

MIT Open Access Articles

*Photonic-electronic integration with
polysilicon photonics in bulk CMOS*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Ram, Rajeev J. "Photonic-Electronic Integration with Polysilicon Photonics in Bulk CMOS." Edited by Graham T. Reed and Michael R. Watts. Silicon Photonics X (April 3, 2015). © 2015 Society of Photo-Optical Instrumentation Engineers (SPIE)

As Published: <http://dx.doi.org/10.1117/12.2175462>

Publisher: SPIE

Persistent URL: <http://hdl.handle.net/1721.1/101003>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Photonic-electronic integration with polysilicon photonics in bulk CMOS

Rajeev J Ram

Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA, USA 02139

ABSTRACT

Here, I review the development of a polysilicon photonic platform that is optimized for integration with electronics fabricated on bulk silicon wafers. This platform enables large-scale monolithic integration of silicon photonics with microelectronics. A single-polysilicon deposition and lithography mask were used to simultaneously define the transistor gate, the low-loss waveguides, the depletion modulators, and the photodetectors. Several approaches to reduce optical scattering and mitigate defect state absorption are presented. Waveguide propagation loss as low as 3 dB/cm could be realized in front-end polysilicon with an end-of-line loss as low as 10 dB/cm at 1280nm. The defect state density could be enhanced to enable all-silicon, infrared photodetectors. The resulting microring resonant detectors exhibit over 20% quantum efficiency with 9.7 GHz bandwidth over a wide range of wavelengths. A complete photonic link has been demonstrated at 5 Gbps that transfers digital information from one bulk CMOS photonics chip to another.

INTRODUCTION

For high-volume, high-density VLSI applications (e.g. microprocessors, systems-on-chip, field-programmable gate arrays (FPGAs)), and memory (DRAM), bulk-silicon wafers remain the dominant production platform. However, thick-buried-oxide, silicon-on-insulator (SOI) wafers have been the dominant monolithic silicon photonic platform due to the easy formation of low-loss waveguides.

We review recent progress of an effort led by the Ram and Orcutt (MIT), Meade and Sandhu (Micron), Stojanovic (UC Berkeley), and Popovic (CU Boulder) research groups to enable the integration of photonic components within the front-end of a bulk CMOS VLSI process. This approach allows tight and large-scale monolithic integration of silicon photonics with microelectronics. A complete photonic link has been demonstrated at 5 Gbps that transfers digital information from one bulk CMOS photonics chip to another [1,2]. A single-polysilicon deposition and lithography mask were used to simultaneously define the transistor gate, the low-loss waveguides, the depletion-mode optical modulators, and the infrared photodetectors. The process was finely tuned to adjust the density of localized electronic states associated with the grain-boundaries in polysilicon to facilitate each of these functions.

E-mail: rajeev@mit.edu, Telephone: +1 (617) 253-4182, Webpage: <http://www.rle.mit.edu/sclaser/>

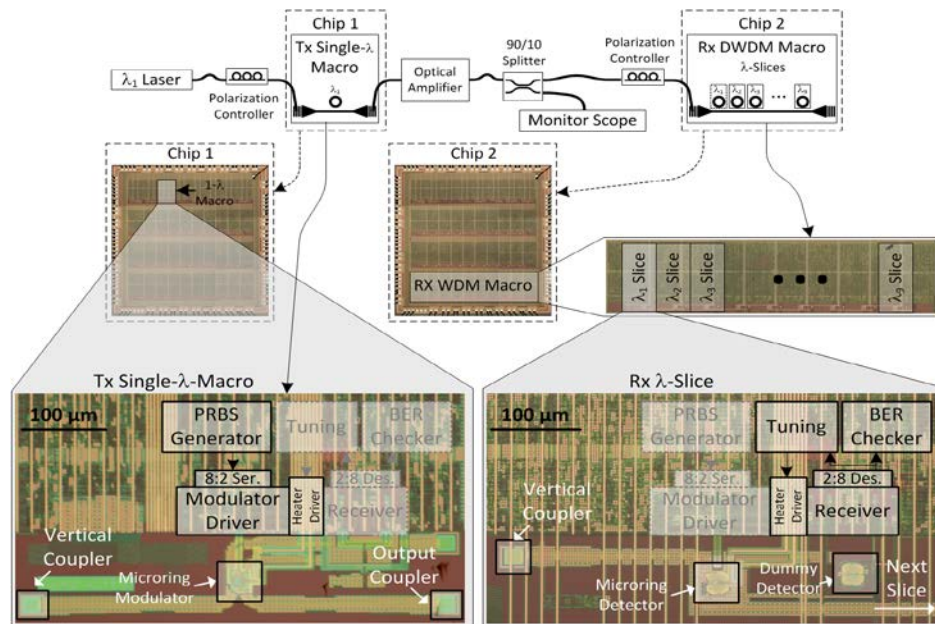


Figure 1. Monolithically integrated transmitters and receivers were used to communicate information between two bulk CMOS chips at 5 Gbps. [1]

The design and fabrication challenge for monolithic integration of photonics within an existing process is to identify a suitable set of layers that provide enough flexibility to produce all of the required devices. The DRAM process manufactures a single product and therefore process technology is optimized to that purpose. This specificity allows for process modifications to integrate photonic devices with optimized processes by introducing steps compatible with the existing process flow. These changes do not present the problem they would in CMOS since there is no DRAM foundry system that forces process uniformity.

ZERO-CHANGE POLYSILICON WAVEGUIDES

Before discussion of polysilicon optimized for photonics, it can be instructive to consider the optical properties of deposited polysilicon that serves as the gate in CMOS electronics. A key advantage of zero-change integration is that it preserves the performance and integrity of the electronics while adding simple photonic functionality for little effort and cost. As mentioned earlier, bulk CMOS represents the largest fraction of microprocessors, foundry ASICs, FPGAs, and all of DRAM processes. Zero-change polysilicon photonics allows for rapid implementation of photonic designs in advanced CMOS. An illustration of this is the figure below. This chip (MIT tapeout June 25, 2007 at Texas Instruments) was the first photonics realized in 300 mm CMOS process and also happened to be the first CMOS photonics in a <100nm technology node [3]. Within this 65nm process, a design infrastructure for integration of photonics within a traditional CMOS design environment was developed and demonstrated.

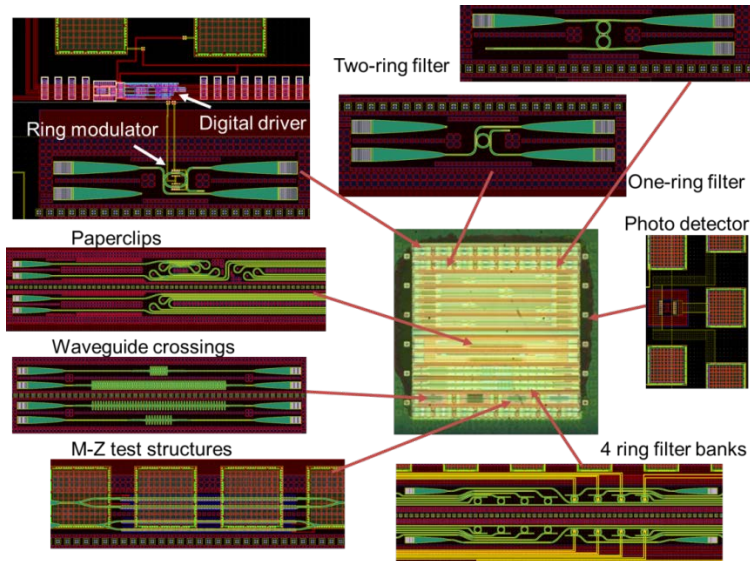


Figure 2. Micrograph of first zero-change CMOS photonic chip fabricated in a Texas Instruments 300mm, 65nm-technology facility. All waveguides utilized thin polysilicon [3].

The similarity of CMOS logic processes allowed the rapid development of polysilicon photonics for more advanced nodes. A 28nm bulk CMOS process (again with a 300mm) was used to demonstrated microring resonator based filter banks and tunable components in 2010 within a multiproject electronics chip [4]. In both cases, polysilicon photonics enabled photonic functionality in processes capable of yielding 2 billion transistor circuits.

The fundamental problem for zero-change photonic integration is the optical loss associated with the polysilicon developed for transistor gate processes. The polysilicon waveguide loss in 65 nm CMOS is observed to be 55 dB/cm (near 1550nm) and 76 dB/cm (at 1280 nm) in 28 nm CMOS. The sources of loss that one can consider include:

- Top surface roughness,
- Line edge roughness (LER),
- Polysilicon material loss associated with defect states,
- Absorption associated with the nitride and other dielectric liners used in the CMOS gate process.

The line edge roughness is well-characterized in a CMOS process. In fact, the ITRS roadmap includes LER as a target metric for the development of technology nodes. Based on the estimated roughness – for example the ITRS 2011 target is 1.9 nm, the LER is expected to contribute less than 1.5 dB/cm of loss for a silicon based waveguide [5]. Nitride material loss could be as high as 20dB/cm but the small overlap between the optical mode and the nitride liner layers suggests an upper limit of loss closer to 9dB/cm. In fact, the dominant loss process is the top surface roughness.

The roughness of the top surface of the polysilicon has little importance to electrical functionality. The polysilicon grain structure is deposited in the columnar grain structure that is optimized to promote vertical diffusion of dopants along grain boundaries. Since the grains are growing during deposition, the surface is formed by the differential growth rates for the crystal orientations of each grain. The resulting rough surface is also clearly visible. Extrapolating theoretical loss curves from a published sidewall roughness analysis, the waveguide sensitivity to top-surface roughness is approximately 10 dB/cm per nm RMS roughness for the measured 100 nm correlation length [5]. The close agreement of the 50 dB/cm prediction offered further confirmation that the top surface roughness was the dominant contribution to overall integrated polysilicon waveguide loss.




Figure 3 Optical field simulation and mode parameters in thin polysilicon waveguide.




Figure 4 Cross-sectional TEM of a polysilicon waveguide integrated in a 28nm bulk-CMOS process. The core thickness is approximately 73 nm with the surface roughness of 5 nm RMS with a 100 nm correlation length [4].

Beyond illustrating the influence of polysilicon deposition conditions on surface roughness and optical loss, zero-change polysilicon waveguides in the bulk-CMOS process may still offer a suitable photonic platform for many applications. Weakly-coupled rings with quality factors approaching 8,000 demonstrate that the waveguide loss is low enough to enable resonant photonic devices [4]. Modulators, detectors and wavelength division multiplexing filters have been demonstrated in this platform. This platform may be useful for systems that do not require long distance on-chip routing where waveguide loss would cause the largest performance impact.

OPTIMIZED POLYSILICON WAVEGUIDES

It is possible to optimize the deposition and processing conditions used for waveguide fabrication. Such an optimization step may be inserted within a CMOS process during the development stage if the future integrated photonic application becomes of sufficient importance to justify the investment. Similar process adjustments that enable functionality beyond basic digital circuits has precedence in the analog circuits community that has resulted in the inclusion of high quality capacitors, triple-well implants and modified threshold transistors.

Waveguide performance was assessed under multiple polysilicon deposition and anneal conditions through a series of experiments in collaboration with Micron Technology. The primary fabrication optimization condition was to reduce the top surface roughness of the polysilicon layer that was determined to be the limiting loss mechanism in the bulk-CMOS work. Two approaches have been taken to reduce the top surface roughness to integrate low loss polysilicon waveguides:

- optimizing deposition and anneal conditions; and
- post-deposition polish recipe development.

dB/cm	Thin (120nm) End-of-Line		Thick (220nm) Crystallized		Thick (220nm) End-of-Line	
	1280	1550	1280	1550	1280	1550
Wavelength	1280	1550	1280	1550	1280	1550
Zero Change *80 nm thick [4]	76*	55*	-	-	-	-
amorph-Dep + Anneal [7]	33	18	6	3	-	18
Columnar Dep+ Polish	-	-	14	6	18	-
Columnar Dep+ Polish + Implant + Anneal [6]	-	-	-	-	14	-
Columnar Dep+ Polish + Implant + Anneal + LPCVD Liner [6]	-	-	-	-	10	-

Table 1 Summary of measured polysilicon waveguide losses for various process conditions. All losses are reported for 2000 nm wide waveguides. The process evolved from 76 dB/cm (1280nm) loss in zero-change CMOS [4] to an optimized polysilicon with 10 dB/cm (1280nm) loss in a memory periphery process [6]. The amorphous deposited polysilicon was deposited at low-temperature and annealed at 950C for 20 seconds. For this process alone, the end-of-line loss is reported with thermal processes that emulate a full-flow. All other data is processed to end of line with integrated electronics and a complete back-end stack-up.

To avoid a rough top surface formed by grain growth during deposition, an amorphously deposited silicon layer that is crystallized during later thermal processing was considered. The amorphous nature of the initial deposition results in a smooth film analogous to a traditional dielectric deposition. The crystallization process then results in a square grain structure that does not significantly roughen the film top surface as has been shown in past work. Although the amorphous silicon in the standard process was incidentally crystallized by the later thermal cycling of the electronics process, specific

anneals were added to the process to control crystallization conditions. Since the polysilicon is deposited at an early point of the process with a very high thermal budget, the insertion of additional anneals does not pose process integration concerns. The initial study focused on two types of crystallization conditions: a rapid $\sim 950^\circ\text{C}$, 20 second spike anneal and a longer $\sim 600^\circ\text{C}$, 30 minute anneal. Due to the higher activation energy of crystal nucleation as compared to crystal growth, it is expected that the longer, lower temperature anneal would result in a larger grain size than the higher temperature anneal. In agreement with prior literature, both anneal conditions resulted in smooth films. As measured by atomic force microscopy (AFM), the post-crystallization roughness remains between 0.3 nm and 0.5 nm RMS for both anneal conditions.

The other roughness-reduction approach has been to polish the polysilicon. Due to the differential crystal plane etch rates, we removed the chemical component of the polish. Instead, a purely mechanical polish was explored for producing smooth polysilicon films. The resulting loss for wide waveguides has been summarized in the table above for 1280nm and 1550nm.

In comparison to the high temperature crystallization condition, the lower temperature anneal and columnar poly-Si recipe are observed to have higher loss. Since the top surface roughness of these films were measured to be equal to or lower than the $\sim 950^\circ\text{C}$ anneal condition, the difference is expected to result from a change in the bulk material loss. A loss that scales with the bulk of the polysilicon material can result from either defect state absorption or scattering at the grain boundaries. These two phenomena are expected to have inverse scaling relations with grain size. Defect state absorption should be proportional to the relative volume of grain boundaries in the material and should decrease with larger grain sizes. For grain sizes below the wavelength of light in the media, scattering is expected to result in larger propagation losses as the grain size and therefore correlation length of the index heterogeneity increases.




Figure 5 Cross-sectional TEM of optimized polysilicon waveguide [6].

The losses of the various test cells were then measured. Several trends emerge from the data. First, reduced waveguide widths enable confinement factor scaling to reduce the contribution of the propagation loss from the polysilicon material. Measured propagation loss scales with decreasing confinement factor down to approximately: A waveguide width of 350 nm enables the lowest reported propagation loss of 6.2 dB/cm for a thin-core polysilicon waveguide at 1550 nm. A final trend is a consistent increase in waveguide propagation loss at shorter wavelengths for wide waveguides. Since the majority of the optical power is guided in the polysilicon core region at these wide widths, this increase can be attributed to the bulk material loss scaling with wavelength. The confinement factor scaling trends observed for all wavelengths suggest that this assumption is justified. The material loss was calculated by accounting for electronic transitions between mid-gap states localized at grain boundaries and Bloch states in the conduction and valence bands.




Figure 6 (a) Extracted polysilicon material loss with data from 2.0 μm wide waveguides plotted as squares and 1.5 μm waveguides as diamonds. (b) Polysilicon density of states, plotted as a function of energy, used for the fit calculations - 23% difference in the height of the peak (density of defects) around 0.35 eV produced the difference between the two fits in (a). The pinned Fermi level and example optical transitions are indicated. [7]

DRAM PROCESS COMPARISON

Although the general process flow and required tools of the DRAM manufacturing process greatly resembles that of CMOS, there are several key differences. As a primary constraint, the low cost model of memory manufacturing prohibits the use of the expensive starting substrates that are required for SOI-based processes. Instead, the memory process utilizes similar starting silicon substrates to bulk-CMOS processes. High density capacitors are fabricated in the pre-metal dielectrics separating the transistors from the metal layers. The capacitors themselves and the connections to them are constructed by up to five different polysilicon layers to create a front end process of far greater complexity than that found in traditional CMOS processes. In contrast, the back end is much simpler containing only two or three metal layers due to the simpler circuit topologies.

Since the business model of the memory manufacturers differs greatly from the foundry model common in the CMOS industry, the integration constraints differ greatly. First, processing masks are never shared among different projects. Second, the processing steps are customized for optimal performance of every product. These factors enable process modifications to optimize photonic devices. The chief challenges are to minimize the additional cost of integration and to prevent modifications of electronic device characteristics. Any additional processing steps required for photonic integration should then be designed as the minimum number of simple modules that can be inserted into the existing memory process without shifting transistor characteristics.

DRAM PROCESS INTEGRATION

Similar to a bulk-CMOS process, the polysilicon used for the transistor gate formation is used as the high index waveguide core. Additional process modules can be added to aid in device creation. For example, instead of requiring local substrate removal, deep trench isolation may optionally be added to the process to eliminate the requirement of further post-processing to enable optical functionality. Further functionality for device creation can be added by adding a partial etch to enable two thicknesses of the waveguide core. A basic set of the required photonic devices that must be integrated in the process is shown below.

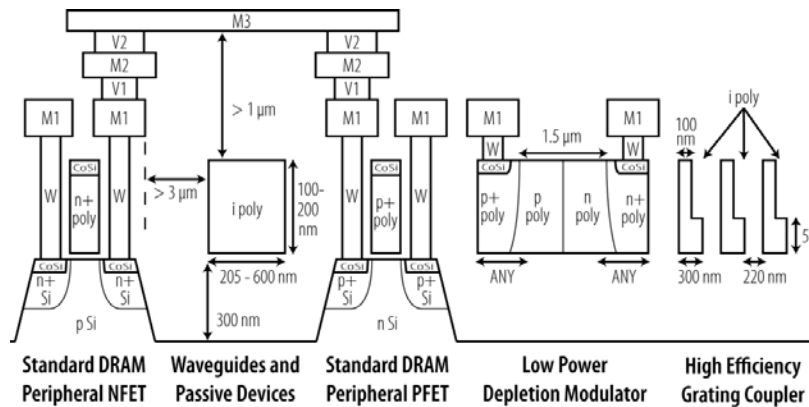


Figure 7 Monolithic DRAM photonic cross-sectional integrated device.

To test the required processes and demonstrate monolithic integration, a modified memory process has been developed by adding photonics-specific processing modules into an existing product manufacturing line. We have demonstrated photonics fabricated in a commercial memory wafer fab, operated by Micron Technology, running a short-flow polysilicon process that emulates a full-flow commercial DRAM process [6]. This DRAM emulation process was subject to many of the constraints of the full-flow process: Photonics devices are fabricated on a 1.2 μm -thick, deposited Deep Trench Isolation (DTI) oxide layer to provide optical confinement. Nitride and oxide dielectrics are present in the near field of the devices, and a single copper layer contacts active silicon layer through tungsten plugs. Instead, only the steps associated with the peripheral transistor process are explicitly performed. A technology computer aided design (TCAD) simulation of the complete proposed process is shown below




Figure 8 Monolithic DRAM integration process TCAD simulation.

ACTIVE PHOTONIC DEVICES IN POLYSILICON

Active photonic functions require us to facilitate charge injection and extraction from the polysilicon devices. The waveguide structures are embedded with p-n junctions that are connected to circuits via metal contacts. The overlap of both metals and dopants with the optical mode must be managed so as to minimize optical loss. For both modulator and detector structures, the contacts are separated from the optical mode by employing a ridge waveguide as illustrated in Figure 8. The ridge is achieved by partially etching the deposited polysilicon. As the dopants must have close proximity with the optical field – for both large modulation efficiency and efficient carrier extraction, the doping density should be adjusted to balance the opposing desire to minimize optical loss and provide effective carrier transport.




Figure 9 Cross-sections for polysilicon modulators and photodetectors deposited on a 1200 nm thick deep trench isolation (DTI) oxide. A partial etch is used to realize ridge waveguides for both structures. A p-n junction is used for efficient depletion width modulation and a p-i-n junction is used for efficient extraction of photogenerated carriers.

For transistors fabricated in polysilicon, high doping levels ($>10^{19} / \text{cm}^3$) are usually utilized. For depletion-mode carrier plasma modulators and high-speed photodetectors, moderate doping levels closer to $10^{18} / \text{cm}^3$ are useful due to the reduced free-carrier absorption, and larger depletion region width modulation. However, the activation of dopants as a function of implant concentration is non-unity in polysilicon and is a highly non-linear function [8]. If the doping concentration is comparable to the density of defect states associated with the grain boundaries, then the dopant activation can approach zero. Essentially the defect states prevent the dopants from contributing free-carriers to the polysilicon. As the doping concentration increases above the defect density, the activation can abruptly rise to 100%. Hence, the activated carrier concentration is a very sensitive function of grain structure and thermal processing for the mid-level carrier concentrations necessary for active photonics. Achieving mid-level doping densities in polysilicon requires careful calibration of the activated carrier concentration for the specific processing conditions of the polysilicon. Hall measurements were used to confirm that activated dopant concentrations were achieved in the ranges $7\text{--}9 \times 10^{17} / \text{cm}^3$ n-type, and $6\text{--}8 \times 10^{17} / \text{cm}^3$ p-type.

The diodes in Figure 9 were embedded in similar resonant microring structures. The microring resonantly enhances the effect of carrier density to modulation to yield a large transmission modulation. Figure 10 shows a measured 5 Gbps eye-

diagram with 7.6 dB of extinction. The optical insertion loss of the modulator was measured to be 1.6 dB. This same modulator is integrated with drive electronics and assembled into a WDM transmitter array as illustrated in Figure 1 [1,2].

Embedding the polysilicon within a resonator enabled us to use the small defect state absorption (illustrated in Figure 6) for efficient photodetection. The free-carrier (either electron or hole) that is associated with defect state absorption generates a displacement current within the intrinsic region of the diode. The PIN ridge-waveguide microring photodetectors with responsivities of $\sim 0.2\text{A/W}$. The device exhibits 3 dB bandwidths of 1.5 GHz and 9.7GHz at -1V and -15V biases, respectively [9]. The detectors were monolithically integrated with sense electronics as illustrated in Figure 1 [1,2]. The resulting receiver and back-end electronics enabled on-chip measurement of bit error rate (BER). The receive electronics had a sensitivity of $<10\ \mu\text{A}$ for data rates from 1-3 Gbps at a BER of 10^{-10} the sensitivity rose to $64\ \mu\text{A}$ at 5 Gbps as the bandwidth limits of the electronics was approached.

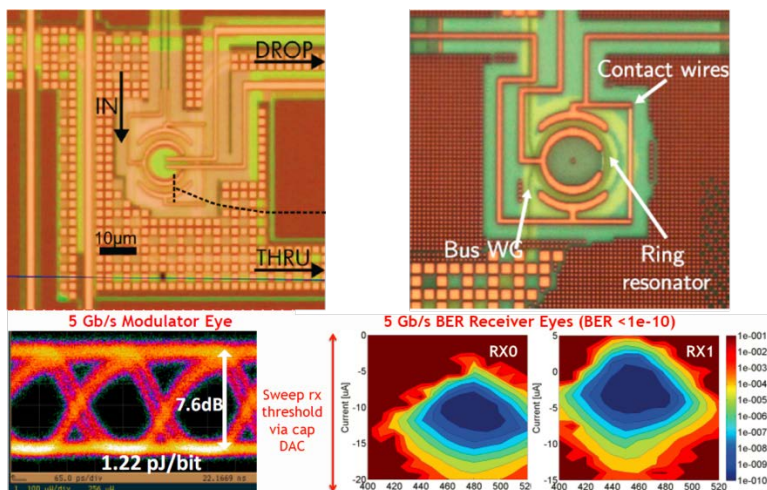


Figure 10 Polysilicon ridge-waveguide modulators and detectors in Micron bulk CMOS process. Also shown are eye diagrams for 5 Gbps operation of both transmitters and receivers [1,2].

The devices demonstrate that energy-efficient, depletion-mode optical modulators and photodetectors can be made in a polysilicon device layer, and fully compatible with a bulk silicon, including DRAM, process flow, thereby circumventing the need for hybrid integration

ACKNOWLEDGEMENTS

We acknowledge support from the DARPA POEM program under award HR0011-11-C-0100.

REFERENCES

- [1] Chen Sun, et al. VLSI Sym, 2014
- [2] Chen Sun, et al. J. Solid-State Circuits, 2015
- [3] Jason Orcutt, et al. Conf. on Lasers and Electro-Optics, 2008
- [4] Jason Orcutt, et al. Optics Express, 2011
- [5] Tymon Barwicz, et al. J. Lightwave Technology 2005
- [6] Roy Meade, et al. VLSI Symposium, 2014
- [7] Jason Orcutt, et al, Opt. Exp. 2012
- [8] Seto, Appl. Phys. Lett, 1975
- [9] Karan Mehta, et al. Optics Letters 2014