Photonic Device Layout Within the Foundry CMOS Design Environment

Jason S. Orcutt and Rajeev J. Ram

Abstract—A design methodology to layout photonic devices within standard electronic complementary metal–oxide–semiconductor (CMOS) foundry data preparation flows is described. This platform has enabled the fabrication of designs in three foundry scaled-CMOS processes from two semiconductor manufacturers.

Index Terms—Complementary metal–oxide–semiconductor (CMOS) integrated circuits, design automation, optical device fabrication, optical planar waveguide components.

I. INTRODUCTION

DENSE integration of photonic devices with electronic integrated circuits has been actively pursued by researchers for several decades. Recently, applications ranging from many-core processors [1] to interleaved analog-to-digital converters [2] have created a demand for chips where hundreds to thousands of photonic devices are integrated with millions of transistors. This scaling requires high bandwidth transistors and a manufacturing infrastructure that can only be found in silicon electronic technology [3]–[5].

Recently, we demonstrated strong-confinement waveguides and wavelength-division-multiplexing (WDM) filter banks in the polysilicon layer within a 65-nm bulk-complementary metal–oxide–semiconductor (CMOS) technology on a 30-cm silicon wafer [6], [7]. The embedded photonics shared a mask with 10^7–10^9 transistors with f_T and $f_{\max}$ of 150 and 80 GHz, respectively. The realization of photonics in standard foundry-CMOS requires strict compliance with the established process flow. As a result, the silicon photonic devices must be designed in existing physical layers and patterned with existing process masks created within the standard electronic computer-aided design (ECAD) environment used by all CMOS foundry customers.

The scaled-CMOS processes of interest employ 40–50 masks to fabricate dense transistor patterns on a 100- to 200-nm pitch connected by 6–10 interconnect layers. To ensure reliable fabrication, physical verification of device layout is required. This verification process is structured as a set of greater than 40 000 software checked design rules based on the physical constraints
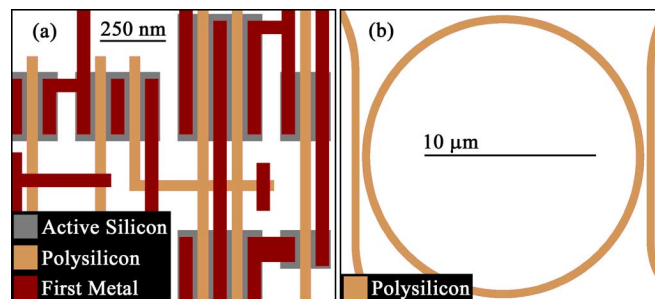
Fig. 1. Layouts showing (a) electronic and (b) photonic size scale disparity.
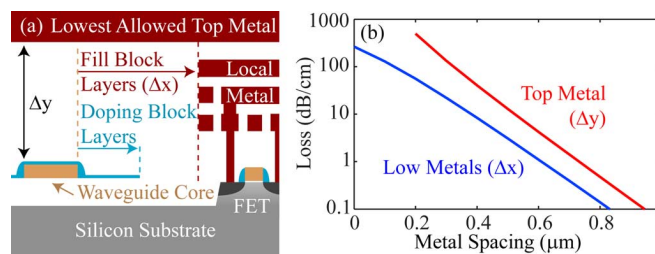

Fig. 2. (a) Cross section showing the relevant design layers, and (b) the impact metal spacing on waveguide loss. Simulation performed at $\lambda = 1550$ nm for a 600×100 nm core. Simulation details: 300-nm cladding under core; $n_{\text{substrate}} = 1$ (local substrate removal, see [7]), $n_{\text{core}} = 3.55$, $n_{\text{clad}} = 1.46$, $n_{\text{cu}} = 0.76$, $k_{\text{cu}} = 10.4$.

of the wafer processes, lithographic projection distortions, and mask fabrication. Multilayer copper interconnect processes require strict adherence to rules for the areal density of metal. These finely tuned CMOS processes are not optimized for the geometries, size scales, and local density variations of photonic devices illustrated in Fig. 1.

Relevant design layers for photonic devices, shown in Fig. 2(a), can be divided as follows: core, doping block, and fill block. Waveguide core layers include both polysilicon and active silicon in the silicon-on-insulator (SOI) process or just polysilicon in the bulk process. Doping block layers are the combination of design layers needed to block doping implants and metal silicidation of the waveguide core material. Fill block layers prevent the automatic insertion of unwanted metal layers. The extent of these metal exclusions, governed by the losses shown in Fig. 2(b), are typically 2–3 μm laterally and only 1 μm vertically. The low required vertical exclusion enables global metals (typically $>$1.5 μm above the polysilicon) to be routed irrespective of the photonics, simplifying the high level design and allowing optimal power and clock distribution. Of greater concern is that the required lateral exclusion on low metals interrupts uniform pattern density. However, the length scales (25–150 μm depending on process) over which the density

must be uniform for the chemical mechanical polishing and patterning steps are much longer than these ∼5-μm regions. As a result, higher metal density surrounding photonic regions can be added to meet design rule targets for minimum pattern density (15%–25% depending on process). By allowing the layout of photonics using standard process layers compliant with all design rules, existing scaled-CMOS foundries can be leveraged to produce highly integrated electronic–photonic circuits.

II. DESIGN ENVIRONMENT INTEGRATION

Although the most important enabler of increased integrated circuit performance has been the progress of fabrication technology, improvements in the scope and sophistication of ECAD tools have also been critical to enable the reliable production of increasingly complex integrated circuits [8]. The modern ECAD-based paradigm partitions design complexity into well-defined stages. The geometric objects placed on all design layers are checked against the accepted set of design rules in an automated process known as design rule checking (DRC). This critical step separates the ideal geometry represented on the design layers from the physical manufacturing process. As long as the geometry meets these constraints, the generation of the complex projection lithography masks and subsequent wafer processes will produce suitable silicon features. As such, the logic designer is not required to understand computational lithography and the mask manufacturing house doesn't have to consider the functionality of input geometry.

This paradigm must be broken to some level to integrate photonic devices within the existing ECAD flow. Since the later stages of the manufacturing are designed and optimized for fabricating electronic device geometries, the photonic designer acting as a foundry customer must consider the later data preparation and fabrication steps. This has two major implications: first, care must be taken in designing photonic structures to yield acceptable performance; and, second, the photonic structures must not affect process yield for other standard electronic customers running on the same wafers.

A. Full-Custom Layout of Photonic Structures

We chose to develop our photonic design platform within the dominant full-custom very large scale integration (VLSI) layout tool, Cadence Design Systems' Virtuoso. The first obstacle in this platform is that the vertices of all objects must conform to a discrete grid. Additionally, all relevant shapes must be made of lines at either 0° or 90°, i.e., Manhattan geometry. This is the result of extensive optimization of the subsequent manufacturing processes for traditional electronic designs. This presents a significant challenge to produce the smooth curves required for many photonic devices. Additionally, proper design hierarchy and automation must be enabled within this platform to enable complex integration.

B. Efficient Photonic Device Layout on Design Layers

Instead of relying on manufacturing processes optimized for curves, the resolution of scaled-CMOS lithography can be used to discretize the required shapes on the 1- to 5-nm design grid.
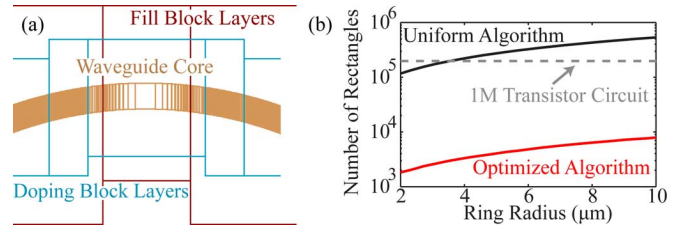

Fig. 3. (a) Optimized discretization algorithm on design layers utilizing different grid resolutions; and (b) the necessary quantity of rectangles to represent a 0.5-μm-wide ring with uniform and optimized algorithms. The rectangle quantity of a 1 million transistor circuit is also shown for reference. Optimized algorithm grids: 1 nm (core), 0.2 μm (doping block), and 0.8 μm (fill block).

Due to the large size of the photonic structures of interest compared to the design grid, an efficient representation of the curves is necessary. Options for building block objects are limited to many-vertex polygons and rectangles. Polygons have the benefit of minimizing the total number of objects in the design database. However, since integration into the existing data preparation and mask manufacturing flow is required, we have chosen to use rectangles as the basic building block to more closely resemble standard electronic geometry databases. Since mask manufacturing is not done in-house by CMOS foundries, the risk of complex polygons generating patterning errors cannot be easily assessed by foundry engineers available to a mask share customer.

Uniform rectangular slicing of a 500-nm width, 10-μm radius ring on a 1-nm grid results in ∼40 000 rectangles per layer. Since 14 layers are required to define the core and surrounding material stack up, the rectangle quantity exceeds 500 000. This would require more than double the number of rectangles, 197 600, present in a >1 million transistor electronic design recently fabricated in a 32-nm process. The compression of rectangle volume, due to the design hierarchy discussed in Section II-C, is required to reduce the computation required for foundry mask pattern generation to a practical length of time [9]. Therefore, a design submitted with many photonic devices represented in such a simple form will not complete the pattern generation flow in an acceptable time.

To minimize data volumes, an efficient rectangle discretization algorithm is then required. The discretization grids are chosen to be as large as possible for each design layer. Waveguide cores such as the polysilicon gate layer and single-crystalline body layer in SOI processes are discretized on the minimum available design grid of 1 or 5 nm, depending on process generation. The other design layers used to ensure the correct material surroundings, such as the doping and fill blocking layers, do not require smooth edges and are discretized coarsely far from the waveguide core. An example of the resulting structure is shown in Fig. 3(a) with the resulting data volume as a function of ring radius plotted in Fig. 3(b).

C. Hierarchical Parameterized Cell Design

In VLSI layout, a careful design hierarchy reduces the data volume of complex integrated circuits. The full chip is comprised of many building block cells replicated and positioned in higher level cells. Importantly, cell layouts can be generated

by scripts based on designer input variables. These parameterized cells, or p-cells, typically describe electronic components such as transistors and resistors.

To integrate the process of photonic layout within Virtuoso, a new hierarchy of photonic device p-cells was developed. At the lowest level, non-Manhattan shapes such as circular arcs, tapers, and sinusoidal offsets are discretized based on geometric input parameters for an arbitrary design layer. Next, basic photonic object cells, e.g., rings and waveguides, place these geometries on the correct layers for finely discretized waveguide cores and coarsely discretized surrounding stackup layers. By isolating the highest data volume geometric objects at the lowest level in the hierarchy, basic building blocks are referenced across the design instead of replicated. This is applied even within an object with symmetry axes such as a ring where a single quarter arc is referenced in four locations. Since Virtuoso creates only a single instance of a p-cell for a given set of input parameters, data volume is greatly reduced to allow manageable compute times.

At the final level of the design hierarchy, the basic object p-cells are placed into standardized layouts with waveguides connecting desired devices to fixed input and output port locations. At this level, the designer can easily place the photonic components into complex networks and test structures through the standard Virtuoso user interface. Real-time user input device parameters and port sizing then allow automatic generation the physical device.

D. Geometry Design Rule Compliance

For the foundry to accept the design for fabrication, successful final automatic verification of design rule compliance is required. Since the typical separation of circuit design from process knowledge has been broken by the inclusion of novel devices, the photonic designer must work closely with the foundry representatives to ensure that rules that misinterpret photonic devices as malformed transistors are waived. More importantly, the photonic designer must ensure that no aspects of the submitted design violates rules, such as minimum geometry feature rules and areal density rules, that would affect the process yield of other customers on the wafers. In our experience, the state-of-the-art lithography used in scaled-CMOS technology allowed all geometric size rules for waveguide core layers to be met with design minimal modifications to desired photonic structures. For surrounding material stackup layers, the coarse discretization, larger than minimum spacing and notch rules, results in zero violations within a basic object p-cell. Locking these grids to a global fixed grid regardless of p-cell location or orientation then eliminates spacing violations generated between p-cells.

III. CONCLUSION

Since the initial integration demonstration, this design platform has enabled the fabrication of photonic devices in two other scaled-CMOS processes from two major semiconductor manufacturers. Due to the modular nature of the Virtuoso environment, process information such as the names of required design layers can be isolated into a single technology file to enable
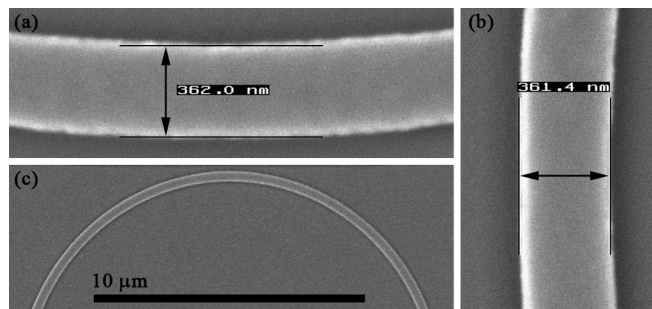

Fig. 4. Scanning electron micrograph (SEM) of a polysilicon ring waveguide core fabricated within a 32-nm bulk-CMOS process showing (a) east–west linewidth, (b) north–south linewidth, and (c) no visible angular distortion.

code reuse. Several design rule violations existed for all of these designs, however, none had any impact on process yield. As a result all violations were waived and successfully completed the standard foundry data preparation flow. Final photonic device geometries such as the curve shown in Fig. 4 show no angular distortion and tight linewidth control from the highly Manhattan geometry optimized lithography and data preparation flow.

Now that the layout of the photonic devices has been integrated within the standard electronic design environment, the advanced design tools developed for complex circuits can be leveraged for silicon photonics. Future work includes automatic verification of connectivity, known as layout-versus-schematic (LVS), by using existing tools such as Mentor Graphics Calibre to extract the photonic p-cells.

REFERENCES

[1] C. Batten *et al.*, "Building manycore processor-to-DRAM networks with monolithic CMOS silicon photonics," *IEEE Micro.*, vol. 29, no. 4, pp. 8–21, Jul./Aug. 2009.
[2] C. W. Holzwarth *et al.*, "High speed analog-to-digital conversion with silicon photonics," *Proc. SPIE*, vol. 7220, p. 72200B, Feb. 2009.
[3] C. Gunn, "CMOS photonics for high-speed interconnects," *IEEE Micro.*, vol. 26, no. 2, pp. 58–66, Mar./Apr. 2006.
[4] Y. Vlasov, W. M. J. Green, and F. Xia, "High-throughput silicon nanophotonic wavelength-insensitive switch for on-chip optical networks," *Nat. Phot.*, vol. 2, pp. 242–246, Mar. 2008.
[5] W. Bogaerts *et al.*, "Nanophotonic waveguides in silicon-on-insulator fabricated with CMOS technology," *J. Lightw. Technol.*, vol. 23, no. 1, pp. 401–412, Jan. 2005.
[6] J. S. Orcutt *et al.*, "Demonstration of an electronic photonic integrated circuit in a commercial scaled bulk CMOS process," in *Proc. Conf. Lasers and Electro-Optics*, San Jose, CA, 2008, Paper CTuBB3.
[7] C. W. Holzwarth *et al.*, "Localized substrate removal technique enabling strong-confinement microphotonics in bulk Si CMOS processes," in *Proc. Conf. Lasers and Electro-Optics*, San Jose, CA, 2008, Paper CThKK5.
[8] J.-T. Kong, "CAD for nanometer silicon design: Challenges and success," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 12, no. 11, pp. 1132–1147, Nov. 2004.
[9] T. Jhaveri *et al.*, "OPC simplification and mask cost reduction using regular design fabrics," *Proc. SPIE*, vol. 7274, p. 727417, Mar. 2009.