Subtractive Photonics in Bulk CMOS

Craig Ives, *Member, IEEE*, Debjit Sarkar, *Graduate Student Member, IEEE*, and Ali Hajimiri, *Fellow, IEEE*

***Abstract*— Subtractive photonics is presented as a method for implementing photonic waveguides into any bulk CMOS or electronics process. Metal and glass are patterned in the backend layers by the foundry to reveal suspended dielectric waveguides when the metal is etched away. This method requires a simple wet etch and provides waveguiding of light up to the visible regime using the broad transparency windows of silicon oxides. Mechanical, chemical, and photonic considerations are discussed, and photonic design is extensively detailed in the context of a 180-nm CMOS process. Example waveguides are constructed and measured, with losses as low as 4.1 dB/cm for a multimode waveguide at 1550 nm. In addition, waveguides are measured in the visible range, waveguide-photodiode couplers are detailed, and electronic–photonic systems are demonstrated to be unaffected by the etching.**

***Index Terms*— CMOS, optoelectronics, photonics, waveguides.**

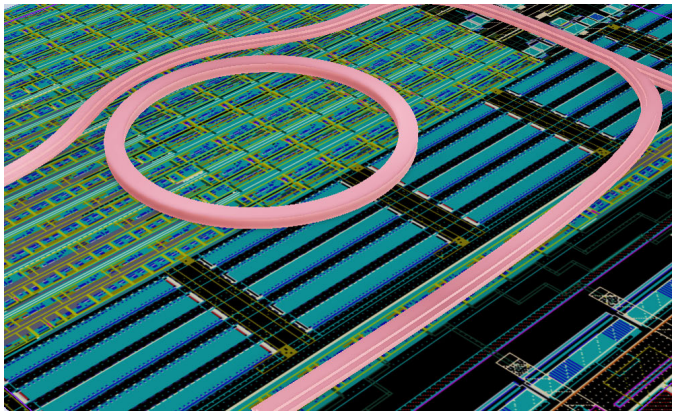

Fig. 1. Integration of photonic waveguides in a digital CMOS chip.

I. Introduction

THE dream of integrating photonics with silicon electronics (Fig. 1) has been pursued since the dawn of silicon photonics almost four decades ago [1]. A marriage of the maturity and ubiquity of CMOS with the communications, sensing, and rich nonlinear phenomena available in photonics may enable more tightly integrated systems than ever before. This may have a large impact on areas such as communications, sensing, and lidar, among many others.

Monolithic integration has been realized in foundry silicon on insulator (SOI) platforms [2] and through post-processing of foundry platforms, which has yielded the waveguide losses of 3.7 dB/cm in 45-nm SOI and 25 dB/cm in 32-nm SOI [3], but it has been difficult to realize in bulk CMOS. While SOI platforms offer a silicon dioxide layer for confining light in the thin crystalline silicon layer used for waveguides, bulk CMOS offers only a silicon substrate with glass and metal interconnect layers, at first glance a seemingly impossible platform to realize photonics in. However, there is an important reason to keep pursuing photonics in bulk CMOS. Bulk CMOS dominates electronics manufacturing with over 90% of the logic market [4], which has resulted in a large number of process options for electronics designers, from low-cost processes with fast turnaround and high tape-out availability, to state-of-the-art processes with the highest performance. Driven by this motivation, there have been a number of attempts to integrate photonics into bulk CMOS, which include front-end-of-the-line (FEOL) methods without process modification, FEOL methods with process modification, and back-end-of-the-line (BEOL) methods.

The basis of FEOL methods without process modification is the observation that many bulk CMOS processes use polysilicon for the transistor gates and that polysilicon can be isolated from the silicon substrate using shallow trench isolation (STI), a layer of silicon dioxide that is typically several hundred nanometers thick. Although this isolation is insufficient to prevent light from coupling into the substrate, it is sufficient to enable the substrate underneath the polysilicon to be etched without any effect on the polysilicon. The local etching of the substrate then provides sufficient optical isolation for the polysilicon waveguides [5]. This "zero-change" method, so-called because the foundry bulk CMOS process is not modified, initially resulted in waveguides with 55-dB/cm loss [4]. In later work, it was recognized that the high losses were due to the polysilicon material itself, which inspired the use of subwavelength grating waveguides that achieved a loss of 38 dB/cm [6].

It was recognized that the high losses of bulk CMOS waveguides could be mitigated using process modification. For example, changing the standard gate deposition from polysilicon to amorphous silicon, followed by an anneal step, can drop the zero-change waveguide losses to 6.2 dB/cm [7]. Alternatively, in the method of deep trench isolation (DTI), a module is added to the fabrication process for the deposition of thick oxide trenches in select portions of the bulk CMOS chip. In one variant of this method, polysilicon waveguides are then deposited on these thick trenches and can be doped using the same steps in the fabrication process as the transistors [8], [9]. This variant yields waveguides with lower loss than the zero-change method, with the best results in a recent demonstration of 21 dB/cm for waveguides fabricated with electronics in a 65-nm process [8]. Previous demonstrations in a 180-nm memory process have yielded waveguides with

Manuscript received 6 May 2023; revised 10 July 2023; accepted 13 August 2023. Date of publication 13 September 2023; date of current version 24 October 2023. This article was approved by Associate Editor Kaushik Sengupta. *(Corresponding author: Craig Ives.)*

The authors are with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: cives@caltech.edu).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/JSSC.2023.3308915.

Digital Object Identifier 10.1109/JSSC.2023.3308915

a loss of 10.5 dB/cm [9]. In the other variant of this method, amorphous silicon is deposited over the deep trench and then crystallized using solid phase epitaxy [10]. The crystallized silicon is then selectively etched to form waveguides with losses as low as 3 dB/cm in standalone operation, but when combined with electronics, the waveguide loss rises to 20 dB/cm [11].

Clearly, it is possible to produce low-loss waveguides in bulk CMOS using process modification. However, process modification necessarily restricts the number of processes a designer can choose from. Ideally, a designer should be able to design monolithic electronic–photonic integrated ciruits (EPICs) in any bulk CMOS process. The zero-change method comes close to this goal, since the process is not modified, but even the best waveguide loss of 38 dB/cm [6] is problematic for large photonic systems like phased arrays. Moreover, both FEOL methods share with SOI processes the problem of the size of photonic components. Due to the wavelength of light, photonic devices consume far more area than electronic devices, and in advanced electronic processes the photonics may simply become cost-prohibitive.

One way around the cost problem, in the context of monolithic integration, is to integrate photonic devices in the backend of the chip. Several variants of this method have been explored. In one method, silicon nitride waveguides, electrooptic polymer modulators, and polycrystalline germanium metal-semiconductor-metal photodetectors are deposited on glass on a silicon wafer. The waveguides attained a loss of 1 dB/cm in a demonstration platform [12] and show promise for integration on the backend of a bulk CMOS chip. Another possibility for this platform is to deposit amorphous silicon waveguides instead of silicon nitride, which can have losses as low as 5 dB/cm [13]. Since the amorphous silicon is ill-suited for modulation using the plasma dispersion effect [14], it must be used in conjunction with an external material. An alternative that does not require electrooptic polymer involves depositing amorphous silicon and then annealing it using excimer laser annealing, a common process in display fabrication [14]. The annealing results in polycrystalline silicon with large grain sizes, which could potentially provide low-loss waveguides, although an initial experimental waveguide had a loss of 65 dB/cm [14]. Modulation could be provided by the plasma dispersion effect in this platform, which would require doping during the backend fabrication process.

Although BEOL methods do not strictly require process modification, since the photonics components could conceivably be deposited after chip fabrication, the post-processing is extensive enough that it is similar in cost to process modification. This motivates the search for a method of BEOL photonic integration that is low-cost and can be applied to any electronics process. Subtractive photonics is one potential result of this search.

The method of subtractive photonics requires patterning the metal and glass in the BEOL of an electronics chip such that suspended dielectric waveguides are formed when the metal is etched away [15], [16], as shown in Fig. 2 and described in Appendix A. This method entails a number of advantages. First, the electronics process is not modified, so any electronics process with multiple interconnect layers of metal and glass can be used with this method. Second, the wet etching process is simple and low cost, lowering the barrier to entry. Third, the multilayer nature of the BEOL stackup naturally opens up a large design space for waveguide geometry,[1] which will be studied in the following. Fourth, as with all BEOL methods, subtractive photonics does not require the consumption of expensive FEOL area, maximizing cost-effectiveness. Fifth, as with some BEOL methods, the waveguide material is broadly transparent from the visible through the infrared. All of these advantages make subtractive photonics a compelling method of integration.

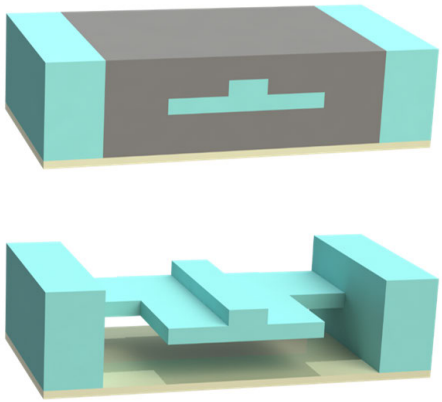

Fig. 2. Visualization of a suspended rib waveguide before and after etching.

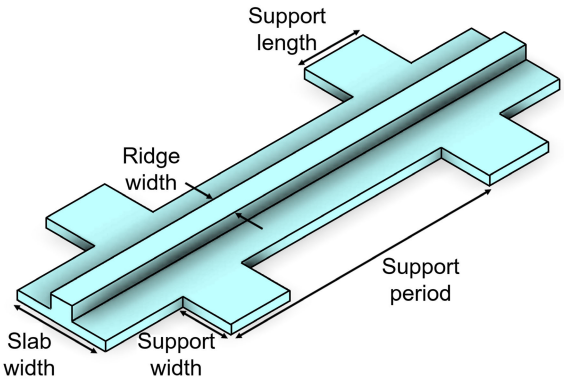

Fig. 3. Dimensions of significance in the rib waveguide.

II. Waveguides in 180-nm Bulk CMOS

Previous work in subtractive photonics [15] focused on suspended channel waveguides, in which the dielectric of the waveguide is surrounded on all sides by air, except for periodic supports. The waveguides in [15] indicate the need for a large number of supports per unit length, but significant reflections at each support will cumulatively limit the practicality of the suspended channel waveguide. To overcome this problem, the waveguide geometry is changed from channel to rib. The rib geometry introduces a rib protruding from a slab, with supports located at ends of the slab (Fig. 3). The propagating mode

[1]In contrast, other photonic platforms possess at most three waveguide layers, and there is little flexibility for waveguide design within the constraints of those layers.

TABLE I
COMPARISON OF BULK CMOS/DRAM INTEGRATION METHODS

Ref	Year	Method	Method Type	Process	Waveguide Material	Loss (dB/cm)	λ (nm)
This work	2023	Subtractive	Unmodified BEOL	180 nm bulk CMOS	Glass	4.1*	1550
[4]	2011	Zero-change	Unmodified FEOL	28 nm bulk CMOS	Polysilicon	55	1550
[6]	2020	Zero-change	Unmodified FEOL	90 nm bulk CMOS	SWG polysilicon	38	1550
[9]	2014	DTI	Modified FEOL	180 nm bulk CMOS	Polysilicon	10.5	1280
[8]	2018	DTI	Modified FEOL	65 nm bulk CMOS	Polysilicon	21**	1300
[11]	2014	DTI	Modified FEOL	65 nm bulk DRAM	Crystalline Si	20**	1550
[7]	2012	Zero-change	Modified FEOL	Emulated DRAM	Polysilicon	6.2	1550
[12]	2010	Deposition	BEOL	Photonics only	Silicon nitride	1	1310
[13]	2006	Deposition	BEOL	Photonics only	Amorphous Si	5/4	1300/1550
[14]	2013	Deposition	BEOL	Photonics only	ELA polysilicon	65	1550

*Best loss for the multimode waveguide
**Comparison is made only to waveguides fabricated with electronics
SWG: Subwavelength grating
DTI: Deep trench isolation
ELA: Excimer laser annealed
Photonics only: A BEOL demonstration that did not use SOI or electronics

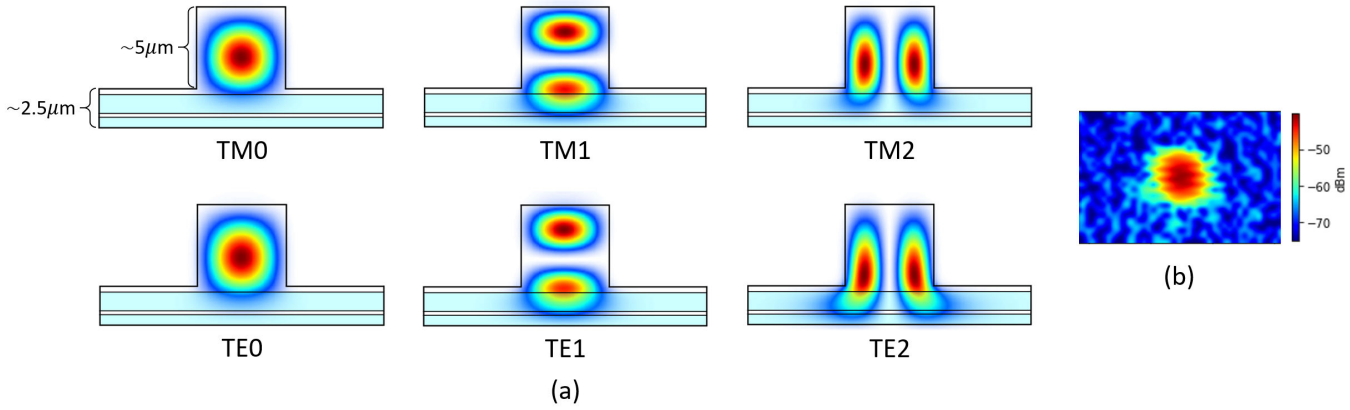

Fig. 4. (a) Simulated confined modes of the M6-V5M5V4 waveguide at 1550 nm. The light blue represents low-k dielectric layers. (b) Image of the waveguide output at 1550 nm obtained by scanning a lensed fiber across the waveguide facet.

is confined near the ridge and distanced from the support, thereby minimizing losses due to support reflections.

To demonstrate the method of subtractive photonics, we present simulations and measurements for two rib waveguide geometries that we have fabricated in a 180-nm bulk CMOS process optimized for RF devices.

A. M6-V5M5V4 Geometry

This waveguide geometry uses the dielectric associated with the ultrathick top metal of the RF process for the ridge of the waveguide and standard dielectric layers for the slab (Figs. 4 and 5). For the process we used, the ridge dielectric is the chip passivation, with the ridge patterning occurring on the metal 6 layer, and the slab patterning occurring on the via 5, metal 5, and via 4 layers (abbreviated as an M6-V5M5V4 geometry). Although the ultrathick top layer can reasonably only be used with a multimode waveguide, its thickness, and the resulting ridge size, provides two immediate benefits: 1) it guarantees that the fundamental mode will experience low loss relative to other waveguide geometries in this platform and 2) it enables low-loss fiber coupling to either the core of a visible wavelength fiber or the spot of a lensed infrared wavelength fiber. Both of these benefits make this geometry well suited for initial investigations. Moreover, the ease of fiber coupling provides an advantage over other photonic processes since the mode mismatch between a fiber and a nanoscale waveguide often causes significant coupling losses. Fig. 6 corroborates this with simulation results.

The M6-V5M5V4 geometry supports six confined modes at a wavelength of 1550 nm and a ridge width of 5 μm, as shown in Fig. 7, which plots the normalized index contrast (see Eq. 1) as ridge width varies. These mode profiles may be compared to an image of the waveguide facet (Fig. 4), obtained by scanning a fiber across the waveguide facet. The loss of the fundamental mode was simulated to be 3.48 dB/cm with 50-nm roughness and 0.14 dB/cm with 10-nm roughness (see Fig. 8).

The spiral fabricated for estimating the loss of this waveguide geometry is about 4.3 cm long and shown illuminated with 635-nm light in Fig. 9. We measured the upper bound (see Appendix B) of the waveguide loss across five samples and at the wavelengths of 635, 780, and 1550 nm. These results are plotted in Fig. 10 alongside the only two previous papers that have reported waveguide loss measurements in unmodified bulk CMOS [4], [6]. Our best result of 4.1 dB/cm at 1550 nm is to be compared against the losses of 55 dB/cm in [4] and 38 dB/cm in [6]. The measurements of waveguide loss at visible wavelengths in unmodified bulk CMOS have, to the best

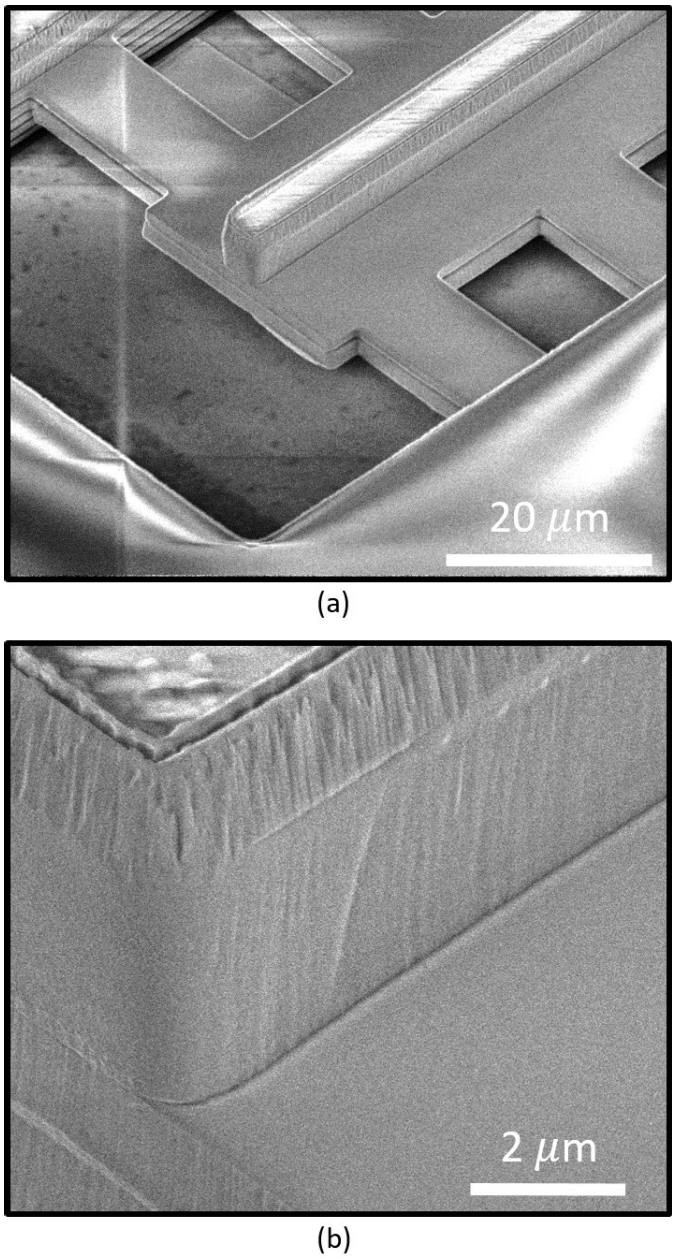

Fig. 5. (a) SEM image of an M6-V5M5V4 waveguide. (b) Close-up view of the ridge of the same waveguide.

of our knowledge, no precedent. The variations in insertion loss at 780 nm may be caused by dust on the waveguides or unetched regions of barrier layer, which may cause greater loss at visible wavelengths than infrared wavelengths since different modes are excited at each wavelength, which may interact with the dust or regions of barrier layer differently.

B. V5M5-V4M4V3 Geometry

Another geometry uses the via 5 and metal 5 layers for a ridge, and the via 4, metal 4, and via 3 layers for the slab. The choice of a two-layer ridge and a three-layer slab means that the ridge-slab index contrast is fairly low, even for rib waveguides. Consequently, the waveguide is single mode over a broad range of wavelengths for ridges as wide

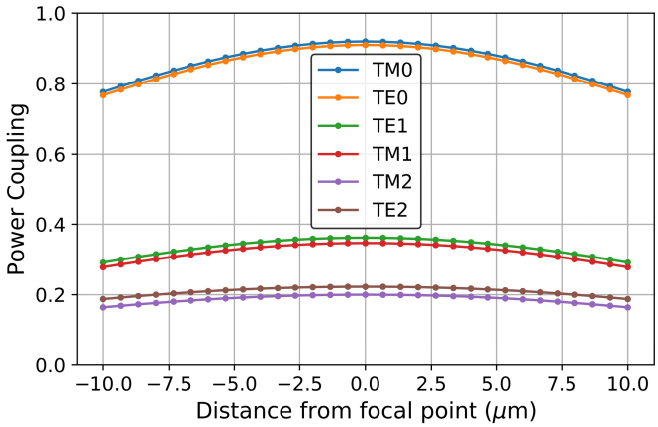

Fig. 6. Simulated fiber coupling efficiency for all confined modes of an M6-V5M5V4 waveguide and a lensed fiber with a 5-μm-diameter spot at 1550 nm.

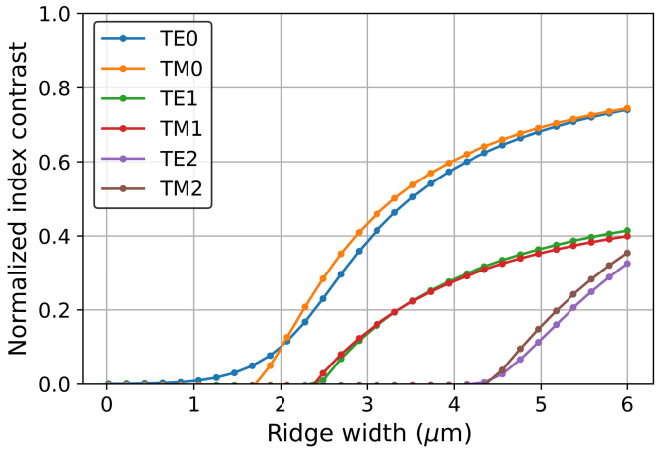

Fig. 7. Normalized effective index of the confined modes of the M6-V5M5V4 waveguide across ridge widths at 1550 nm.

as 4 μm, as shown in Fig. 11. Single-mode waveguides are useful for phase-sensitive applications such as coherent communications or interferometry since it can be difficult to excite and maintain only one mode in a multimode waveguide. The downside of single-mode rib waveguides is the large bend radius, a consequence of the low index contrast, so a multimode waveguide might be more useful for applications that require compact layouts but are not phase-sensitive.

The TE and TM confined modes at a wavelength of 780 nm and a ridge width of 4 μm are shown in Fig. 12. The simulated fiber coupling efficiency is plotted in Fig. 13 for a lensed 780-nm fiber with a spot size of 2 μm. For the same geometry and wavelength, Fig. 14 shows the slab width optimization using an eigenmode expansion (EME) simulator (see Section IV-B3).

The loss of the V5M5-V4M4V3 geometry with a 3-μm-wide ridge was simulated to be 4.35 dB/cm with 50-nm roughness and 0.15 dB/cm with 10-nm roughness (see Fig. 8, which depicts a waveguide with comparable roughness). The upper bound (see Appendix B) for this geometry with a 3-μm-wide ridge was measured to be 20.7 dB/cm at a wavelength of 1550 nm and using a spiral about 3.5 cm long. Although this waveguide loss compares favorably to many other methods of integrating photonics with electronics

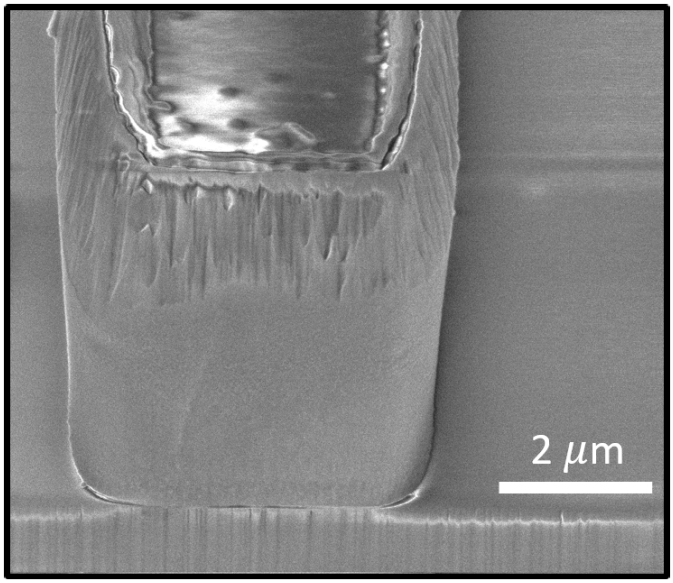

(a)

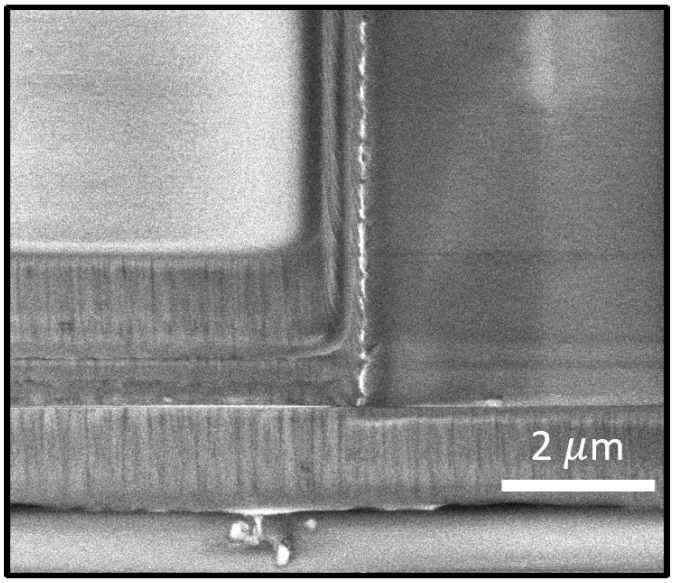

(b)

Fig. 8. SEM images for estimating sidewall roughness. Note that the horizontal surfaces are very smooth so that only the vertical sidewalls contribute to scattering loss. (a) M6-V5M5V4 waveguide. (b) V5M5V4-M4V3 waveguide.

(see Table I), the insertion loss of the spiral is significantly greater than that of the M6-V5M5V4 waveguide, which can be reasonably attributed to the large bend losses of the V5M5-V4M4V3 geometry. At visible wavelengths, a lack of suitable amplification prevented loss measurements due to the large insertion loss, again most likely due to the bend loss of this geometry.

III. Photodiode Couplers

CMOS processes offer a convenient method of interfacing between optics and electronics in the form of CMOS photodiodes, with the understanding that such silicon photodiodes can only detect light below a wavelength of about 1.1 μm. Although CMOS photodiodes are well studied and can be fabricated in standard bulk CMOS processes with zero design rule check (DRC) violations [17], research using these devices has so far been encumbered in many applications by a lack of waveguides in bulk CMOS, a barrier that subtractive photonics can potentially address.

In the 180-nm bulk CMOS process used in this work, a deep n-well (DNW) layer is available, which provides the option of an additional photodiode junction. Fig. 15 shows schematically that there are junctions available at the interface between the p-substrate (PSUB) and the DNW, the DNW and the p-well (PW), and the PW and a highly n-doped region (N+). It is of course also possible to forgo the DNW and use the junctions between the PSUB and n-well (NW), and the NW and a highly p-doped region (P+), but the analysis of [17] shows that the N+/PW junction provides the highest bandwidth, and the additional junction provided by DNW affords flexibility, so in this work, we predominantly use the PSUB/DNW/PW/N+ stackup for testing.

The main question to be answered in the context of subtractive photonics is how the light in the waveguide should be coupled into the substrate where the CMOS photodiode is located. One option is to simply terminate the waveguide in a reflector, although only a crude approximation to a reflector can be formed using the discrete thicknesses of the dielectric stackup. Moreover, the low index contrast between the glass and air will reduce the effectiveness of the reflector, and the high index contrast between the glass and silicon substrate will cause further losses. An alternative may be to couple adiabatically between the waveguide and the substrate with a long taper. Although this technique may encounter lower reflection losses, the larger required area increases photodiode capacitance.

A. Tapering Coupler

Fig. 16 shows the structure of our experimental waveguide taper. The V5M5-V4M4V3 waveguide has a 4-μm-wide ridge and a 25-μm-wide slab, both of which are tapered down to zero width over a length of 100 μm. The CMOS photodiode underneath the coupling region uses the full PSUB/DNW/PW/N+ stackup for maximum flexibility. The metal 1 layer is used to contact all four regions of the photodiode, and the metal 2 layer is used to connect each region to exterior pads or devices. This metal interconnect must be protected from the etching process used to form the waveguides, so the via 2 layer is reserved as a fully glass insulating layer. This leaves only the metal 3 layer to form an air gap between the waveguide and the photodiode insulation. Although the waveguide taper alone will cause light to couple into the photodiode due to the expanding mode, filling in the air gap on the metal 3 layer with a glass taper guarantees even stronger coupling to the substrate since the light loses all confinement from the following.

We designed several photodiodes with tapering waveguide couplers for demonstration purposes. One photodiode has an active area of 16 $\times$ 150 μm, while another has an active area of 50 $\times$ 150 μm. The measured responsivities of the N+/PW

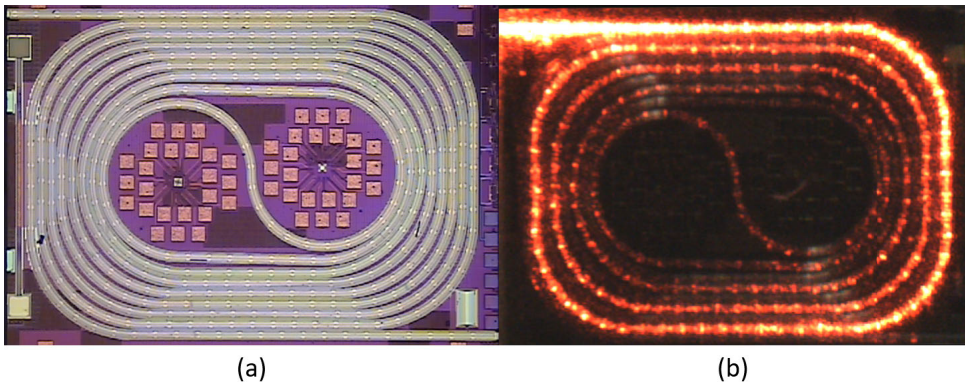

Fig. 9. (a) Die photograph of the M6-V5M5V4 spiral. (b) Spiral illuminated with 635-nm light.

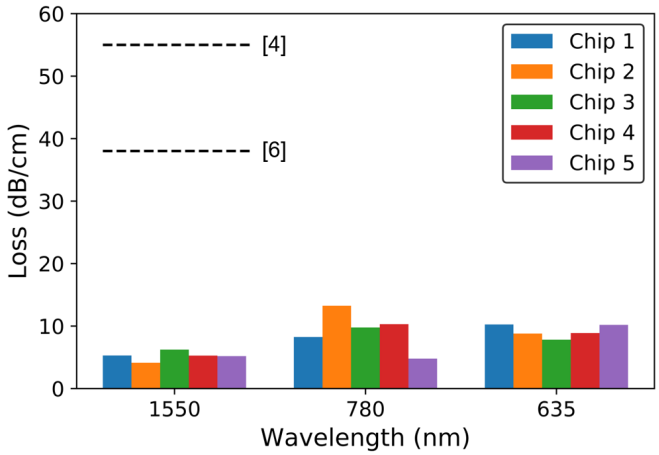

Wavelength (nm)	1550	780	635
Average loss (dB/cm)	5.2	9.3	9.2
Standard deviation (dB/cm)	0.7	2.8	0.9

Fig. 10. Measurements of the upper bound on waveguide loss for the M6-V5M5V4 spiral. The results are compared to the two previous publications reporting waveguide loss in unmodified bulk CMOS.

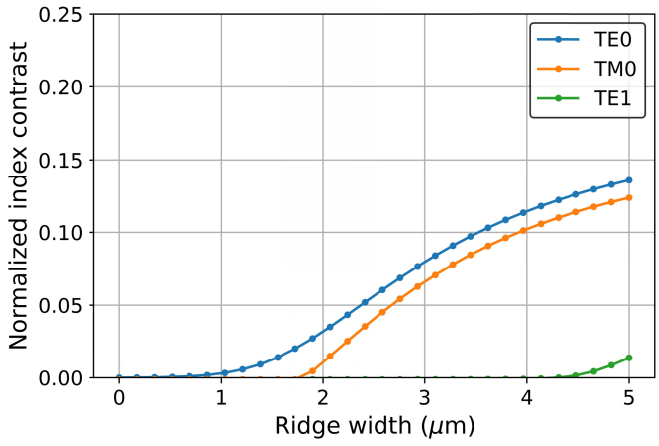

Fig. 11. Normalized effective index of the confined modes of the V5M5-V4M4V3 waveguide across ridge widths at 780 nm.

junctions are plotted in Figs. 17 and 18 against reverse bias, at the wavelengths of 635 and 780 nm, as well as in a comparison of fiber coupling and waveguide coupling at 780 nm.

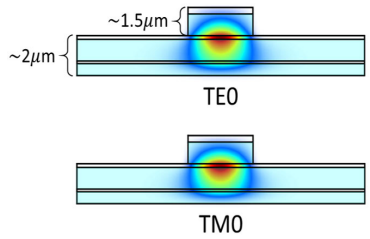

Fig. 12. Simulated confined modes of a V5M5-V4M4V3 waveguide with a 4-μm-wide ridge at 780 nm. The light blue represents low-k dielectric layers.

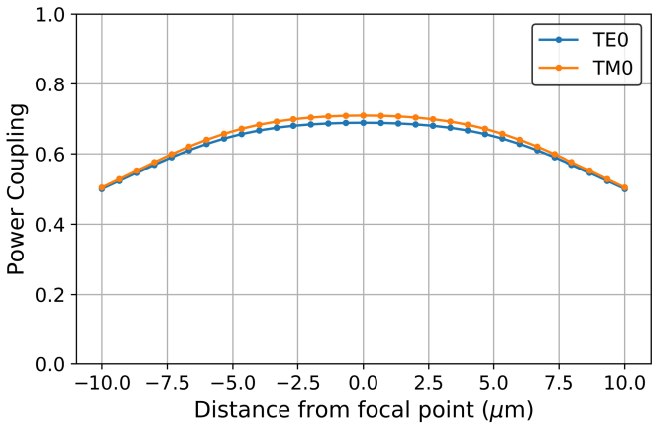

Fig. 13. Simulated fiber coupling efficiency for all confined modes of a V5M5-V4M4V3 waveguide and a lensed fiber with a 2-μm-diameter spot at 780 nm.

The finite-difference time-domain (FDTD) simulated coupling efficiency for the 16 $\times$ 150 μm photodiode is 30% at 780 nm, while the simulated coupling efficiency for the 50 $\times$ 150 μm photodiode is 34% at 780 nm. These responsivities can be improved by switching to the DNW/PSUB junction, which has the largest depletion region, at the cost of bandwidth due to the greatly increased diffusion current. Another method of increasing responsivity is to increase the photodiode area, which will reduce the bandwidth due to the increased capacitance. In order to show that electronics on chip is unaffected by post-processing, the smaller photodiode was also connected to a simple transimpedance amplifier (TIA) that exhibited a bandwidth of about 200 MHz with an estimated photocurrent

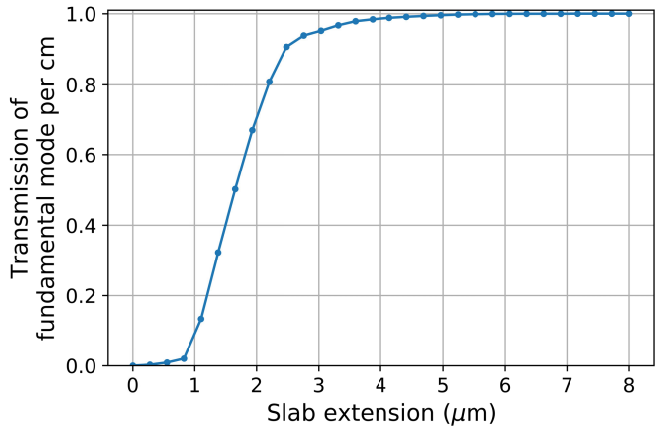

Fig. 14. Simulated support coupling efficiency at 780 nm for a V5M5-V4M4V3 waveguide with a 4-μm-wide ridge.

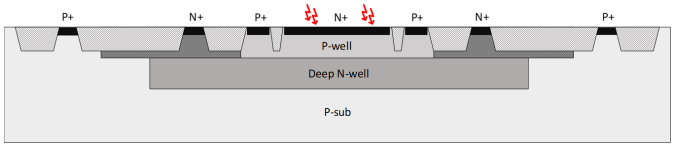

Fig. 15. CMOS photodiode with a PSUB/DNW/PW/N+ stackup.

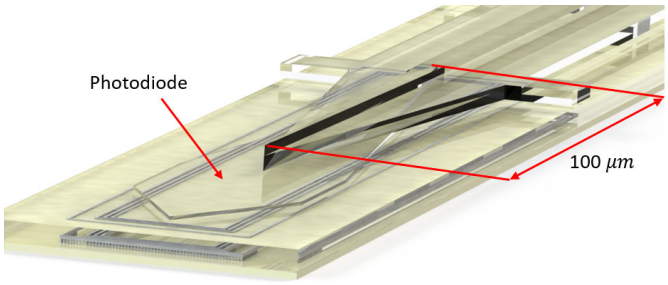

Fig. 16. 3-D model of the tapered waveguide-photodiode coupler. The ridge and slab taper to a point over a distance of 100 μm. The photodiode is visible inside the insulating glass.

of about 1 μA and $V_{DD} = 1.8$ V, measured with a 50-Ω probe and shown in Figs. 19 and 20 (the measurement setup is shown in Fig. 21). The bandwidth of the TIA can be improved through many means, for example, by adding an equalizer or negative capacitance compensation [18]. We also designed a segmented photodiode for the purpose of understanding the true area requirements of a waveguide taper. It can be seen in Fig. 22 that the responsivity is highest at segments 5 and 6, indicating that the photodiode area can be optimized for low capacitance even for long waveguide tapers. The high response at segment 1 is most likely due to light that is not confined in the waveguide due to imperfect fiber coupling.

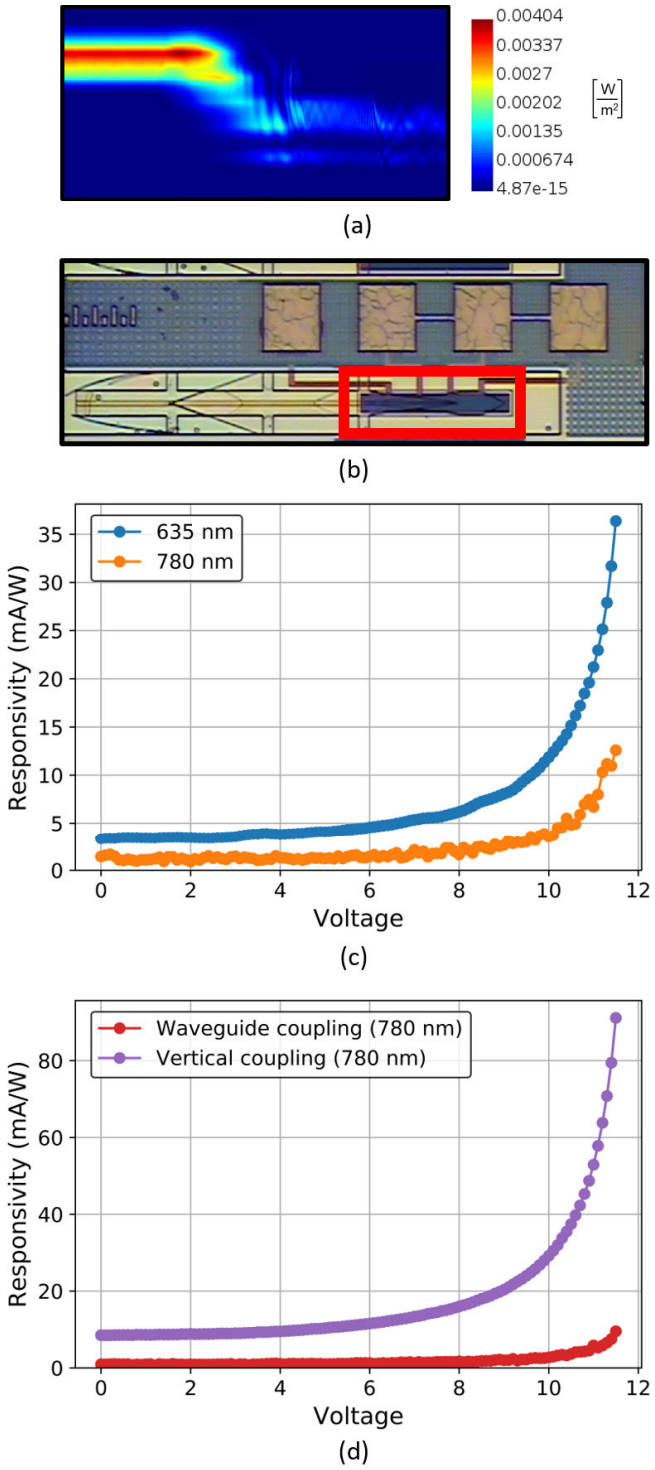

Fig. 17. Results of a 16 × 150 μm photodiode. (a) FDTD simulation using the full dielectric stackup. (b) Die photograph. (c) Measured responsivity of the waveguide-coupled photodiode at wavelengths of 635 and 780 nm. (d) Measured responsivity using waveguide coupling and vertical fiber coupling at 780 nm.

B. Reflecting Coupler

Despite the shortcomings of the reflecting coupler mentioned above, it minimizes chip area, which is useful both for the design of large systems and for minimizing photodiode capacitance, so this option should be optimized rather than abandoned. We began the design of a reflector by arranging the layers of the waveguide as an approximation of an angled reflector with a perfectly flat surface. We then applied particle swarm optimization in a commercial FDTD tool to obtain the optimized endface locations for each layer. With the addition of another reflecting layer, we simulated a coupling efficiency

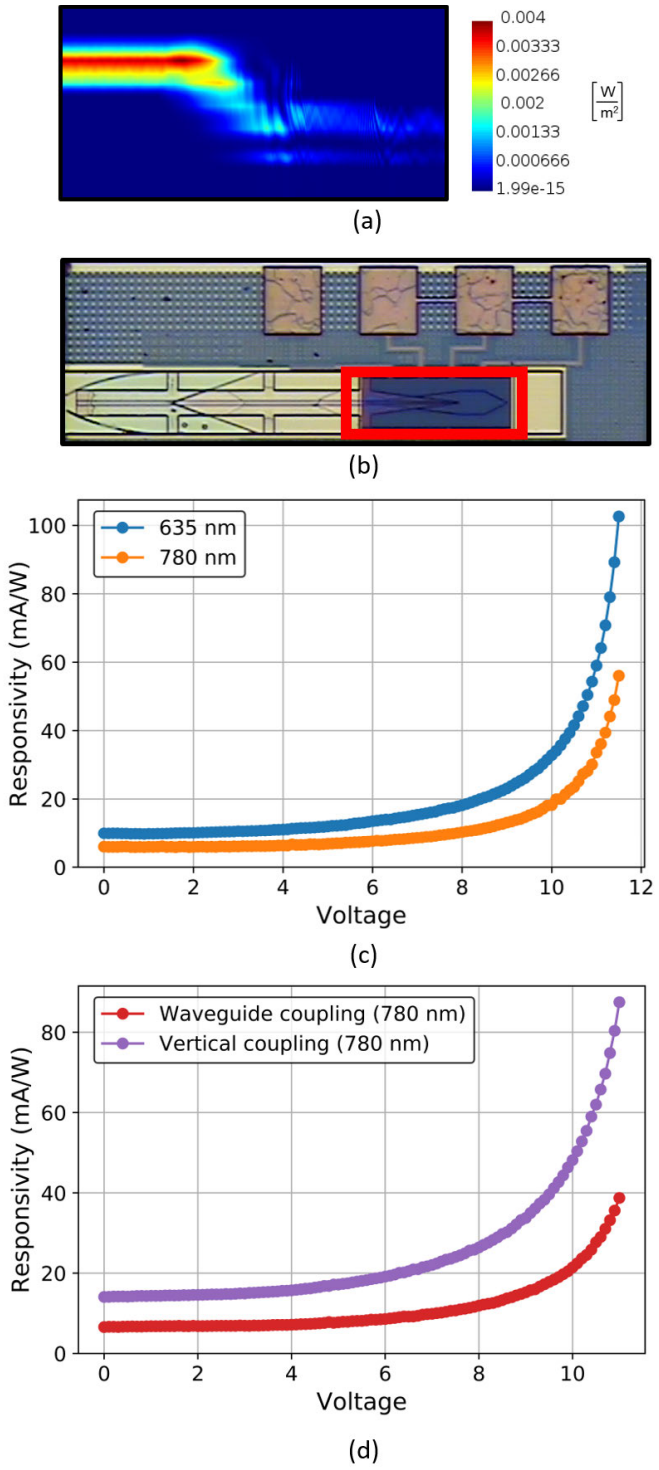

Fig. 18. Results of a 50 × 150 μm photodiode. (a) FDTD simulation using the full dielectric stackup. (b) Die photograph. (c) Measured responsivity of the waveguide-coupled photodiode at wavelengths of 635 and 780 nm. (d) Measured responsivity using waveguide coupling and vertical fiber coupling at 780 nm.

of 54% into the silicon substrate (Fig. 23). This is superior to what can be obtained with a standard grating coupler due to the low refractive index contrast of the glass–air system.

The insertion loss of the reflecting coupler was measured relative to a reference photodiode. Note that the insertion loss here includes the loss of the attached waveguide, the loss of the waveguide bend with a 100 μm radius, and the loss due to fiber coupling. At 635 nm, the insertion loss was measured to be 9.6 dB, while at 780 nm, the insertion loss was measured to be 10.7 dB. The simulated fiber coupling loss is about 1.5 dB at 780 nm (Fig. 13), which puts an upper bound of 9.2 dB on the loss of the reflector. Considering that the bend loss of the V5M5-V4M4V3 waveguide is high due to the poor ridge-slab index contrast, it is likely that the loss of the reflector is relatively close to the simulated loss of 2.7 dB, although further measurements are necessary to provide conclusive proof.

IV. Design Considerations

One challenge in subtractive photonics is that the index contrast between silicon dioxide-based BEOL dielectrics and the surrounding air is small compared to silicon photonics, which will necessarily increase the size of photonic structures, both in cross section and layout area. Another challenge is that the suspended photonic structures face mechanical limits, which constrains the design space. These constraints are studied in the following.

A. Discussion of Low-k Dielectrics

Since the waveguides are suspended, they must have enough mechanical strength to survive the etching process and subsequent use, so the properties of the dielectric layers used in the BEOL are critical. For processes with low-k dielectric layers, the dielectrics that need to be considered are the low-k layers, silicon dioxide layers, and possibly etch stop layers. In a typical 180-nm process, the low-k dielectric is fluorosilicate glass (SiOF) [19], [20], [21], which has a refractive index of 1.43 [20], while a typical 65-nm process uses organosilicate glass (SiCOH) [21], [22], which has a refractive index of 1.40 at a wavelength of 633 nm [23]. Fortunately, these refractive indices are not far off from silicon dioxide's refractive index of 1.46 at a wavelength of 633 nm [24], [25], allowing simple modeling to be done using silicon dioxide alone, but more careful modeling with all layers will produce waveguide modes that are slightly distorted in comparison to the modes of an equivalent silicon dioxide waveguide (Fig. 24).

The low-k dielectrics are weaker than silicon dioxide, with greater susceptibility to cracking and peeling. These properties have been thoroughly investigated by the CMOS interconnect research community, which has produced a large number of papers reporting the results of mechanical tests [26]. The low-k layers may also have poor adhesion to dielectric layers above and below them. Again, this issue has been thoroughly investigated [26]. One property of the dielectric stackups that is of particular relevance to suspended waveguides is the mechanical stress induced in the dielectric layers due to the fabrication process [20]. Differing coefficients of thermal expansion between materials result in stress as the chip experiences thermal cycling during fabrication [27]. This built-in stress is released when metal is etched away, and sufficient stress can lead to dielectric warping. Consequently, waveguide design must take all these effects into account.

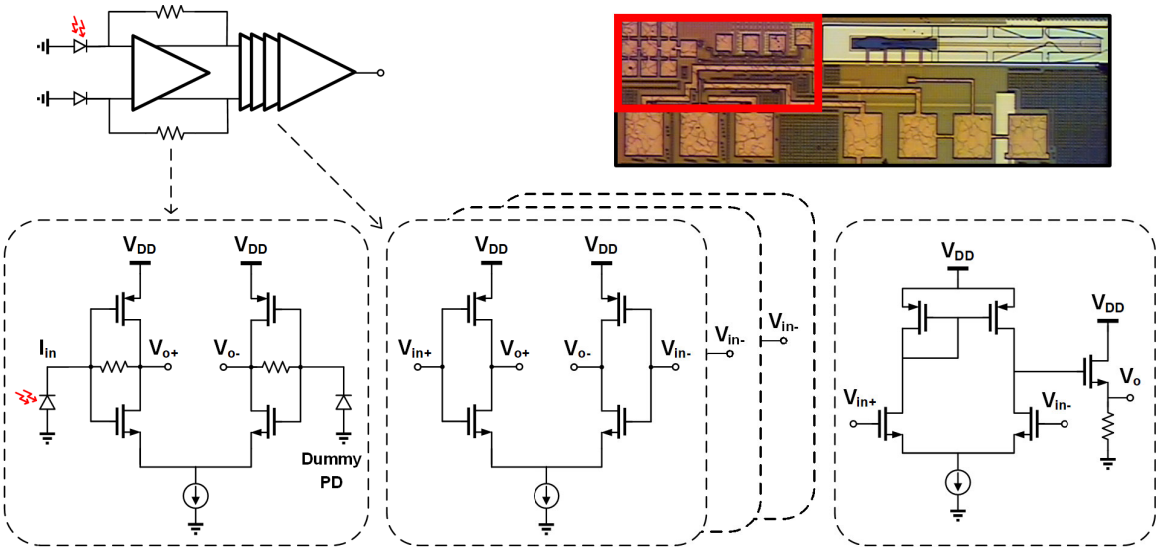

Fig. 19. Schematic of the TIA, which is highlighted in red in the die photograph.

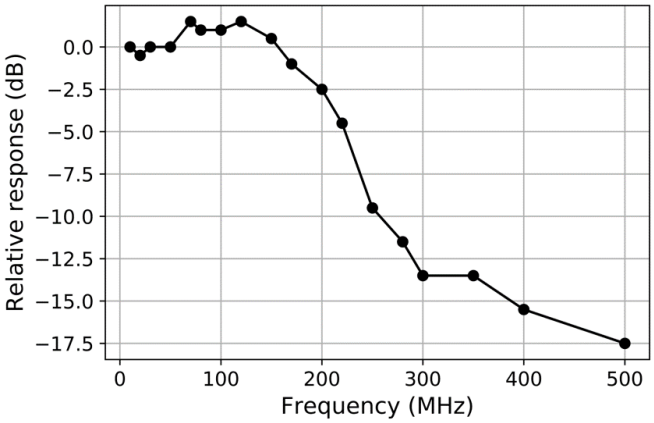

Fig. 20. Measured bandwidth of the TIA.

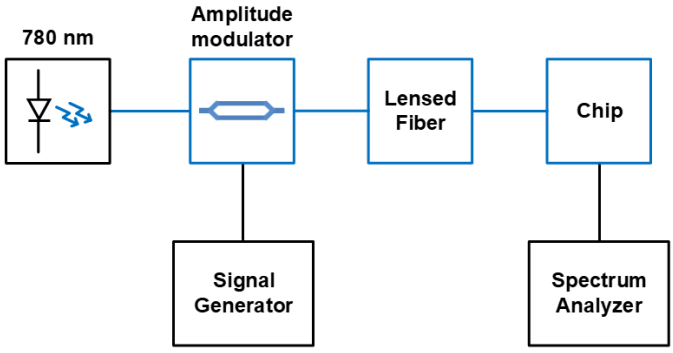

Fig. 21. Measurement setup for the TIA bandwidth measurement. The lensed fiber was coupled to the photodiode vertically and not through the waveguide.

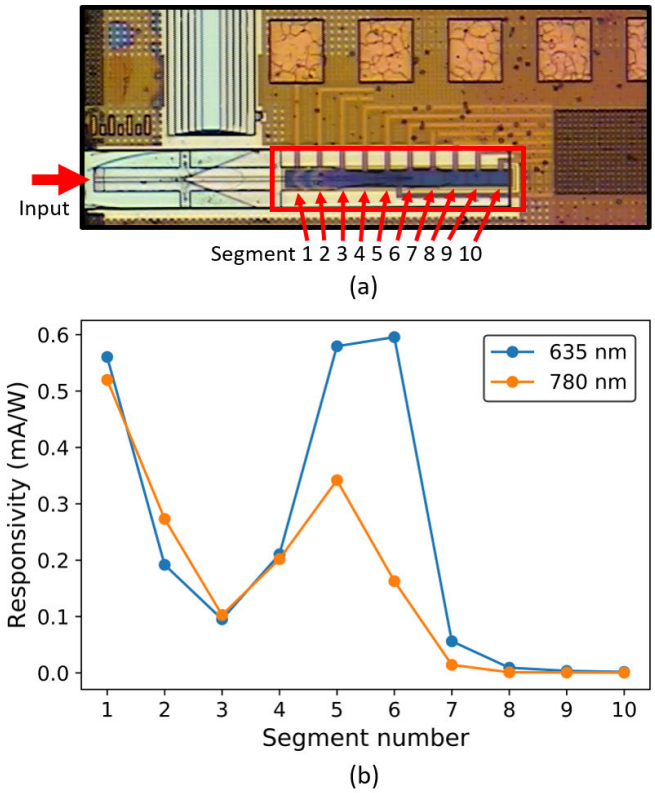

Fig. 22. (a) Die photograph and (b) responsivity of a segmented CMOS photodiode. Note that the input arrow denotes the direction of incident light.

B. Rib Waveguide Design

The modes of a waveguide may generally be divided into confined modes and radiation modes. In the case of a suspended rib waveguide, however, the slab introduces nuance. In cross sections of the slab waveguide that include supports, the slab is unbounded in width and so supports radiation modes. In cross sections of the slab that do not include supports, the slab is finite in width and supports confined modes. However, the periodic presence of the supports renders the confined slab modes lossy, especially in bends, so designers can reasonably consider the slab as infinite during analysis. In this case, the confined modes are defined as those modes with an effective index greater than the effective index of the surrounding slab. We consider "slab modes" to be those

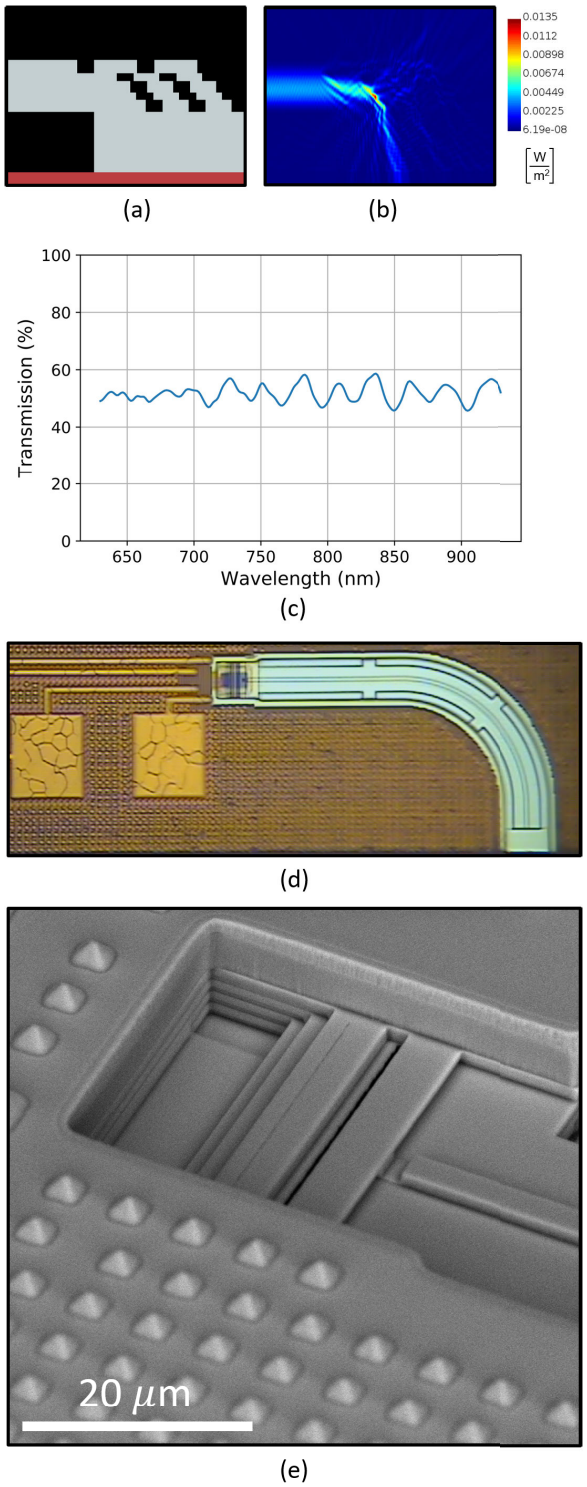

Fig. 23. (a) Schematic side view of the reflecting coupler. (b) FDTD simulation of light coupling into the silicon substrate at 780 nm. (c) Plot of the FDTD simulated transmission of the reflecting coupler into the silicon substrate. (d) Die photograph of the measured reflecting coupler with pads and input waveguide. (e) SEM image of the top of a reflecting coupler.

radiation modes that are confined in one dimension in the infinite slab and possess an effective index higher than the medium surrounding the waveguide.

Since the effective index of a confined mode in a rib waveguide must lie between the highest material index and

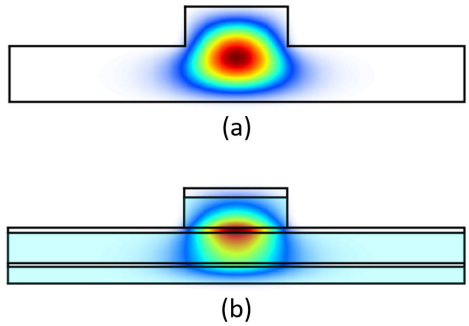

Fig. 24. (a) Waveguide mode simulated using glass only. (b) Waveguide mode simulated using the full dielectric stackup, including low-k materials denoted by light blue. Although low-k materials vary from process to process, the low-k material simulated here is SiOF with a refractive index of 1.43.

the effective index of the slab, we define the normalized index contrast as

$$\frac{n_{\text{eff}} - n_{\text{slab}}}{n_{\text{glass}} - n_{\text{slab}}} \tag{1}$$

which is used to characterize mode confinement when analyzing waveguide geometry.

1) Slab Thickness: The design of a rib waveguide in subtractive photonics begins with knowledge of the minimum acceptable thickness for the waveguide slab. We have found in our investigations of a 180-nm process that a slab composed of two layers (the dielectric associated with a metal layer and the dielectric associated with a via layer) provides sufficient mechanical strength for reasonable waveguide structures. However, if the width of the supports is large enough or if the length of the supports is small enough, a waveguide with a two-layer slab will collapse. Consequently, waveguide designers must follow certain design rules when laying out supports, just as in the design of other chip components.

A three-layer slab is stronger than a two-layer slab and can tolerate larger support widths and shorter support lengths without waveguide collapse. Moreover, the reduced index contrast enables the design of single-mode waveguides with relatively large cross section. However, the thicker slab comes at the cost of reduced mode confinement and therefore increased bend radius.

2) Ridge Geometry: As mentioned above, the index contrast between the ridge and the slab of the waveguide determines a tradeoff between bend loss and the cross section of a single-mode waveguide. The thickness of the slab and the thickness of the ridge may be chosen to optimize this index contrast for a given application, subject to the constraints of a minimum slab thickness for mechanical strength, and a maximum ridge thickness for single-mode operation. For example, given the two-layer slab criterion outlined above, a designer may choose to implement a V5M5-V4M4 geometry. This provides minimum slab thickness, while the two-layer ridge provides adequate index contrast while avoiding multimode operation.

Once the slab thickness and the ridge thickness are chosen, the designer must select a ridge width. The ridge width is constrained to both a minimum for adhesion requirements and a maximum for ensuring single-mode operation. In the

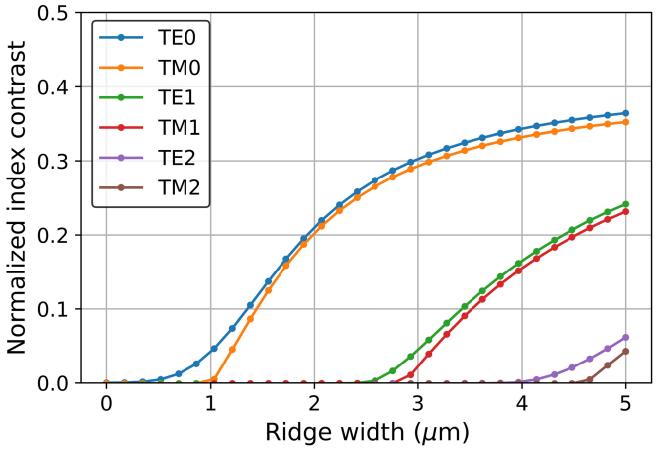

Fig. 25. Effective index of the confined modes of the V5M5-V4M4 waveguide versus ridge width at 780 nm.

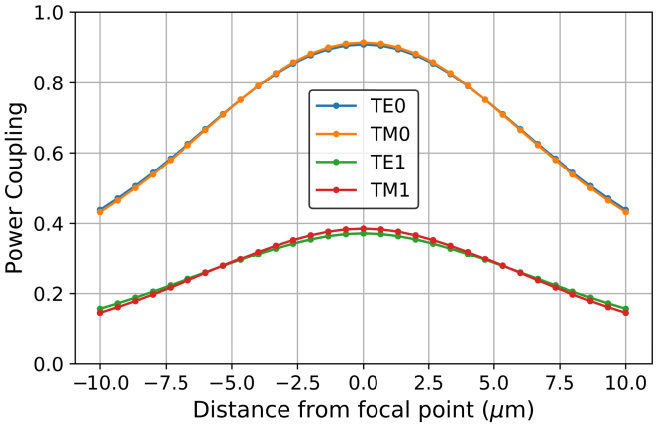

Fig. 26. Simulated fiber coupling efficiency for all confined modes of a V5M5-V4M4 waveguide and a 780-nm lensed fiber with a 2-μm-diameter spot.

180-nm process used in this work, we have found that a 3-μm-wide ridge ensures adequate ridge-slab adhesion so that may be taken as the minimum ridge width for the example V5M5-V4M4 waveguide. Fig. 25 shows the modes for this geometry as the ridge width varies, demonstrating that higher order modes are supported at the minimum ridge width of 3 μm. However, since these modes are poorly confined at that ridge width, they will be highly lossy in bends, and thus, the waveguide can be considered as quasi-single mode.

An additional consideration for ridge thickness is the efficiency of edge coupling. Fig. 26 shows the mode overlap between the 2-μm-diameter spot of a 780-nm lensed fiber and the confined modes of the V5M5-V4M4 geometry. Coupling of 90% to the fundamental TE and TM modes can be achieved, indicating the suitability of this ridge geometry for edge coupling.

3) Slab Width: As alluded to above, the slab width must be chosen to minimize both the waveguide area and the reflections due to supports. Since the slab width changes between a section of the waveguide with support and a section without, the goal is to maximize mode overlap between the fundamental modes of both sections. In Fig. 27, we used an EME simulator to plot this mode overlap over a range of slab widths for

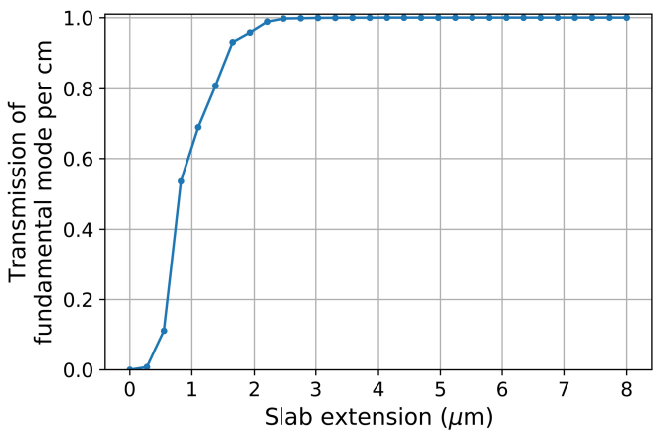

Fig. 27. Simulated support coupling efficiency for a V5M5-V4M4 waveguide at 780 nm.

the V5M5-V4M4 geometry. We can see that a one-sided slab extension of 3 μm is sufficient to minimize the loss due to support reflections.

4) Support Sizing: In general, a tradeoff exists between the size of waveguide supports and the speed of etching since the thickness, width, and length of supports determine the size of the openings that etchant flows through. Without the exact material properties of the dielectric stackup, however, it is difficult to model the structural integrity of the waveguide, especially during the etching procedure, and thus, an empirical approach is easiest. In the 180-nm process used in this work, we have found that 20 μm long and 5 μm wide supports placed every 100 μm provides excellent mechanical strength with reasonable etching times.[2]

C. Waffling

In some subtractive aluminum-oxide semiconductor processes, such as the bulk 180-nm CMOS process used to fabricate our waveguides, large, contiguous areas of metal are forbidden on the top via layer. In these cases, the designer may instead circumscribe a section of the via layer with slot vias and enclose it with metal on the layers above and below.[3] In order to facilitate the removal of the metal-enclosed dielectric during the AETA etch, we break the enclosed sections into smaller "waffles," where via arrays are placed inside the slot via perimeter. The strong adhesion of the barrier layers to the dielectric pieces can be observed in Fig. 28. While electronic design tools may still flag waffling as a violation of DRC rules, it presents a reasonable compromise between design and fabrication limitations.

Waffling is not necessary in all processes. Of the commercial semiconductor processes that subtractive photonics has been tested in, only the bulk 180-nm process optimized for RF devices required it. Other processes, such as a 180-nm process

[2]Typical etching times for many structures are about 4 h in Aluminum Etch Type A (AETA) at 80 °C, followed by 2 h in hydrogen peroxide and EDTA at 80 °C (see Appendix A). These times will vary depending on the size of the structure, the stirring rate, how often the etchant is replaced, and so on.

[3]Slot vias are often found in the seal rings provided by foundries' process development kits (PDKs), in which case their manufacturability is assured.

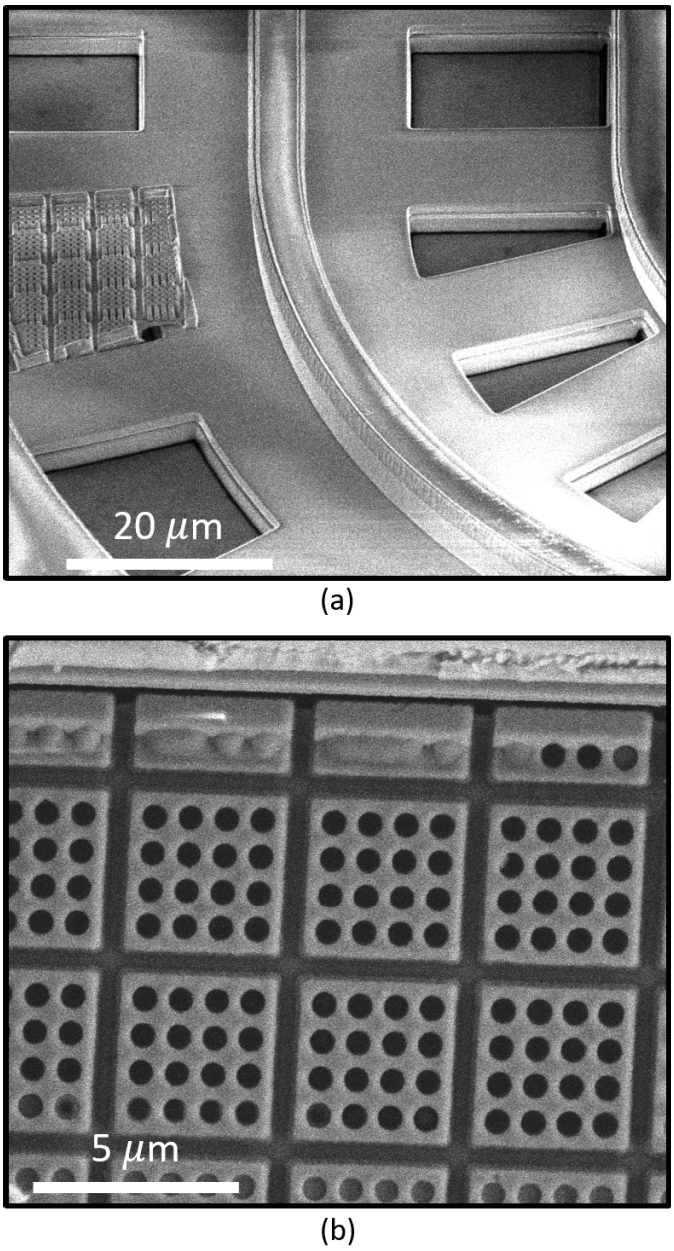

Fig. 28. (a) M6-V5M5V4 waveguide with waffling partially removed. (b) SEM image of waffling after etching. Note that the barrier layer is holding the glass pieces together.

for digital circuits and a 65-nm process for RF devices, did not require it due to variations in the fabrication process.

V. Future Outlook

As mentioned in Section I, the method of subtractive photonics is intended to supply low-cost BEOL photonic integration to arbitrary electronics processes. The waveguides demonstrated in this article can be complemented by photodiodes that are available in the CMOS process for detecting both visible and infrared light. For the former, CMOS photodiodes offer some flexibility in the type and number of junctions used for detection, potentially providing bandwidths up to 12 GHz [17]. For the latter, Schottky photodiodes have been explored that offer responsivities up to several mA/W [28].

Thus far, with minimal post-processing, we have means in a bulk CMOS process of propagating optical signals and converting those optical signals to electrical signals, in addition to all of the capabilities afforded by integrated electronics. Clearly, we are still missing a component that can convert electrical signals into optical signals—a modulator. Two possible intrinsic mechanisms are thermo-optic modulation and modulation using MEMS-like structures. With regard to the former, silicon dioxide possesses a thermo-optic coefficient comparable to silicon nitride. However, the low bandwidth of a thermal modulator would limit applications, which is also the main drawback of any potential MEMS modulator. Examining the modulation mechanisms in [29], we may consider oxide-bonded lithium niobate or deposited materials such as electrooptic polymers, transparent conducting oxides, 2-D materials, or phase-change materials. These are all viable options for high-speed modulation, but the cost of post-processing these materials may render the economic advantages of subtractive photonics untenable. In consideration of all these options, thermal modulation offers an immediate path toward full monolithic electronic–photonic systems in bulk CMOS, while further work would be necessary to implement higher speed modulation options.

The primary fiber coupling mechanism used in this article is edge coupling because the waveguides inherently have a large cross-sectional area that is advantageous for edge coupling, and the glass waveguide material is inherently index matched to the fiber. Grating couplers offer another possibility for fiber coupling, but the low index contrast of the glass–air material system gives them low coupling efficiencies.

Advanced CMOS nodes offer new tradeoffs. An immediate benefit is the larger number of wiring layers, which is expected to increase up to 19 layers in the next few years [30] and provides increased flexibility for waveguide design. Thinner metal layers used in intermediate interconnect will most likely need to be combined together for mechanical strength but might also enable single-mode waveguides with smaller cross-sectional area. Reduced sidewall roughness due to improvement in lithography is another benefit of advanced nodes. Although cobalt and ruthenium may be included in intermediate wiring in future nodes [30], there are a variety of etchants that can be used to remove those metals without affecting glass [31]. CMOS photodiodes may be possible in FinFET processes due to the well-substrate junctions available [32], [33], though to our knowledge, there has been no reported demonstration.

As mentioned in Section IV-A, a typical 65-nm process uses SiCOH for the low-k dielectric. Subsequent nodes continued using SiCOH with a focus on increasing porosity, with some advanced nodes offering air gaps [21]. The interconnect community has faced considerable difficulty integrating ultralow-k materials ($k < 2.6$) into the BEOL for a variety of reasons, including mechanical weakness and poor adhesion properties [21], [30]. This slowdown in k reduction means that the low-k dielectric has not changed much in the last couple decades, giving waveguide designers some predictability in material properties when moving to more advanced nodes.

However, the poor mechanical strength and adhesion of lightly porous or highly porous SiCOH incorporated in various processes may restrict the waveguide design space opened up by a large number of layers.

VI. Conclusion

In this article, we have presented measurement results and general design considerations for the method of subtractive photonics. We have demonstrated waveguides formed using BEOL materials that present significant waveguide loss improvements over previous demonstrations of photonics in unmodified bulk CMOS. Photonic properties are analyzed in the context of material constraints and foundry limitations. Photodiodes native to CMOS are used for detecting light, with both tapering and reflecting waveguide-photodiode couplers measured. With these devices and the low barrier of entry due to the minimal post-processing, subtractive photonics provides a promising avenue for the monolithic integration of electronics and photonics in CMOS.

Appendix A
Etching Procedure

The etchant used and the etching protocol will generally depend on the structure of the backend interconnect. In a typical 180-nm process, the backend interconnect metal is aluminum with a small percentage of copper and silicon, with tungsten plugs for vias. The tungsten plugs are insulated from the surrounding dielectric using a Ti/TiN bilayer that acts as a barrier and adhesion layer. TiN is also used as an antireflective coating for patterning of the metal layers [34]. In a typical 65-nm process, the backend interconnect is formed using a copper dual damascene process, which requires a diffusion barrier to isolate the copper from the surrounding dielectric. A Ta/TaN bilayer is typically used as the diffusion and adhesion layer [21].

The typical etching procedure starts the chips in AETA [35] at 80 °C for several hours. At this point, a visual check is performed to determine the next etching step. In some processes, such as the 65-nm process used in [15], the barrier layers inhibit continued etching after the metal above the barrier has been removed. In this case, it has been found that a 1:1 solution of hydrogen peroxide and EDTA at 80 °C for 1 h or two works well for etching the barrier layers [36], [37].[4] Etching with AETA can then proceed, with further etching of barrier layers as necessary. In other processes, such as the 180-nm process used in this work, the barrier layers do not inhibit full etching of the metal, allowing all the aluminum to be etched before the barrier layer remnants are cleaned up with hydrogen peroxide and EDTA etch.

Appendix B
Waveguide Loss Measurement

The waveguide loss measurements consist of upper bounds determined by measuring the insertion loss of long waveguide spirals relative to a free-space fiber-to-fiber baseline. The insertion loss consists of fiber-chip coupling losses, waveguide bend losses, and all other waveguide losses, which may be formulated as

$$IL = L_{\text{coupling}} + L_{\text{bends}} + L_{\text{waveguide}}. \tag{2}$$

This implies that if W is the waveguide length

$$\frac{IL}{W} \geq \frac{L_{\text{waveguide}}}{W} \tag{3}$$

which is the upper bound presented for the fabricated waveguides.

[4]For TiN in particular, another option is a solution of hydrogen peroxide and ammonium hydroxide [34].

Acknowledgment

The authors would like to thank Aroutin Khachaturian for the useful discussions and for providing SEM images. They would also like to thank the Caltech CI2 Grant and the Space Solar Power Project for their support.

References

[1] R. Soref and J. Larenzo, "All-silicon active and passive guided-wave components for λ = 1.3 and 1.6 μm," *IEEE J. Quantum Electron.*, vol. QE-22, no. 6, pp. 873–879, Jun. 1986.

[2] M. Rakowski et al., "45 nm CMOS—Silicon photonics monolithic technology (45CLO) for next-generation, low power and high speed optical interconnects," in *Proc. Opt. Fiber Commun. Conf. Exhib. (OFC)*. Washington, DC, USA: Optica Publishing Group, Mar. 2020, Paper no. T3H.3.

[3] V. Stojanović et al., "Monolithic silicon-photonic platforms in state-of-the-art CMOS SOI processes," *Opt. Exp.*, vol. 26, no. 10, pp. 13106–13121, May 2018.

[4] J. S. Orcutt et al., "Nanophotonic integration in state-of-the-art CMOS foundries," *Opt. Exp.*, vol. 19, no. 3, pp. 2335–2346, Jan. 2011. [Online]. Available: https://opg.optica.org/oe/abstract.cfm?URI=oe-19-3-2335

[5] J. S. Orcutt et al., "Demonstration of an electronic photonic integrated circuit in a commercial scaled bulk CMOS process," in *Proc. Conf. Lasers Electro-Optics Conf. Quantum Electron. Laser Sci.*, May 2008, Paper no. CTuBB3. [Online]. Available: https://opg.optica.org/abstract.cfm?URI=CLEO-2008-CTuBB3

[6] Y.-J. Hung, C.-T. Tang, T.-H. Chen, T.-H. Yen, M.-J. Tsai, and S.-L. Lee, "Low-loss polysilicon subwavelength grating waveguides and narrowband Bragg reflectors in bulk CMOS," *Opt. Exp.*, vol. 28, no. 6, pp. 7786–7798, 2020.

[7] J. S. Orcutt et al., "Low-loss polysilicon waveguides fabricated in an emulated high-volume electronics process," *Opt. Exp.*, vol. 20, no. 7, pp. 7243–7254, 2012.

[8] A. H. Atabaki et al., "Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip," *Nature*, vol. 556, no. 7701, pp. 349–354, Apr. 2018, doi: 10.1038/s41586-018-0028-z.

[9] R. Meade et al., "Integration of silicon photonics in bulk CMOS," in *Proc. Symp. VLSI Technol. VLSI-Technol., Dig. Tech. Papers*, Jun. 2014, pp. 1–2.

[10] H.-C. Ji et al., "Bulk silicon photonic wire for one-chip integrated optical interconnection," in *Proc. 7th IEEE Int. Conf. Group IV Photon.*, Sep. 2010, pp. 96–98.

[11] H. Byun et al., "Bulk-Si photonics technology for DRAM interface [invited]," *Photon. Res.*, vol. 2, no. 3, pp. A25–A33, Apr. 2014, doi: 10.1364/prj.2.000a25.

[12] I. A. Young et al., "Optical I/O technology for tera-scale computing," *IEEE J. Solid-State Circuits*, vol. 45, no. 1, pp. 235–248, Jan. 2010.

[13] J. Fedeli, R. Orobtchouk, C. Seassal, and L. Vivien, "Integration issues of a photonic layer on top of a CMOS circuit," *Proc. SPIE*, vol. 6125, Mar. 2006, Art. no. 61250H.

[14] Y. H. D. Lee and M. Lipson, "Back-end deposited silicon photonics for monolithic integration on CMOS," *IEEE J. Sel. Topics Quantum Electron.*, vol. 19, no. 2, Mar./Apr. 2013, Art. no. 8200207.

[15] R. Fatemi, C. Ives, A. Khachaturian, and A. Hajimiri, "Subtractive photonics," *Opt. Exp.*, vol. 29, no. 2, pp. 877–893, 2021.

[16] C. Ives and A. Hajimiri, "13.5 subtractive photonic waveguide-coupled photodetectors in 180 nm bulk CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2023, pp. 12–14.
[17] M.-J. Lee and W.-Y. Choi, "Performance optimization and improvement of silicon avalanche photodetectors in standard CMOS technology," *IEEE J. Sel. Topics Quantum Electron.*, vol. 24, no. 2, pp. 1–13, Mar. 2018.
[18] Q. Pan, Y. Wang, Y. Lu, and C. P. Yue, "An 18-Gb/s fully integrated optical receiver with adaptive cascaded equalizer," *IEEE J. Sel. Topics Quantum Electron.*, vol. 22, no. 6, pp. 361–369, Nov. 2016.
[19] S. Yang et al., "A high performance 180 nm generation logic technology," in *IEDM Tech. Dig.*, Dec. 1998, pp. 197–200.
[20] W. Chang, S. M. Jang, C. H. Yu, S. C. Sun, and M. S. Liang, "A manufacturable and reliable low-*k* inter-metal dielectric using fluorinated oxide (FSG)," in *Proc. IEEE Int. Interconnect Technol. Conf.*, May 1999, pp. 131–133.
[21] M. Rudan, R. Brunetti, and S. Reggiani, *Springer Handbook Semiconductor Devices*. Cham, Switzerland: Springer, 2023, doi: 10.1007/978-3-030-79827-7.
[22] S.-K. JangJian, C.-P. Liu, W.-S. Hwang, S.-W. Chen, K.-H. Wei, and Y.-L. Wang, "Eletrical and mechanical properies of nitrogen and fluorine incorporated organosilicate glass prepared by plasma enhanced chemical vapor deposition," in *Proc. Electrochem. Soc.*, vol. 10, pp. 188–194, Jan. 2003.
[23] S.-K. Kwak, K.-H. Jeong, and S.-W. Rhee, "Nanocomposite low-*k* SiCOH films by direct PECVD using vinyltrimethylsilane," *J. Electrochem. Soc.*, vol. 151, no. 2, p. F11, 2004, doi: 10.1149/1.1636738.
[24] W. A. Pliskin and R. P. Esch, "Refractive index of SiO_2 films grown on silicon," *J. Appl. Phys.*, vol. 36, no. 6, pp. 2011–2013, Jun. 1965.
[25] R. Kitamura, L. Pilon, and M. Jonasz, "Optical constants of silica glass from extreme ultraviolet to far infrared at near room temperature," *Appl. Opt.*, vol. 46, no. 33, pp. 8118–8133, Nov. 2007. [Online]. Available: https://opg.optica.org/ao/abstract.cfm?URI=ao-46-33-8118
[26] S. P. Murarka, M. Eizenberg, and A. K. Sinha, *Interlayer Dielectrics for Semiconductor Technologies*. Amsterdam, The Netherlands: Elsevier, 2003, doi: 10.1016/b978-0-12-511221-5.x5000-4.
[27] P. S. Ho, J. J. Leu, and W. W. Lee, *Low Dielectric Constant Materials for IC Applications*. Berlin, Germany: Springer, 2003, doi: 10.1007/978-3-642-55908-2.
[28] W. Diels, M. Steyaert, and F. Tavernier, "Schottky photodiodes in bulk CMOS for high-speed 1310/1550 nm optical receivers," *IEEE J. Sel. Topics Quantum Electron.*, vol. 24, no. 6, pp. 1–8, Nov. 2018.
[29] G. Sinatkas, T. Christopoulos, O. Tsilipakos, and E. E. Kriezis, "Electro-optic modulation in integrated photonics," *J. Appl. Phys.*, vol. 130, no. 1, Jul. 2021, Art. no. 010901.
[30] *International Roadmap for Devices and Systems: More Moore 2022 Update*, IEEE, Piscataway, NJ, USA, 2022.
[31] *Etchant Compatibility Chart*. Accessed: Jun. 28, 2023. [Online]. Available: https://transene.com/etch-compatibility/
[32] K. Domanski, "Latch-up in FinFET technologies," in *Proc. IEEE Int. Rel. Phys. Symp. (IRPS)*, Mar. 2018, pp. 2C.4-1–2C.4-5.
[33] M. Eberlein and H. Pretl, "A no-trim, scaling-friendly thermal sensor in 16 nm FinFET using bulk diodes as sensing elements," *IEEE Solid-State Circuits Lett.*, vol. 2, no. 9, pp. 63–66, Sep. 2019.
[34] C. S. Yoo, *Semiconductor Manufacturing Technology*. Singapore: World Scientific, Mar. 2008, doi: 10.1142/6123.
[35] *Aluminum Etchants*. Accessed: Apr. 11, 2023. [Online]. Available: https://transene.com/aluminum/
[36] *Tantalum Wet Etching for Photolithography*. Accessed: Aug. 15, 2019. [Online]. Available: https://www.finishing.com/56/48.shtml
[37] *Etching Solution for Tantalum/Tantalum Nitride*. Accessed: Aug. 15, 2019. [Online]. Available: https://www.finishing.com/178/71.shtml

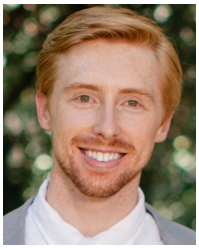

Craig Ives (Member, IEEE) received the B.S. degree in electrical engineering from the University of California at San Diego (UCSD), La Jolla, CA, USA, in 2017, and the M.S. degree in electrical engineering from the California Institute of Technology, Pasadena, CA, USA, in 2020, where he is currently pursuing the Ph.D. degree.

Debjit Sarkar (Graduate Student Member, IEEE) received the B.S. degree in electrical engineering from the University of Michigan, Ann Arbor, MI, USA, in 2019, and the M.S. degree in electrical engineering from the California Institute of Technology, Pasadena, CA, USA, in 2022, where he is currently pursuing the Ph.D. degree.

He was an Electronics Engineer with the Naval Research Laboratory, Washington, DC, USA, in 2020. His research interests include optoelectronic systems, high-speed electronics, and biophotonics.

Mr. Sarkar was awarded the National Science Foundation Graduate Research Fellowship in 2020.

Ali Hajimiri (Fellow, IEEE) received the B.S. degree in electronics engineering from the Sharif University of Technology, Tehran, Iran, in 1994, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 1996 and 1998, respectively.

He has been with Philips Semiconductors, Sunnyvale, CA, USA, where he worked on a BiCMOS chipset for GSM and cellular units from 1993 to 1994. In 1995, he was with Sun Microsystems, Santa Clara, CA, USA, working on the UltraSPARC microprocessor's cache RAM design methodology. In 1997, he was with Lucent Technologies (Bell Labs), Murray Hill, NJ, USA, where he investigated low-phase-noise integrated oscillators. In 1998, he joined the faculty of the California Institute of Technology, Pasadena, CA, USA, where he is currently the Bren Professor of electrical engineering and medical engineering, the Director of the Microelectronics Laboratory, and the Co-Director of the Space Solar Power Project. In 2002, he co-founded Axiom Microdevices Inc., whose fully integrated CMOS PA has shipped around 400 000 000 units, and was acquired by Skyworks Inc. in 2009. He is the author of *Analog: Inexact Science, Vibrant Art* (2020, Early Draft), a book on fundamental principles of analog circuit design, and *The Design of Low Noise Oscillators* (Boston, MA, USA: Springer). He has authored and coauthored more than 250 refereed journal articles and conference technical articles and has been granted more than 160 U.S. patents with many more pending applications. His research interests are high-speed and high-frequency integrated circuits for applications in sensors, photonics, biomedical devices, and communication systems.

Dr. Hajimiri is a fellow of the National Academy of Inventors (NAI). He was selected to the TR35 top innovator's list. He won the Feynman Prize for Excellence in Teaching, Caltech's most prestigious teaching honor, and Caltech''s Graduate Students Council Teaching and Mentoring Award and the Associated Students of Caltech Undergraduate Excellence in Teaching Award. He was the winner of the Gold Medal at the National Physics Competition and the Bronze Medal at the 21st International Physics Olympiad, Groningen, Netherlands. He was recognized as one of the top-10 contributors to ISSCC. He was a co-recipient of the IEEE JOURNAL OF SOLID-STATE CIRCUITS Best Paper Award of 2004, the International Solid-State Circuits Conference (ISSCC) Jack Kilby Outstanding Paper Award, the RFIC Best Paper Award, the CICC Best Paper Award two times, and a three-time winner of the IBM Faculty Partnership Award, the National Science Foundation CAREER Award, and the Okawa Foundation Award. He has served on the Technical Program Committee of the International Solid-State Circuits Conference (ISSCC). He has served as an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS (JSSC) and the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: EXPRESS BRIEFS, a member for the Technical Program Committee of the International Conference on Computer Aided Design (ICCAD), a Guest Editor for the IEEE TRANSACTIONS ON MICROWAVE THEORY AND TECHNIQUES, and a Guest Editorial Board Member of *Transactions of Institute of Electronics, Information and Communication Engineers of Japan* (IEICE). He has served as a Distinguished Lecturer for the IEEE Solid-State and Microwave Societies.