased, the bias DAC can also be used for bias tuning to compensate for the resonance drifts due to the fabrication variation [7], allowing for both improved tuning power efficiency and speed relative to heater-based tuning [8].

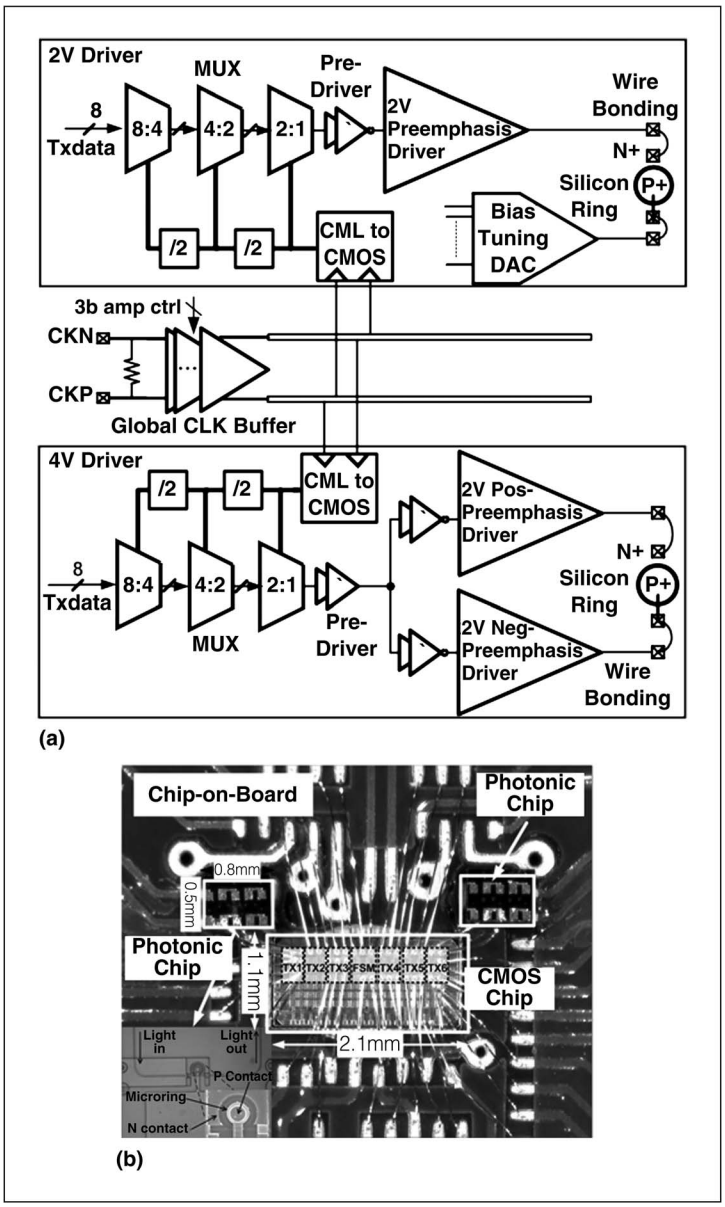

Figure 5. (a) WDM transmitter architecture, and (b) prototype of optical transmitter circuits bonded for optical testing (CLK: clock).

Figure 4 shows the cosimulation result of the second-generation ring prototype based on the proposed ring model with CMOS $2V_{pp}$ pre-emphasis driver circuits. An optimum 9-Gb/s optical eye is achieved when the pre-emphasis duration is set to 80 ps, and the anode is biased at 1.45 V.

WDM transmitter architecture

Figure 5a shows a block diagram of the CMOS WDM photonic transmitter prototype integrating five transmitter modules in a 1-mm^2, 65-nm CMOS area, with one transmitter used as forwarded clocking, and the other four used as data transmission in the four-channel data WDM link. Applying a forwarded-clock architecture in a photonic WDM system offers the potential for improved high-frequency jitter tolerance with minimal jitter amplification, because the clock and data signals experience the same delay over the common low-dispersive optical channel.

Two versions of the CMOS drivers are implemented to modulate the two generation designs of the carrier-injection ring modulators. A differential driver, with an average bias level of approximately 0 V, provides a $4V_{pp}$ output swing to allow for high-speed operation of the first-generation ring modulator with a relatively large series contact resistance (about 2 KΩ). A single-ended driver delivers a $2V_{pp}$ output swing on the second-generation ring modulator cathode and uses a nonlinear bias-tuning DAC on the anode for an adjustable dc-bias level. A half-rate current mode logic (CML) clock is distributed to the five CMOS transmitter modules, where 8-b parallel data are multiplexed to the full output data rate by cascade 2 : 1 multiplexers before being buffered by the modulator drivers. The distributed CML clock is converted to CMOS levels by the local CML-to-CMOS buffer. These CMOS drivers are wire bonded to carrier-injection silicon ring resonator modulators (Figure 5b). A continuous-wavelength light near 1300 nm from a tunable laser is vertically coupled into the photonic device's input port via the grating coupler. The modulated light is then coupled out from the modulator's through port into a multimode fiber, for routing to the optical oscilloscope for high-speed data recovery and eye measurement.

Nonlinear pre-emphasis modulator driver

The $4V_{pp}$ and $2V_{pp}$ pre-emphasis drivers employ similar circuit architectures. An on-chip $(2^7 - 1)$

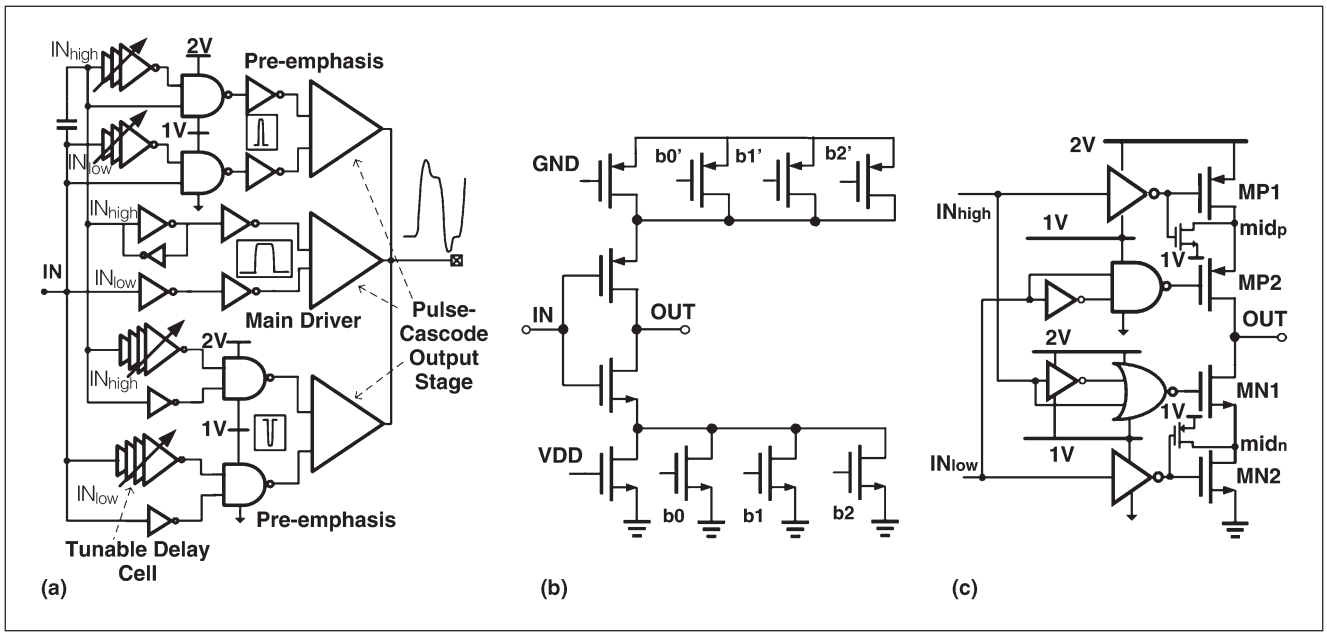

Figure 6. Nonlinear pre-emphasis modulator driver transmitters: (a) per-terminal 2-V pre-emphasis driver, (b) tunable delay cell, and (c) pulsed-cascode output stage.

pseudorandom binary sequence (PRBS) source generates 8-b parallel outputs. Serialization of 8-b data is performed in both transmitter versions with three 2 : 1 multiplexing stages. The serialization clocks are generated from a half-rate CML clock, which is distributed to five transmitter modules, converted to CMOS levels, and subsequently divided to switch the multiplexer stages. The modulator drivers then transmit the serialized data, with both output stage versions using a main driver, and positive- and negative-edge pre-emphasis pulse drivers in parallel (Figure 6a) to generate the pre-emphasis output waveform. Tunable delay cells, implemented with digitally adjustable current-starved inverters (Figure 6b), allow for independent control of the rising- and falling-edge pre-emphasis pulse duration over a range of 20–100 ps. Finally, pulsed-cascode output stages (Figure 6c) with only thin-oxide core devices reliably provide a final per-terminal output swing of twice the nominal 1-V supply.

A capacitive-level shifter and parallel logic chain generate the signals that drive the final pulsed-cascode output stages: IN_{low}, swinging between ground (GND) and the nominal V_{DD}; and IN_{high}, level-shifted between V_{DD} and $2 \times V_{DD}$. During an output transition from high to low, the IN_{low} input switches MN_2 to drive node mid_n to near GND, and the IN_{high} input triggers a positive pulse from the level-shifted nor-pulse gate. This nor-pulse gate drives the gate of MN_1 so that the output can begin discharging at roughly the same time that the MN_1 source is being discharged. Similarly, during an output transition from low to high, the IN_{high} input switches MP_1 to drive node mid_p to near 2 V, and the IN_{low} input triggers a negative pulse from the nand-pulse gate that drives the gate of MP_2 so that the output can begin discharging at roughly the same time that the MP_2 source is being charged. This scheme guarantees that the drain-source voltage does not stress the output PMOS and NMOS transistors with the nominal 1-V supply.

Experimental results

The CMOS driver circuits were fabricated in a 65-nm CMOS general-purpose process. As the photographs in Figure 5b show, a chip-on-board test setup is used, with the CMOS drivers wire bonded to the two generations of silicon ring resonator chips for optical signal characterization. For high-speed optical testing, a continuous-wavelength laser is first coupled into a single-mode fiber probe, which is then coupled to a waveguide connected with a silicon ring resonator through a grating coupler. The current version of the grating coupler used in this work exhibits 7-dB loss because of the simplified structure of the grating that is etched at the same

time as the waveguide. In future work, further improvement can be achieved with more sophisticated two-mask gratings, which have demonstrated losses down to 2–3 dB. The waveguide loss is measured to be 3 dB/cm, which is negligible for the 480-μm waveguide. Overall, with 1-mW optical power from the CW laser source, around 40 μW is detected at the ring's through-port output when the ring is on off-resonance.

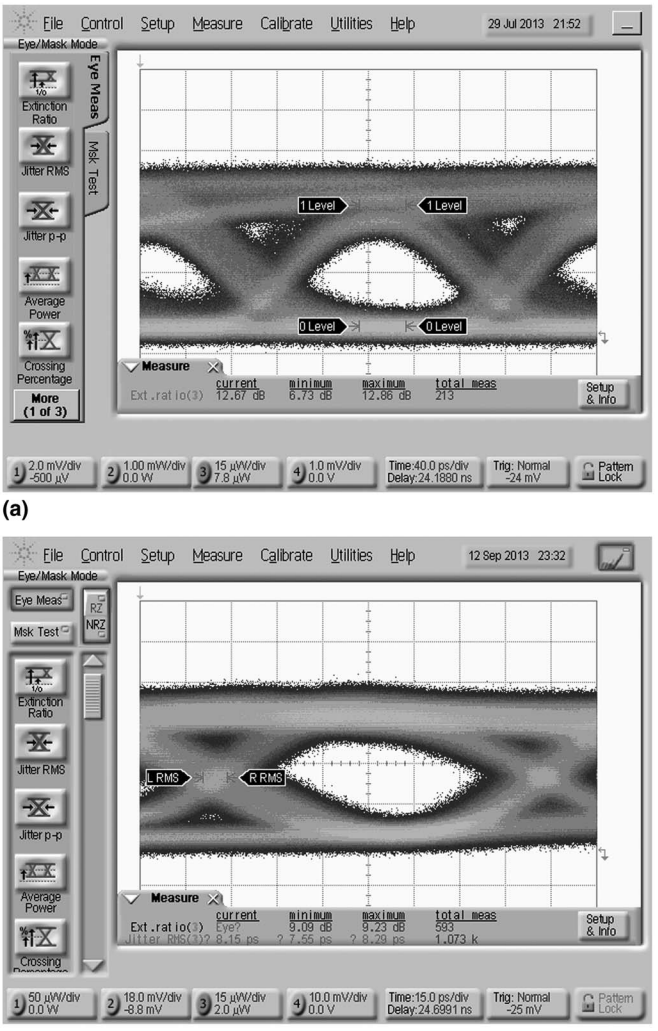

Figure 7. Measured ring modulator optical eye diagrams: (a) 5-Gb/s optical eye diagrams of first-generation ring modulators driven by the $4V_{pp}$ pre-emphasis driver; and (b) 9-Gb/s optical eye diagrams of second-generation ring modulators driven by the $2V_{pp}$ pre-emphasis driver.

After the modulated light is vertically coupled out into a multimode fiber, the light is observed with an optical oscilloscope. The $4V_{pp}$ CMOS driver modulates the first-generation ring modulator to improve the carrier dynamics in the p–i–n junction and compensate for the voltage overhead due to the large series contact resistance. Optimizing the pre-emphasis settings allows for an open eye with a 12.7-dB extinction ratio. Here, the maximum optical data rate is limited to 5 Gb/s because of the unanticipated ring resonator modulator's excess contact resistance of about 2 kΩ.

A residual oxide layer in the first-generation modulator created a high-resistance barrier between the metal contacts and the silicon devices. We removed this layer in the second generation, yielding a lower series contact resistance of about 200 Ω. For this device, the cathode is modulated by the energy-efficient $2V_{pp}$ CMOS driver, and the anode is biased at an adjustable dc level through a nonlinear voltage DAC for pre-emphasis optimization and bias-based ring resonance wavelength tuning. This driver configuration provides the potential for ring-tuning power and speed improvements in the context of proposed coarse thermal and fine bias-tuning systems [9].

Figure 7a and b shows the measured optical eye diagrams of the first- and second-generation prototypes, respectively. The second-generation prototype achieves an extinction ratio of 9.2 dB at a modulation speed of 9 Gb/s. Both generations of ring modulators exceed the approximately 7-dB extinction ratios achieved with the depletion-mode devices of Liu et al. [8], Buckwalter et al. [10], and Rakowski et al. [11]. The modulation efficiency of the two generations of prototypes are 808 fJ/b (5 Gb/s) and 480 fJ/b (9 Gb/s), respectively. Thanks to its dramatically reduced modulation swing (half of the first prototype) and low junction series resistance, the second-generation prototype improved the modulation energy efficiency by 40% and increased the modulation speed by 80% relative to the first-generation prototype. Overall, these energy-efficiency levels and the small optical device footprint (less than a 5-μm radius) provide strong motivation to leverage this photonic I/O architecture in a WDM system with multiple 10-Gb/s channels on a single waveguide.

Table 1 compares our design with recent carrier-injection and depletion ring transmitter designs that use hybrid integration via face-to-face microsolder bonding [8], [11] and monolithic integration [10], [12].

The second-generation prototype uses nonlinear pre-emphasis modulation to achieve 9-Gb/s operation at an energy efficiency better than or comparable to the designs operating between 2.5 and 10 Gb/s. Buckwalter et al. achieved 25-Gb/s modulation with monolithic integration of CMOS circuits and depletion rings in a 130-nm CMOS silicon-on-insulator (SOI) platform [10], but at the cost of both poor energy efficiency due to the relatively slow transistor speed and large area. The area of our transmitters is higher relative to the 40- and 45-nm CMOS designs, due to the 65-nm implementation and the use of a higher 8 : 1 serialization factor relative to the 2 : 1 designs [8], [12] and the nonserializing design [11].

Table 1 Performance summary and comparisons.

Parameter	This Work	Moss et al.[12]	Rakowski et al.[11]	Buckwalter et al.[10]	Liu et al.[8]
Technology (photonics)	130-nm SOI	45-nm CMOS SOI	130-nm SOI	130-nm CMOS SOI	130-nm SOI
Technology (CMOS)	65-nm CMOS	45-nm CMOS SOI	40-nm CMOS	130-nm CMOS SOI	40-nm CMOS
Integration type	Hybrid	Monolithic	Hybrid	Monolithic	Hybrid
Wavelength	1,310 nm	1,550 nm	1,550 nm	1,560 nm	1,550 nm
Ring radius	2.5 μm (1st gen), 5 μm (2nd gen)	20 μm	40 μm	7.5 μm	12 μm
Ring mode	Carrier injection	Carrier injection	Depletion	Depletion	Depletion
Ring Q factor	~8,000	4,290	~13,000	~13,000	~15,000
TX optical data rate	5 Gbps (1st gen), 9 Gbps (2nd gen)	2.5 Gbps	4 Gbps	25 Gbps	10 Gbps
ER	12.7 dB (1st gen), 9.2 dB (2nd gen)	3 dB	6.9 dB	6.9 dB	7 dB
Driver area	5,500 μm^2 (1st gen), 4600 μm^2 (2nd gen)	1,950 μm^2	238 μm^2	480,000 μm^2	120 μm^2
TX energy efficiency	0.81 pJ/bit (1st gen), 0.48 pJ/bit (2nd gen)	1.23 pJ/bit	0.34 pJ/bit	8.30 pJ/bit	0.4 pJ/bit

* ER: extinction ratio; Q: quality factor; SOI: silicon on insulator.

The optical modulation speed is limited up to 9 Gb/s, mainly for two reasons. First, the lack of driver output impedance control and the relatively long bond wires introduce some additional reflection-induced ISI, degrading the signal quality. Second, attenuation in the on-chip global clock distribution path limits the CMOS driver operation speed. Adopting advanced integration techniques such as flip-chip bonding or monolithic integration can alleviate the first issue. However, monolithically integrating the CMOS circuits with the photonic devices on a single die inevitably leads to compromises in the performance of each device, as well as increased fabrication costs. For example, it is inexpensive to implement both CMOS circuits and photonic devices in 130-nm CMOS SOI technology [10], which is sufficient for the photonic devices to confine the optical mode in the waveguides. However, this comes at the expense of the electrical circuits' energy and area efficiency. Using advanced CMOS SOI technology could alleviate these circuit issues, but would increase fabrication cost because of potential process changes and the relatively large footprint of photonic devices. These challenges have motivated a flip-chip bonding hybrid integration strategy for our next round of planned prototypes.

THE SILICON RING modulator-based photonic WDM transmitter presented in this paper incorporates high-swing nonlinear pre-emphasis drivers to overcome the limited bandwidth of carrier-injection ring resonator modulators. These prototypes provide the potential for silicon photonic links that can deliver distance-independent connectivity whose pin bandwidth scales with the degree of WDM. ■

■ References

[1] G. Li et al. "High-efficiency 25 Gb/s CMOS ring modulator with integrated thermal tuning," in *Proc. 8th IEEE Int. Conf. Group IV Photon.*, 2011, pp. 8–10.

[2] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson, "12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators," *Opt. Exp.*, vol. 15, no. 2, pp. 430–436, 2007.

[3] G. Wojcik et al., "A single Comb laser source for short reach WDM interconnects," *Proc. SPIE—Int. Soc. Opt. Eng.*, vol. 7230, DOI: 10.1117/12.816278.

[4] R. A. Soref and B. R. Bennett, "Electrooptical effects in silicon," *IEEE J. Quantum Electron.*, vol. QE-23, no. 1, pp. 123–129, Jan. 1987.

[5] C.-H. Chen et al., "Hybrid integrated DWDM silicon photonic transceiver with self-adaptive CMOS

circuits,'' in *Proc. IEEE Opt. Interconnects Conf.* 2013, pp. 122–123.

[6] C. Chen et al., "A WDM silicon photonic transmitter based on carrier-injection microring modulators,'' HP Laboratories, Tech. Rep. HPL-2014-11, 2014.

[7] C. Li et al., ''A ring-resonator-based silicon photonics transceiver with bias-based wavelength stabilization and adaptive-power-sensitivity receiver,'' in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2013, pp. 124–125.

[8] F. Y. Liu et al., ''10-Gbps, 5.3-mW optical transmitter and receiver circuits in 40-nm CMOS,'' *IEEE J. Solid-State Circuits*, vol. 47, no. 9, pp. 2049–2067, Sep. 2012.

[9] A. Titriku, C. Li, A. Shafik, and S. Palermo, ''Efficiency modeling of tuning techniques for silicon carrier injection ring resonators,'' in *Proc. IEEE Opt. Interconnects Conf.*, 2014, pp. 13–14.

[10] J. F. Buckwalter, Z. Zheng, G. Li, K. Raj, and A. V. Krishnamoorthy, ''A monolithic 25-Gb/s transceiver with photonic ring modulators and Ge detectors in a 130-nm CMOS SOI process,'' *IEEE J. Solid-State Circuits*, vol. 47, no. 6, pp. 1309–1322, Jun. 2012.

[11] M. Rakowski et al., ''Low-power, low-penalty, flip-chip integrated, 10 Gb/s ring-based 1 V CMOS photonics transmitter,'' in *Proc. IEEE Opt. Fiber Commun. Conf./Nat. Fiber Opt. Eng. Conf.*, 2013.

[12] B. R. Moss et al., ''A 1.23 pJ/b 2.5 Gb/s monolithically integrated optical carrier-injection ring modulator and all-digital driver circuit in commercial 45 nm SOI,'' in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2013, pp. 126–127.

Cheng Li is a Research Scientist in the Systems Research Laboratory, HP Laboratories, Palo Alto, CA, USA. His research interests include design of high-speed transceiver circuits for optical and photonic links, and photonic interconnects for network-on-chip systems. Li has a PhD in electrical engineering from Texas A&M University, College Station, TX, USA.

Chin-Hui Chen is a Research Scientist in the Systems Research Laboratory, HP Laboratories, Palo Alto, CA, USA. Her research interests include III–V semiconductors and silicon photonic ICs for applications in optical communication and interconnects. Chen has a PhD in electrical engineering from the University of California Santa Barbara, Santa Barbara, CA, USA. She is a member of the IEEE and the Optical Society of America (OSA).

Marco Fiorentino is a Research Scientist in the Systems Research Laboratory, HP Laboratories, Palo Alto, CA, USA. His research interests include quantum optics, high-precision measurements, and optical communications. Fiorentino has a PhD in physics from the University of Naples, Naples, Italy. He is a senior member of the Optical Society of America (OSA).

Raymond G. Beausoleil is an HP Fellow in the Systems Research Laboratory, HP Laboratories, Palo Alto, CA, USA, where he leads the Large-Scale Integrated Photonics Research Group. He is also a Consulting Professor of Applied Physics at Stanford University, Stanford, CA, USA. His research interests include microscale and nanoscale quantum optics for classical and quantum information processing, as well as quantum optics and models of laser interferometric gravitational wave interferometers. Beausoleil has a PhD in physics from Stanford University. He is a Fellow of the American Physical Society.

Binhao Wang is currently working toward a PhD in electrical engineering at Texas A&M University, College Station, TX, USA. His research interests include high-speed optical links and RF photonics. Wang has an MS in optical engineering from Zhejiang University, Hangzhou, China.

Samuel Palermo is an Assistant Professor in the Electrical and Computer Engineering Department, Texas A&M University, College Station, TX, USA. His research interests include high-speed electrical and optical links, high-performance clocking circuits, and integrated sensor systems. Palermo has a PhD in electrical engineering from Stanford University, Stanford, CA, USA. He is a member of Eta Kappa Nu and IEEE, and serves as an Associate Editor for the IEEE Transactions on Circuits and System—Part II: Express Briefs.

■ Direct questions and comments about this article to Cheng Li, HP Laboratories, Palo Alto, CA 94304 USA; cheng.li6@hp.com.