

# PHOTONICS Research

## Miniature, highly sensitive MOSCAP ring modulators in co-optimized electronic-photonic CMOS

HAYK GEVORGYAN,<sup>1,4</sup> ANATOL KHILO,<sup>2</sup> MARK T. WADE,<sup>2</sup> VLADIMIR M. STOJANOVIĆ,<sup>2,3</sup> AND MILOŠ A. POPOVIĆ<sup>1,5</sup> 

<sup>1</sup>Department of Electrical and Computer Engineering, Boston University, Boston, Massachusetts 02215, USA

<sup>2</sup>Ayar Labs Inc., Emeryville, California 94608, USA

<sup>3</sup>Department of Electrical Engineering and Computer Sciences, University of California Berkeley, Berkeley, California 94709, USA

<sup>4</sup>e-mail: hayk@bu.edu

<sup>5</sup>e-mail: mpopovic@bu.edu

Received 29 July 2021; revised 21 October 2021; accepted 13 November 2021; posted 15 November 2021 (Doc. ID 438047); published 22 December 2021

Convergence of high-performance silicon photonics and electronics, monolithically integrated in state-of-the-art CMOS platforms, is the holy grail for enabling the ultimate efficiencies, performance, and scaling of electronic-photonic systems-on-chip. It requires the emergence of platforms that combine state-of-the-art RF transistors with optimized silicon photonics, and a generation of photonic device technology with ultralow energies, increased operating spectrum, and the elimination of power-hungry thermal tuning. In this paper, in a co-optimized monolithic electronics-photonics platform (GlobalFoundries 45CLO), we turn the metal-oxide-semiconductor (MOS) field-effect transistor's basic structure into a novel, highly efficient MOS capacitor ring modulator. It has the smallest ring cavity (1.5  $\mu\text{m}$  radius), largest corresponding spur-free free spectral range (FSR = 8.5 THz), and record 30 GHz/V shift efficiency in the O-band among silicon modulators demonstrated to date. With 1  $V_{\text{pp}}$  RF drive, we show an open optical eye while electro-optically tuning the modulator to track over 400 pm (69 GHz) change in the laser wavelength (using 2.5  $V_{\text{DC}}$  range). A 90 GHz maximum electro-optic resonance shift is demonstrated with under 40 nW of power, providing a strong nonthermal tuning mechanism in a CMOS photonics platform. The modulator has a separately optimized body layer but shares the gate device layer and the gate oxide with 45 nm transistors, while meeting all CMOS manufacturability design rules. This type of convergent evolution of electronics and photonics may be the future of platforms for high-performance systems-on-chip. © 2021 Chinese Laser Press

<https://doi.org/10.1364/PRJ.438047>

### 1. INTRODUCTION

In the era beyond Moore's Law, the incorporation of photonics into advanced microelectronic chips and increasing convergence of photonics and electronics may be what sustains the advance of CMOS chip technology in the coming decades. Photonic interconnects show great promise to assist high-performance computer chips to continue to satisfy the growing demand for processing power in machine learning, big data analytics, the Internet of Things, and other emerging applications. Building optical transceivers through the monolithic integration of electronics and photonics achieves the ultimate low parasitics in electrical connections between photonic actives and CMOS circuits. It also leverages high-volume manufacturing through CMOS foundries and reliable packaging techniques [1,2]. Such transceivers have shown a steady

improvement in energy efficiency and bandwidth density as they move to more advanced technology nodes and utilize photonic components with a smaller footprint and lower energy consumption [1]. For instance, transceivers showed a sharp increase in energy efficiency and bandwidth density—key metrics for optical links—as microring resonator modulators (MRMs) [3] came to replace large and power-hungry Mach-Zehnder modulators (MZMs) [4]. This was made possible by overcoming the thermal stability issues of silicon MRMs using active thermo-optic feedback control [5,6] and implementing them in CMOS electronic-photonic platforms that supported such complex systems-on-chip. Such MRM-based links were used to demonstrate the first microprocessor with photonic I/O, comprising millions of transistors and hundreds of photonic devices [3], and more recently, multi-terabits per second

transceiver chiplets have been integrated in the same chip package with a high-performance commercial field-programmable gate array (FPGA) die [7].

The next jump in efficiency and bandwidth density for interconnects could be set off by another disruptive new device technology, or by introducing sufficiently novel or improved features in MRMs, which can enable transformative new functionalities or architectures.

Research on MRMs has been concentrated on improving their switching energy and speed, which currently exceeds 100 Gbps and often involves complex modulation schemes [8,9]. In serial optical links operating at data rates that exceed the system clock rate many times, the switching energy cost of the modulator constitutes only a small fraction of total link energy, which is dominated by the serializer, driver, thermal control, and laser. System-level analyses have shown that utilizing a large number of lower-speed parallel wavelength channels, each operating near the system clock rate, could be a more energy efficient solution for reaching aggregate bandwidths of hundreds of gigabits per second per fiber [10,11]. Such massively parallel architectures would require MRMs with very small radius and wide free spectral range (FSR), where a large number of wavelength channels can be accommodated. Typical MRMs are based on the rib waveguide geometry, which is subject to higher propagation losses in tight bends due to optical mode leakage into the slab. This limits the radius of such MRMs to 5–10  $\mu\text{m}$  and FSR to  $\sim 1.5$ –3 THz [8,9]. It is known that resonators with fully etched outer sidewall, such as microdisks, have orders of magnitude lower bending loss than rib-waveguide-based ring resonators of the same radius [12]. Conversely, modulators based on these types of resonators can be made much smaller in radius than the typical rib-waveguide modulators for an equivalent bend radiation loss ( $Q$ -factor). Microdisk modulators have been demonstrated in silicon photonics processes with radii as small as 2  $\mu\text{m}$  (6.92 THz FSR) [13] and 1.75  $\mu\text{m}$  (10 THz FSR, simulated) [14]. However, the excitation of higher-order modes in microdisk resonators produces spurious resonances that can introduce additional loss and cross talk [14].

To fully realize the potential of massively parallel links, the increased contribution of thermal tuning in the energy per bit, when the channel data rate is reduced and the number of channels is increased, must also be addressed [10,11]. It has been shown that the carrier plasma dispersion effect in highly efficient silicon MRMs [12,15] can provide a tuning range on the order of the dense wavelength division multiplexing (WDM) channel spacing, with near-zero static power consumption. Combined with barrel shifting of the ring resonances [10], this tuning range is sufficient for compensating fabrication errors and locking resonances to a WDM laser source grid. The concept of barrel shifting is useful when systematic fabrication errors misalign the modulator resonators from the laser wavelength grid by much more than one channel spacing. In this case, the laser channels can be spread across the FSR of the resonator, the resonators tuned only to the closest channel, and the next FSR used to wrap around the laser grid to address all channels. This greatly reduces the required maximum tuning range for each ring, and hence the energy cost of tuning.

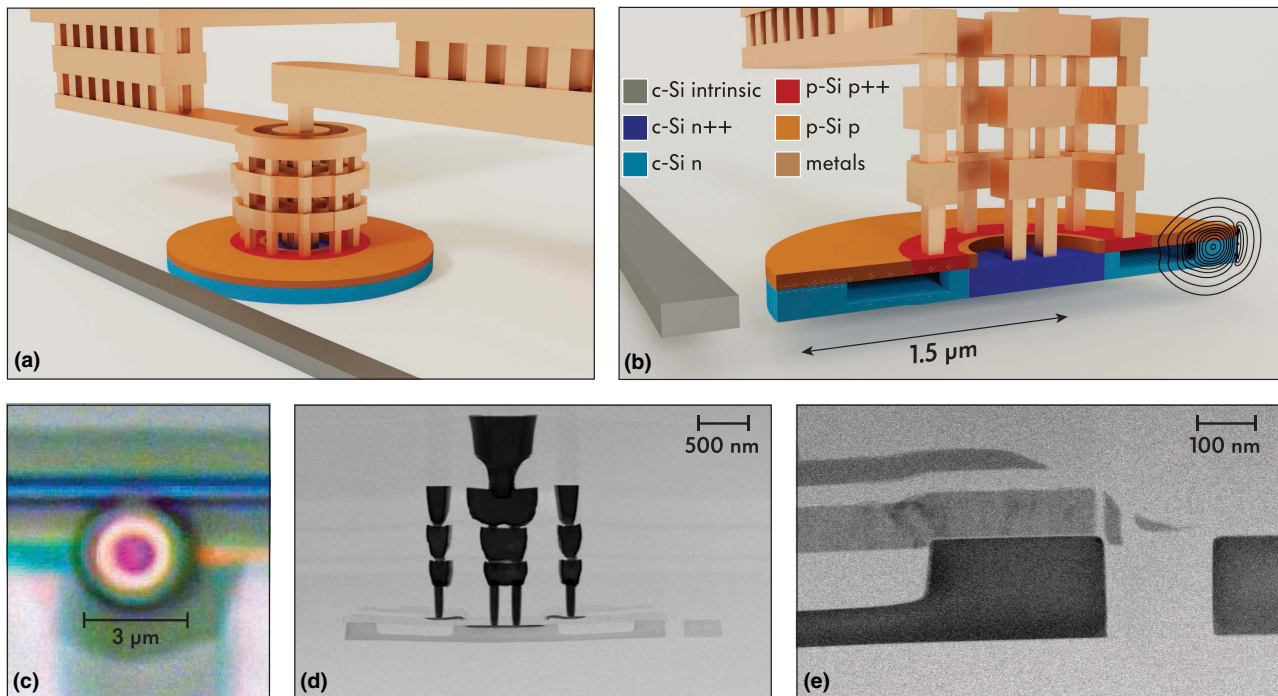
Bit reordering is then performed in the electrical back end. This results in a lower overall energy cost [10].

In this paper, we demonstrate the first MOS capacitor resonant modulator in a CMOS monolithic electronics-photonics platform. Based on a C-shaped cross section in a two-silicon device layer stack [Figs. 1(b), 1(d), and 1(e)], it has the smallest ring cavity (1.5  $\mu\text{m}$  radius), the largest corresponding spur-free FSR (FSR = 8.5 THz), and a record 30 GHz/V shift efficiency in the O-band, among silicon modulators demonstrated to date. It integrates a metal-oxide-semiconductor capacitor (MOSCAP) structure formed as the sandwich of the crystalline, photonics-optimized silicon body and poly-crystalline CMOS transistor gate device layers, separated by a thin gate oxide (also referred to as semiconductor-insulator-semiconductor capacitor or SISCAP [16]). The device is implemented in GlobalFoundries' next-generation 45 nm monolithic silicon-on-insulator (SOI) CMOS electronics-photonics process (45CLO) [2]. Its measured FSR of 8.5 THz or 48 nm spans nearly half of the telecommunications O-band, while showing no unwanted resonances of higher-order modes, owing to the radially inward partial-etch in the lower, body silicon device layer. The device also has the highest resonance shift efficiency in the O-band, 30 GHz/V, on par with our recent demonstration of vertical p-n junction designs of O-band spoked-ring modulators in "zero-change" CMOS [15] and with a C-band microdisk modulator [12]. Using its high voltage sensitivity, we demonstrate nonthermal tuning over a 90 GHz range with 4 V bias. MOSCAP MRMs have previously been demonstrated in silicon photonics (no electronics) and in III-V/Si hybrid platforms with radius 5  $\mu\text{m}$  and larger [17–20]. The devices that use silicon dioxide [17] or aluminum oxide [18–20] as an insulator show lower voltage sensitivity compared to the modulator in this work. The C-band MOS capacitor resonator with high- $k$  dielectric shows higher total resonance shift (1.1 nm or 140 GHz) with 4 V bias; however, only DC operation of this device is demonstrated [18]. In this first implementation, our MRM operates at 5 Gbps with 10 fJ/b and 10 Gbps with 5 fJ/b energy consumption. It can accommodate a large number of WDM channels within one FSR, which can enable wide parallel bus photonic interconnects.

## 2. MODULATOR DESIGN

A three-dimensional rendering of the MRM structure is shown in Figs. 1(a) and 1(b). Crystalline silicon (c-Si) body and poly-crystalline silicon (p-Si) gate device layers are separated by silicon dioxide and constitute the MOSCAP structure. The thickness of the body is 160 nm [2], while the gate is sub-100 nm thick. The total thickness of the modulator is slightly higher than the 220 nm thickness of the c-Si device layer of standard silicon photonics processes. The exact process geometry and parameters are available from the foundry under a nondisclosure agreement. The body is n-type doped and the gate p-type. This choice of doping was determined by the availability of suitable implant recipes in the process. In principle, the opposite doping polarity can be used.

The optical mode propagates along the 350 nm-wide waveguide core of the resonator, as shown in Fig. 1(b), where the thickness of the gate oxide is below 5 nm. The thin



**Fig. 1.** Perspective view of the MOSCAP modulator and microscope images of the fabricated device. (a) Full three-dimensional rendering of the modulator structure as designed; (b) device cross section, showing internal structure of the resonator, doping profiles of c-Si body and p-Si gate device layers, placement of metal contacts, and optical mode profile; (c) optical micrograph of the fabricated device; (d) TEM of the full device cross section, showing the resonator structure, bus waveguide, and lower four metal layers; (e) TEM zoom-in on the guiding core of the resonator and ring-bus coupling region.

oxide ensures high capacitance density. The gate oxide used in the MRM is the same as the gate oxide of one of the field-effect transistor (FET) types of this process. No modifications to the thickness of this standard oxide were made. Because the threshold voltage of FETs in a CMOS depends on the oxide thickness, the foundry controls it with great precision and ensures high repeatability. In addition, the oxide film is of high quality, minimizing leakage current. The vertical stack arrangement of the MOS layers also provides a strong overlap between the optical mode and the accumulated free carriers, similar to vertical p-n-junction modulators [12,15]. These two factors enable efficient resonance frequency modulation.

Radially inwards, the body is partially etched by 110 nm and filled with shallow trench isolation (STI) oxide [2], as shown in Figs. 1(b) and 1(d). The partial etch plays two important functional roles. First, the thick STI oxide significantly reduces parasitic capacitance by creating lower capacitance density in the region where the optical mode is weak and therefore has little contribution to the total resonance frequency modulation. Second, increased refractive index contrast confines the fundamental mode radially outward and prevents optical coupling into higher order modes from the bus waveguide. This is valuable because achieving low-loss phase-matched coupling [21] to resonators this small is practically impossible due to the high radiation loss in curved couplers. Instead, a simple, low-loss straight bus waveguide is used here. The bus width is 350 nm. For different doping designs, the optimum coupling gap (to


provide critical coupling between ring and bus) ranges between 125 and 180 nm.

Figure 1(c) is an optical micrograph of the MRM, showing the silicon device layer. Metal contacts are not visible in this image because it is taken from the back side of the chip after removing the silicon substrate. A transmission electron micrograph (TEM) of the full MRM cross section is shown in Fig. 1(d). The caving-in of the resonator structure in this image is likely caused by stress imposed by the metal contacts on the device layer in the thin lamella sample used for the analysis and is not representative of the actual device. The high-resolution lithography of the 45 nm CMOS node allows placement of vias and a few lower-level metal layers close to the center of the resonator with very tight radial confinement, which allows shrinking of the MRM radius to 1.5 μm without introducing excess losses due to metal absorption. Figure 1(e) shows a TEM zoom-in on the optical guiding core of the resonator with partially-etched c-Si body and p-Si gate layers, and the ring-bus coupling region. The bus waveguide has no p-Si cap to maintain low propagation loss.

### 3. EXPERIMENTAL RESULTS

The optical transmission spectrum of the MRM across the O-band, measured through a pair of grating couplers, is shown in Fig. 2(a). The resonances are spaced by a 8.5 THz (48 nm) FSR. Using a two-point-coupling scheme [22] the FSR could be further extended to cover the entire O-band. This FSR can





**Fig. 2.** Experimental data from passive optical and DC EO device characterization (device with  $32 \text{ k}\Omega/\square$  gate sheet resistance). (a) Optical transmission response of the MRM in O-band showing two resonances spaced by 8.5 THz or 48 nm FSR; (b) modulator optical response for gate voltages between  $-3$  and  $0 \text{ V}$ , showing resonance blueshift due to free carrier inversion; (c) optical response for gate voltages between  $0$  and  $4 \text{ V}$ , showing resonance blueshift due to carrier accumulation.

be used for WDM links with a large number of channels that could rely on a telecom-band optical comb [23] to provide continuous-wave light for each channel. The spectrum between the main resonances is free from spurious higher-order resonances seen in disk cavities owing to the single-mode operation of the MRM ensured by the partial etch along the inner radius in the lower, body silicon portion of the device cross section. The grating couplers used here show a large variation in transmission loss across the FSR. In practical links, they could be replaced by improved grating designs, or broadband edge couplers (available in the 45CLO process as standard library components [2]).

The modulator optical response is shown in Figs. 2(b) and 2(c) for different DC gate bias voltages. The resonance shows a blueshift with either an increasing positive or negative gate voltage. In the case of positive gate voltages, it happens due to accumulation of majority carriers on the p-Si-oxide and oxide-c-Si interfaces. In the case of negative gate voltages, it happens due to carrier inversion.

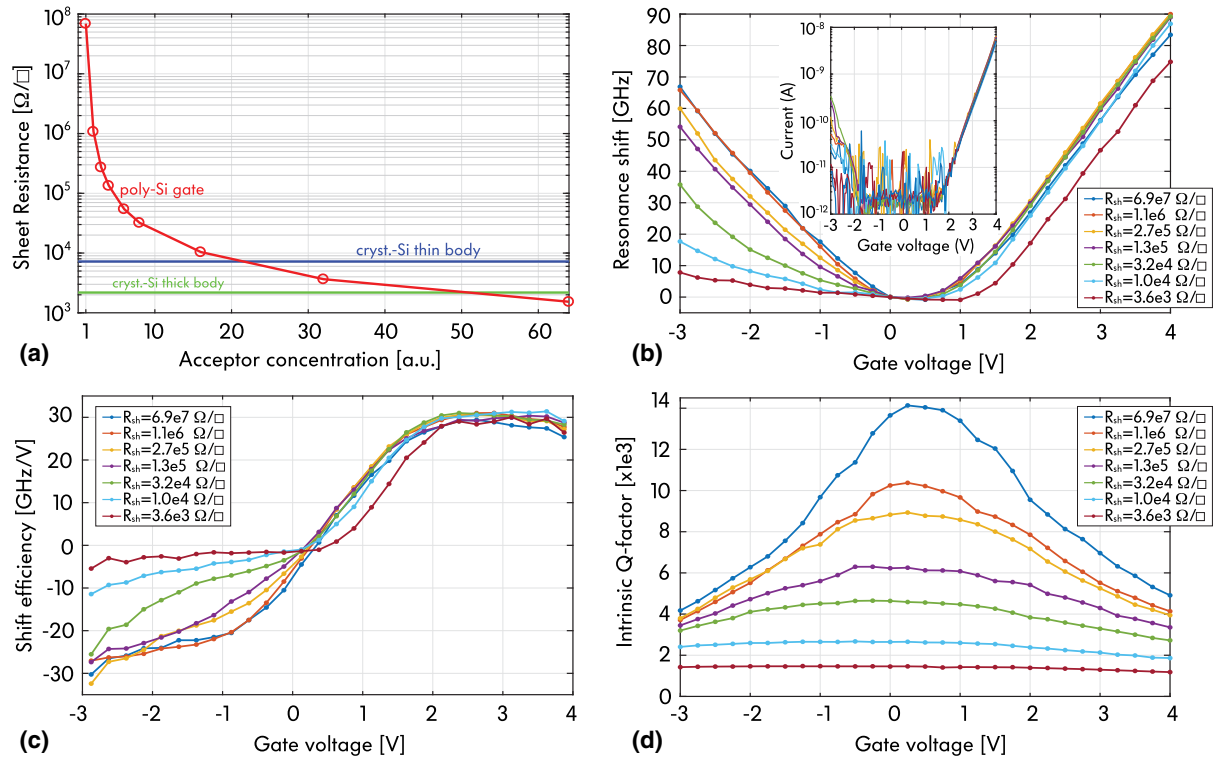
The MRM energy efficiency is proportional not only to the efficiency of shifting the resonance frequency (GHz/V), but also to the resonance quality factor, which drops with increasing free carrier concentration in the optical guiding region. The MRM's speed depends on the parasitic resistance of the conductive paths to the capacitor region. In polycrystalline silicon films, the free carrier concentration and mobility, and hence the resistance, show strong nonlinear dependence on doping concentration [24]. To study this dependence and explore the trade-off between the optical quality factor and resistance, MOSCAP modulators and supporting test structures for measuring resistance were fabricated with a wide range of acceptor concentrations in the p-Si. This was carried out by logarithmically stepping the dose of a foundry-provided doping recipe from  $1\times$  to  $64\times$  the base dose. The measured p-Si sheet resistance is shown in Fig. 3(a). When increasing the acceptor concentration 8 times from its base value, the sheet resistance drops 2150 times. However, an additional  $8\times$  increase in concentration (from  $8\times$  to  $64\times$ ) results in only a 21 times further drop in resistance. This behavior of p-Si resistance has been observed and well studied [24] and is associated with near-complete free carrier trapping in p-Si grain boundaries when impurity atoms

are fewer than the available trap states, with an abrupt jump in carrier concentration after all trap states are filled, explaining Fig. 3(a). At low acceptor concentrations [well below  $8\times$  in Fig. 3(a)], the sheet resistance is extremely high and unsuitable for Gbps operation of the modulator. With doping density increasing beyond  $20\times$  to  $50\times$  the base value, the p-Si sheet resistance drops to a few  $\text{k}\Omega/\square$ , below the resistance of, respectively, the thin and thick parts of the c-Si body part of the resonator [Fig. 3(a), blue and green lines], which then becomes the limiting factor on the resistance.

The DC electro-optic (EO) performance of these different doping variants of the modulator was measured. The  $I$ - $V$  characteristics are shown in the inset of Fig. 3(b). The thin gate oxide can sustain gate voltages between  $-3 \text{ V}$  (inversion) and  $+4 \text{ V}$  (accumulation), with leakage current not exceeding  $10 \text{ nA}$ . No damage to the gate oxide was observed in this voltage range. However, the oxide breaks down irreversibly at around  $4.5 \text{ V}$  as the current reaches a few tens of nanoamps. The optical resonance shift, relative to  $0 \text{ V}$  bias, is shown in Fig. 3(b) for all device variants. The resonances shift up to  $90 \text{ GHz}$  ( $524 \text{ pm}$ ) between  $0$  and  $4 \text{ V}$ . With less than  $10 \text{ nA}$ , this EO tuning is realized with under  $40 \text{ nW}$  of static power consumption.

The ring-to-bus waveguide coupling regime changes significantly with bias voltage, from over to critical to under-coupled, due to varying optical loss. Despite this, the ring-bus coupling gap can be (and was) designed to yield a high resonance extinction ratio (ER), sufficient for intensity modulation at all voltage bias points in this range. For instance, in Fig. 2(c) the ER is greater than  $12 \text{ dB}$  across all biases.

The slope of the shift curves in Fig. 3(b) represents the (low-frequency) resonance shift efficiency, shown in Fig. 3(c). It is proportional to the low-frequency MOS capacitance. The shift efficiency increases from  $0 \text{ GHz/V}$  near  $0 \text{ V}$  gate bias to  $30 \text{ GHz/V}$  in complete accumulation regime above  $2 \text{ V}$  bias. In inversion, the shift efficiency is also close to  $30 \text{ GHz/V}$  for devices with lightly doped gates. However, as the asymmetry of carrier concentrations in the body and gate increases with gate doping, the shift efficiency drops. In this case, when the body is entering into the inversion regime, the gate is still being



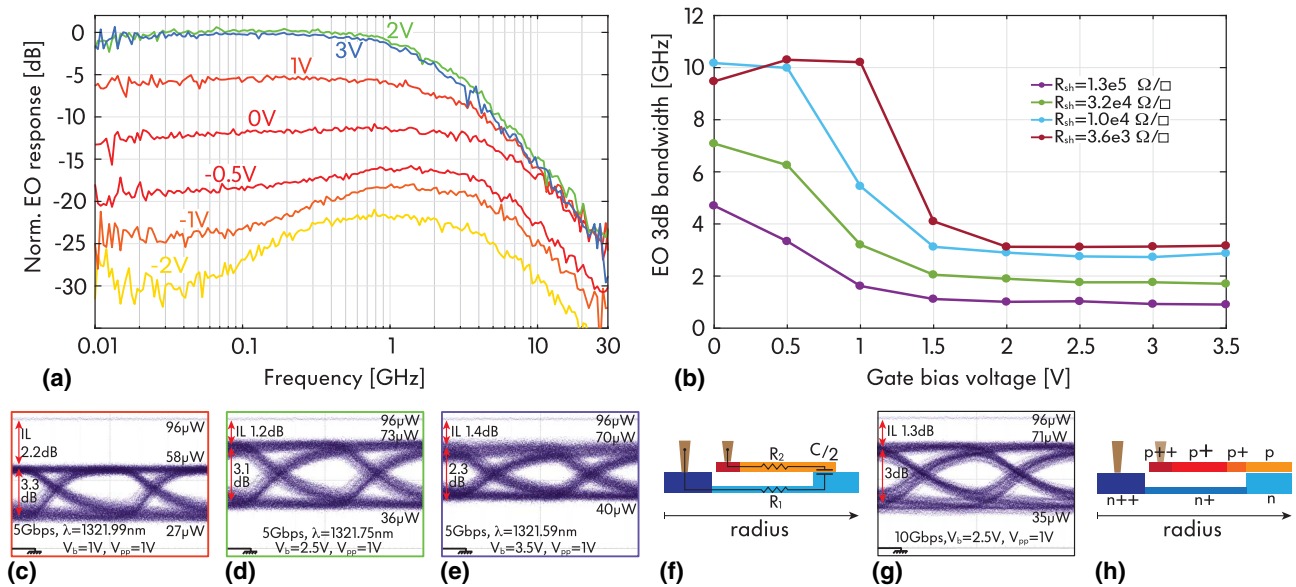
**Fig. 3.** Experimental data from DC electrical and EO device characterization for different p-Si gate doping concentrations. (a) Dependence of gate sheet resistance on acceptor concentration; (b) resonance frequency shift versus gate voltage, showing up to 90 GHz shift for gate voltages between 0 and 4 V; (c) modulator shift efficiency at different bias points, showing 30 GHz/V resonance modulation efficiency above 2 V; (d) resonance intrinsic quality factor versus gate voltage.

depleted of the majority carriers, which are two counteracting effects resulting in a reduced net change in free carrier concentration in the region overlapping the optical field. In other words, the depletion width in the p-Si effectively increases the insulator thickness, thus reducing the MOS capacitance.

It is worth noting that the presented (DC) shift efficiency in the inversion regime does not represent the high-frequency modulation efficiency because in this regime the capacitance of the MOS structure is reduced significantly compared to its low-frequency value. The optical resonance's intrinsic quality factor, extracted from measurement and shown in Fig. 3(d), varies from a maximum of 14,000 for a lightly doped device to a maximum of 1500 for a heavily doped device. This corresponds to linewidths between 15 and 150 GHz, which are good values that allow strong resonant enhancement while not limiting the speed optically. In each case, the  $Q$ -factor drops with increase in magnitude of the gate voltage due to extra loss from added free carriers. The external quality factors, extracted from experimental measurements, corresponding to 125 and 180 nm ring-to-bus coupling gaps are 4100 and 9700, respectively. Doping implants that are standard in the 45CLO process, with profiles designed for MZMs, but with modified doses, were used here. Custom-designed implants for this device, taking into account the nontrivial C-shaped cross section, could allow weaker dopings [higher sheet resistance values in Fig. 3(d)] to be used in the guiding part of the polysilicon to increase the

optical  $Q$  while preserving low resistance outside the optical region.

Moving to high-frequency dynamic tests, the small signal EO response of the modulators was measured at different gate bias voltages. Generally, the EO bandwidth depends on the detuning of the laser wavelength from the ring resonance. However, in situations such as this, when the RC-time-limited bandwidth is much smaller than the resonance linewidth, this dependence is small. Nevertheless, in these experiments, the laser detuning was adjusted to produce the highest optical modulation amplitude, an optimum operation point for data-com ring modulators. Figure 4(a) shows the normalized EO response of the modulator with a gate sheet resistance of  $32 \text{ k}\Omega/\square$ . With the transition of the device from inversion into depletion and then accumulation regimes, the modulation efficiency increases as a result of increasing high-frequency capacitance of the MOS structure. Simultaneously, the 3 dB EO bandwidth drops as a result of increasing RC time constant. Bandwidths of a few device variants are shown in Fig. 4(b) versus gate bias voltage, with the green line showing the bandwidth extracted from Fig. 4(a), which drops from 7 GHz at 0 V to approximately 2 GHz above 1.5 V. The EO response in Fig. 4(a) shows peaking at higher frequencies under negative gate bias voltages. Further investigation is needed to determine the cause of the peaking. Nonetheless, it does not affect the performance of the device, because as a data modulator it will



**Fig. 4.** Experimental results from high-speed EO characterization of modulators. (a) Normalized EO response of the modulator with gate sheet resistance of  $32 \text{ k}\Omega/\square$ ; (b) dependence of EO 3 dB bandwidth on gate bias voltage for a few modulator variants with different gate doping concentrations; (c)–(e) optical eye diagrams at 5 Gbps acquired at several bias points of gate voltage and laser wavelength; (f) illustration of the MOSCAP cross section with half the capacitance of the original device; (g) eye diagram (10 Gbps) of the MRM with 2 times smaller capacitance; (h) illustration of the MOSCAP cross section with additional implants introduced for reducing series resistance while preserving resonance quality factor.

be operated under positive gate bias voltages where no peaking of EO response is observed.

The same device is used for demonstrating on–off key (OOK) modulation at 5 Gbps with  $1V_{pp}$  drive at gate bias voltages between 1 and 3.5 V and corresponding laser wavelengths between 1321.99 and 1321.59 nm. Optical eye diagrams are shown in Figs. 4(c)–4(e). The eye diagrams show a gray line indicating (“through port”) transmission level far off resonance. The insertion loss and ER then indicate the 1 and 0 levels within the modulator resonance, as given by the sampling scope, relative to this off-resonant transmission level. As indicated in Figs. 4(c)–4(e), at different bias points, the modulation insertion loss ranges between 1.2 and 2.2 dB and the ER ranges between 2.3 and 3.3 dB. Using the gate-body MOS structure capacitance density [ $\text{fF}/\mu\text{m}^2$ ] provided in the 45CLO process design kit (PDK), we estimate the total capacitance of the modulator at 40 fF in accumulation, and its  $CV^2/4$  energy consumption at 10 fJ/b. Operation at 10 Gbps with 5 fJ/b energy is demonstrated using a device with half the capacitance, where the gate extends radially over half of the guiding optical region, as shown in Fig. 4(f). The eye diagram with 1.3 dB insertion loss and 3 dB ER is shown in Fig. 4(g).

The MRM speed is primarily limited by resistances of the p-Si gate and thin c-Si slab. A subject of current ongoing work is reduction of the series resistance by using a few additional implants to increase the doping concentration in the gate and slab layers in the region that is away (radially inward) from the optically guiding core of the resonator, as shown in Fig. 4(h). This is expected to significantly reduce resistance with negligible impact on the optical loss due to no overlap

with the optical guided mode and will enable higher RC bandwidths while to a large extent preserving the optical resonance quality factor. To further reduce the optical propagation loss and resistance in the gate, the p-Si deposition for photonic devices can be optimized and performed separately from transistor gate deposition to yield larger grain size and, therefore, higher carrier mobility and reduced optical scattering, as has been shown in previous work in CMOS platforms [24,25].

#### 4. CONCLUSIONS

With its large FSR, compact footprint, and very high shift efficiency, we believe this device and ones like it could be enablers of next-generation wide parallel optical I/O buses implemented in CMOS. This work also demonstrates the judicious coevolution, by design, of electronics and photonics in a state-of-the-art monolithic CMOS platform, to realize gains in performance and function in both domains while preserving the CMOS manufacturability process flow that supports the realization of complex systems comprising myriads of devices. This is done through maximal reuse of device layers, masks, and implants where possible, while introducing minimal new features as necessary. Adopting this constrained approach to device innovation applies more broadly, beyond modulators, and could enable advanced CMOS platforms like the 45CLO process in this paper to be a foundation for the gradual introduction of various novel device functions and technologies into state-of-the-art commercial foundry platforms, and their immediate incorporation and impact in larger CMOS systems-on-chip both in research and industrial application.

**Acknowledgment.** We would like to thank K. Nummy, A. Stricker, K. Giewont, T. Letavic, and the entire GlobalFoundries 45CLO process team for support on an early development run (Haystack); Deniz Onural for help with ring-bus coupling simulations, and Bohan Zhang.

**Disclosures.** HG: Ayar Labs, Inc. (F,E), AK: Ayar Labs, Inc. (E), MTW: Ayar Labs, Inc. (I,E), VMS: Ayar Labs, Inc. (I,E), MAP: Ayar Labs, Inc. (F,I,C).

**Data Availability.** Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

## REFERENCES

1. V. Stojanović, R. J. Ram, M. Popović, S. Lin, S. Moazeni, M. Wade, C. Sun, L. Alloatti, A. Atabaki, F. Pavanella, N. Mehta, and P. Bhargava, "Monolithic silicon-photonics platforms in state-of-the-art CMOS SOI processes," *Opt. Express* **26**, 13106–13121 (2018).
2. M. Rakowski, C. Meagher, K. Nummy, A. Aboketaf, J. Ayala, Y. Bian, B. Harris, K. McLean, K. McStay, A. Sahin, L. Medina, B. Peng, Z. Sowinski, A. Stricker, T. Houghton, C. Hedges, K. Giewont, A. Jacob, T. Letavic, D. Riggs, A. Yu, and J. Pellerin, "45 nm CMOS–silicon photonics monolithic technology (45CLO) for next-generation, low power and high speed optical interconnects," in *Optical Fiber Communication Conference (OSA, 2020)*, paper T3H.3.
3. C. Sun, M. T. Wade, Y. Lee, J. S. Orcutt, L. Alloatti, M. S. Georgas, A. S. Waterman, J. M. Shainline, R. R. Avizienis, S. Lin, B. R. Moss, R. Kumar, F. Pavanella, A. H. Atabaki, H. M. Cook, A. J. Ou, J. C. Leu, Y.-H. Chen, K. Asanović, R. J. Ram, M. A. Popović, and V. M. Stojanović, "Single-chip microprocessor that communicates directly using light," *Nature* **528**, 534–538 (2015).
4. N. B. Feilchenfeld, F. G. Anderson, T. Barwicz, S. Chilstedt, Y. Ding, J. Ellis-Monaghan, D. M. Gill, C. Hedges, J. Hofrichter, F. Horst, M. Khater, E. Kiewra, R. Leidy, Y. Martin, K. McLean, M. Nicewicz, J. S. Orcutt, B. Porth, J. Proesel, C. Reinholm, J. C. Rosenberg, W. D. Sacher, A. D. Stricker, C. Whiting, C. Xiong, A. Agrawal, F. Baker, C. W. Baks, B. Cucci, D. Dang, T. Doan, F. Doany, S. Engelmann, M. Gordon, E. Joseph, J. Maling, S. Shank, X. Tian, C. Willets, J. Ferrario, M. Meghelli, F. Libsch, B. Offrein, W. M. J. Green, and W. Haensch, "An integrated silicon photonics technology for O-band datacom," in *IEEE International Electron Devices Meeting (IEDM)* (IEEE, 2015), pp. 25.7.1–25.7.4.
5. K. Padmaraju, J. Chan, L. Chen, M. Lipson, and K. Bergman, "Thermal stabilization of a microring modulator using feedback control," *Opt. Express* **20**, 27999–28008 (2012).
6. C. Sun, M. Wade, M. Georgas, S. Lin, L. Alloatti, B. Moss, R. Kumar, A. H. Atabaki, F. Pavanella, J. M. Shainline, J. S. Orcutt, R. J. Ram, M. Popovic, and V. Stojanovic, "A 45 nm CMOS-SOI monolithic photonics platform with bit-statistics-based resonant microring thermal tuning," *IEEE J. Solid-State Circuits* **51**, 893–907 (2016).
7. M. Wade, E. Anderson, S. Ardan, P. Bhargava, S. Buchbinder, M. L. Davenport, J. Fini, H. Lu, C. Li, R. Meade, C. Ramamurthy, M. Rust, F. Sedgwick, V. Stojanovic, D. Van Orden, C. Zhang, C. Sun, S. Y. Shumarayev, C. O'Keefe, T. T. Hoang, D. Kehlet, R. V. Mahajan, M. T. Guzy, A. Chan, and T. Tran, "TeraPHY: a chiplet technology for low-power, high-bandwidth in-package optical I/O," *IEEE Micro* **40**, 63–71 (2020).
8. J. Sun, M. Sakib, J. Driscoll, R. Kumar, H. Jayatilaka, Y. Chetrit, and H. Rong, "A 128 Gb/s PAM4 silicon microring modulator," in *Optical Fiber Communications Conference and Exposition (OFC)* (OSA, 2018), paper Th4A.7.
9. Y. Tong, Z. Hu, X. Wu, S. Liu, L. Chang, A. Netherton, C.-K. Chan, J. E. Bowers, and H. K. Tsang, "An experimental demonstration of 160-Gbit/s PAM-4 using a silicon micro-ring modulator," *IEEE Photon. Technol. Lett.* **32**, 125–128 (2020).
10. M. Georgas, J. Leu, B. Moss, C. Sun, and V. Stojanovic, "Addressing link-level design tradeoffs for integrated photonic interconnects," in *IEEE Custom Integrated Circuits Conference (CICC)* (IEEE, 2011), pp. 1–8.
11. D. A. B. Miller, "Attojoule optoelectronics for low-energy information processing and communications," *J. Lightwave Technol.* **35**, 346–396 (2017).
12. E. Timurdogan, C. M. Sorace-Agaskar, J. Sun, E. S. Hosseini, A. Biberman, and M. R. Watts, "An ultralow power athermal silicon modulator," *Nat. Commun.* **5**, 4008 (2014).
13. A. Biberman, E. Timurdogan, W. A. Zortman, D. C. Trotter, and M. R. Watts, "Adiabatic microring modulators," *Opt. Express* **20**, 29223–29236 (2012).
14. M. Gehl, C. Long, D. Trotter, A. Starbuck, A. Pomerene, J. B. Wright, S. Melgaard, J. Sirola, A. L. Lentine, and C. DeRose, "Operation of high-speed silicon photonic micro-disk modulators at cryogenic temperatures," *Optica* **4**, 374–382 (2017).
15. H. Gevorgyan, D. V. Orden, D. Onural, A. Khilo, V. Stojanovic, M. Wade, and M. Popovic, "High shift efficiency O-band spoked-ring modulator allowing fully electro-optic channel tuning in a 45 nm CMOS platform," in *Conference on Lasers and Electro-Optics (OSA, 2021)*, paper SW3C.5.
16. M. Webster, C. Appel, P. Gothoskar, S. Sunder, B. Dama, and K. Shastri, "Silicon photonic modulator based on a MOS-capacitor and a CMOS driver," in *IEEE Compound Semiconductor Integrated Circuit Symposium (CSICS)* (IEEE, 2014), pp. 1–4.
17. J. V. Campenhout, M. Pantouvaki, P. Verheyen, S. Selvaraja, G. Lepage, H. Yu, W. Lee, J. Wouters, D. Goossens, M. Moelants, W. Bogaerts, and P. Absil, "Low-voltage, low-loss, multi-Gb/s silicon micro-ring modulator based on a MOS capacitor," in *Optics InfoBase Conference Papers (OSA, 2012)*, paper OM2E.4.
18. X. Huang, D. Liang, C. Zhang, G. Kurczveil, X. Li, J. Zhang, M. Fiorentino, and R. Beausoleil, "Heterogeneous MOS microring resonators," in *IEEE Photonics Conference (IPC)* (IEEE, 2017), pp. 121–122.
19. S. Srinivasan, D. Liang, and R. G. Beausoleil, "Heterogeneous SISCAP microring modulator for high-speed optical communication," in *European Conference on Optical Communications (ECOC)* (IEEE, 2020), pp. 1–3.
20. Q. Li, C. P. Ho, H. Tang, M. Okano, K. Ikeda, S. Takagi, and M. Takenaka, "Si racetrack optical modulator based on the III-V/Si hybrid MOS capacitor," *Opt. Express* **29**, 6824–6833 (2021).
21. E. S. Hosseini, S. Yegnanarayanan, A. H. Atabaki, M. Soltani, and A. Adibi, "Systematic design and fabrication of high-Q single-mode pulley-coupled planar silicon nitride microdisk resonators at visible wavelengths," *Opt. Express* **18**, 2127–2136 (2010).
22. H. Shoman, H. Jayatilaka, A. H. K. Park, A. Mistry, N. A. F. Jaeger, S. Shekhar, and L. Chrostowski, "Compact wavelength- and bandwidth-tunable microring modulator," *Opt. Express* **27**, 26661–26675 (2019).
23. A. L. Gaeta, M. Lipson, and T. J. Kippenberg, "Photonic-chip-based frequency combs," *Nat. Photonics* **13**, 158–169 (2019).
24. J. Y. W. Seto, "The electrical properties of polycrystalline silicon films," *J. Appl. Phys.* **46**, 5247–5254 (1975).
25. A. H. Atabaki, S. Moazeni, F. Pavanella, H. Gevorgyan, J. Notaros, L. Alloatti, M. T. Wade, C. Sun, S. A. Kruger, H. Meng, K. Al Qubaisi, I. Wang, B. Zhang, A. Khilo, C. V. Baiocco, M. A. Popović, V. M. Stojanović, and R. J. Ram, "Integrating photonics with silicon nano-electronics for the next generation of systems on a chip," *Nature* **556**, 349–354 (2018).