Part I

Introduction

Cambridge University Press
978-1-107-08545-9 — Silicon Photonics Design
Lukas Chrostowski, Michael Hochberg
Excerpt
More Information

1 Fabless silicon photonics

1.1 Introduction

We are on the cusp of revolutionary changes in communication and microsystems technology through the marriage of photonics and electronics on a single platform. By marrying large-scale photonic integration with large-scale electronic integration, wholly new types of systems-on-chip will emerge over the next few years.

Electronic-photonic circuits will play a ubiquitous role globally, impacting such areas as high-speed communications for mobile devices (smartphones, tablets), optical communications within computers and within data centres, sensor systems, and medical applications. In particular, we can expect the earliest impacts to emerge in telecommunications, data centers and high-performance computing, with the technology eventually migrating into higher-volume, shorter-reach consumer applications.

In the emerging field of electronics in the 1970s, Lynn Conway at Xerox PARC and Professor Carver Mead at Caltech developed an electronics design methodology, wrote a textbook, taught students how to design electronic integrated circuits, and had their designs fabricated by Intel and HP as multi-project wafers, where several different designs were shared in a single manufacturing run [1]. These efforts led to the foundation of an organization named MOSIS in 1981 that introduced cost sharing of fabrication runs with public access. The inexpensive design-build-test cycle enabled by MOSIS trained, and continues to train, thousands of designers who are responsible for the ubiquity of electronics we see today. MOSIS got started based on commercial processes that were already in production, and opened them up to the design community for prototyping and research purposes.

One of the keys to the long-term success of the microelectronics community, and in particular of the CMOS community, has been this type of access. By making these volume production processes publicly available for research and development at modest cost, anyone with a very modest level of funding is able to do cutting edge, creative work *in a process that can instantly go into large-scale production*. Training student engineers to use the production tools and processes, and then letting them loose to build cutting-edge circuits which can, with modest funding, be translated into fabless IC start-ups, has been the source of countless successful companies. It is hard to over-emphasize the difference between this and the situation in photonics (and most engineering fields), where getting from research into production involves huge barriers.

Silicon photonics is currently at the same early stage of expansion as electronics was in the 1970s, but with a major advantage for chip fabrication: existing silicon foundries

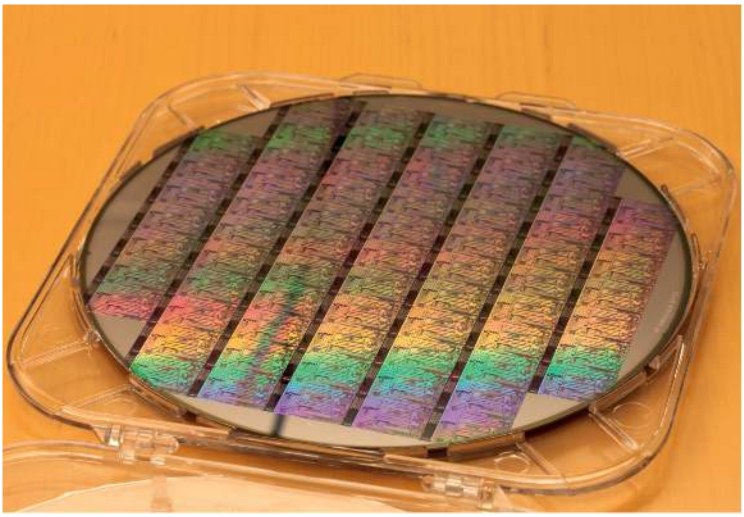

Figure 1.1 Photograph of 8” SOI wafer with various photonic components and circuits. With permission [10].

that produce highly controlled wafers (Figure 1.1) for microelectronics already exist. The microfabrication infrastructure to do silicon photonics already exists, in the microelectronics industry. Several companies are manufacturing silicon photonic chips, e.g., Luxtera’s chips are already used in some high-performance computer clusters [2]. We are presently in an important transition in that academics, students, and industry worldwide now have access to active silicon photonics fabrication, e.g., via the multi-project wafer services offered by ePIXfab [3], IME [4–6], CMC Microsystems [9], and others. However, none of the processes that are currently available to the wider community are production-qualified; they are all prototyping and R&D processes, which support only very limited volumes. The inability to leverage pre-existing commercial processes was a significant impediment to the success of OpSIS, which recently shut down, since commercial users were rightly unwilling to rely on non-production-proven processes for product development. Funding from the research users and funders was not sufficient to keep the effort going, and would not allow the development of processes suitable for commercial use, which was the logical next step for the program.

We are fortunate that silicon allows us to perform all of the key optical functions at a reasonably competitive performance level, as shown in Figure 1.2, with the exception of a laser. There is a lot of recent and very interesting work going on related to growing quantum dots and germanium for lasers monolithically integrated in silicon [7]. And with the bonding technologies that have been inherited from the electronics industry, it is possible to bond lasers at relatively low cost, either through front-end integration similar to the approaches of Intel and Aurrion, or through die-scale bonding of finished laser chips [8]. These approaches are still in development, but it’s clear that there are several practical approaches to making cheap light sources at various levels of integration with the silicon platform. It remains to be seen which approaches will be the most successful.

978-1-107-08545-9 — Silicon Photonics Design
Lukas Chrostowski, Michael Hochberg
Excerpt
More Information

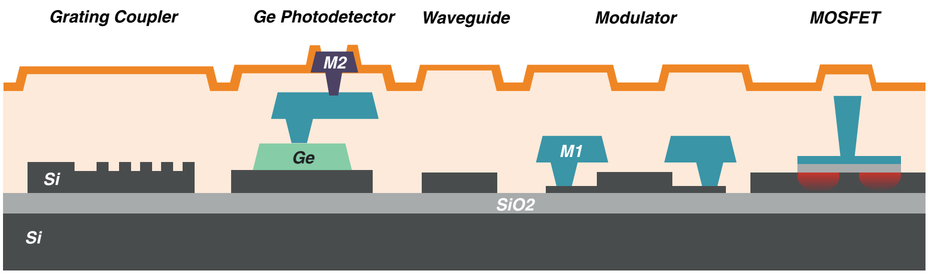

Figure 1.2 Typical process stack representing a silicon photonics platform with grating couplers, germanium photodetectors, waveguides, modulators and MOSFETs, on a silicon-on-insulator wafer. Note that most of the advanced work being done in silicon photonics today does not utilize monolithic integration with transistors, but instead leverages 3d or 2.5d integration. With permission [10].

1.2 Silicon photonics: the next fabless semiconductor industry[1]

The same foundries and processes that were developed to build transistors are being re-purposed to build chips that can generate, detect, modulate, and otherwise manipulate light. This is somewhat counter-intuitive, since the electronics industry spends billions of dollars to develop tools, processes and facilities that lend themselves to building the very best transistors, without any thought about how to make these processes compatible with photonics (with the exception of the processes designed to make devices such as CMOS and CCD camera chips). How are we so lucky that these capabilities can be directly reused for photonics?

In reality, they can't be directly reused. Every attempt to directly integrate photonic functionality into CMOS or bipolar silicon wafers, without making any process changes, has yielded poorly-performing devices. Electronics processes are designed for making electronics; it stands to reason they cannot be used for competitive photonics products. And even if they could, it wouldn't make economic sense; silicon photonics chips require relatively primitive processing (90 nm type capabilities) compared to advanced microelectronics chips (16 nm). Using the tools for truly advanced microelectronics to try to build photonics is a mistake, and would be impossible to justify from a performance or economic standpoint.

There is no reason to expect that the integration flows used to build electronic circuits would be in any way compatible with making components that manipulate light. But, over the last decade, it has emerged that silicon is actually a fantastic material system for building photonic devices, as well as electronic ones. And, even more surprisingly, the silicon photonics community has developed process flows that permit the re-use of CMOS fabrication infrastructure to build complex photonic circuits, where information is transferred seamlessly from the electronic to the optical domain and back again.

[1] A version of this chapter has been published in the *IEEE Journal of Solid-State Circuits* [10]: Michael Hochberg, Nicholas C. Harris, Ran Ding, Yi Zhang, Ari Novack, Choice Xuan, Tom Baehr-Jones. "Silicon photonics: The next fabless semiconductor industry", *IEEE Journal of Solid-State Circuits*, Vol. 5, No. 1, pp. 48–58, March 2013.

While the fully-integrated processes used to make transistors are not reused, modular process steps can be rearranged and reused, with distinct process flows being developed to build silicon photonics. This is not a trivial endeavor, but several organizations have shown that it is possible.

What has emerged is a vibrant community of companies and academics using the materials and techniques that have been developed over the past 50 years in the silicon microelectronics industry, and repurposing them to build photonic devices and circuits. What is particularly compelling about this work is that many of the efforts don't just make use of the same kinds of equipment in separate facilities, but actually use the exact same tools and facilities where CMOS transistors are routinely built. The constraints associated with working in such facilities are significant: materials that are not proven to be compatible with the CMOS processes are banned, and both processes and circuits have to be designed in such a way that processing them will not harm or contaminate the tools. The cost of mask sets and process development in the more advanced CMOS-compatible manufacturing facilities can also be very high [11] in advanced processes. But if the billion-dollar scale investments that go into building modern CMOS facilities can be directly leveraged to build silicon photonic systems on chip, it means that there is an immediate and a rapid path to commercialization and large-scale production is available.

1.2.1 Historical context – Photonics

Up to the present day, there has been very little in the way of opportunities for fabless photonics companies.

One of the key problems in photonics, historically, is that processes have been highly specialized for the particular application, utilizing different materials. With individual devices separately packaged and connected together by fibres, it is not unusual to see communications systems which incorporate chips made in half a dozen different material systems: RF CMOS or bipolar processes for the high-bandwidth electronics (e.g., serializers and deserializers), FPGA's or highly scaled CMOS for the digital parts (e.g., control circuits), diffused waveguides on glass for optical multiplexers (e.g., arrayed waveguide gratings) and passives, lithium niobate for modulators, indium phosphide for lasers, germanium for photodetectors, and MEMS-based switches, for instance. Each of these devices is made in a process that is fundamentally and irreconcilably incompatible with the ones used to make the other components. Each material system is chosen to provide ultimate performance for a single type of device. This means, in most cases, that the photonic components are produced in specialty fabrication facilities, in very low volumes. This results in high-cost components, since very few photonic devices are truly high-volume on the scale of the electronics industry. The only things that come close are VCSELs (which are made on a wafer scale, but are used as discrete devices), and components for PON networks (again, wafer-scale fabrication of directly modulated lasers, but they are used as singulated discrete components).

While discrete photonic devices can be interfaced to one another with standard optical fibres and connectors, a large fraction of the final device cost and yield loss

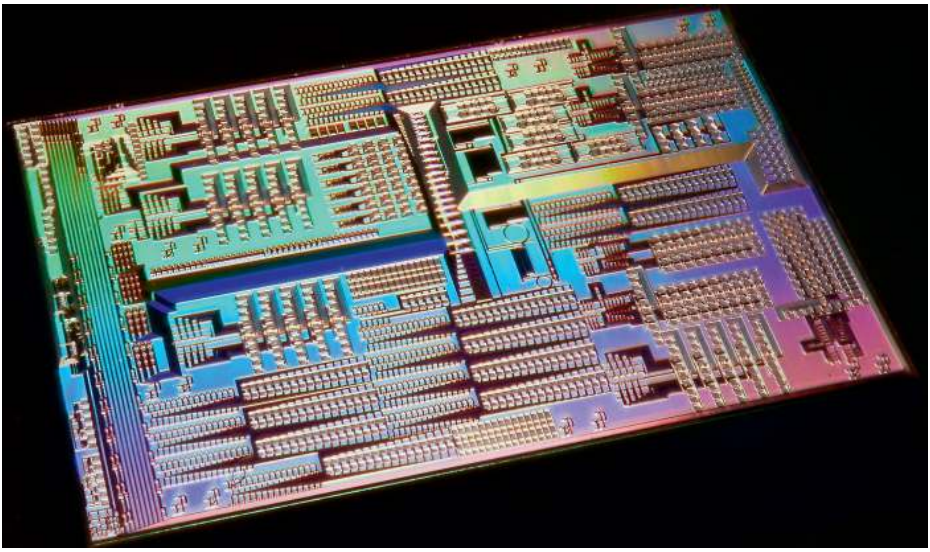

Figure 1.3 Photograph of SOI photonic chip fabricated at IME A*STAR.

emerges from the photonic packaging processes, which generally require 5- and 6-axis alignments with submicron accuracy, and from the packages themselves, which are often hermetically sealed and are sometimes quite literally gold-coated. Again, this contributes to the high cost of photonic components and systems.

The great promise of silicon photonics lies in integrating multiple functions into a single package, and manufacturing most or all of them using the same fabrication facilities that are used to build advanced microelectronics, as part of a single chip or chip stack (see Figures 1.3-1.4). Doing so will radically drive down the cost of moving data through fibres, and will create the opportunity for a variety of fundamentally new applications of photonics, where high-complexity systems can be built at very modest costs.

1.3 Applications

There are a number of applications that are emerging for complex silicon photonic systems, the most common being data communication. This includes high-bandwidth digital communications for short-reach applications, complex modulation schemes and coherent communications for long-reach applications, and so on.

Beyond data communications, there are a huge number of new applications being explored in both the commercial and academic worlds for this technology. These include: nano-optomechanics and condensed matter physics [12], biosensing [13, 14], nonlinear optics [15], LIDAR systems [16], optical gyroscopes [17, 18], radio frequency integrated optoelectronics [19, 20], integrated radio transceivers [21], coherent communications [22], novel light sources [23, 24], laser noise reduction [25], gas sensors [26], very long wavelength integrated photonics [27], high-speed and microwave signal

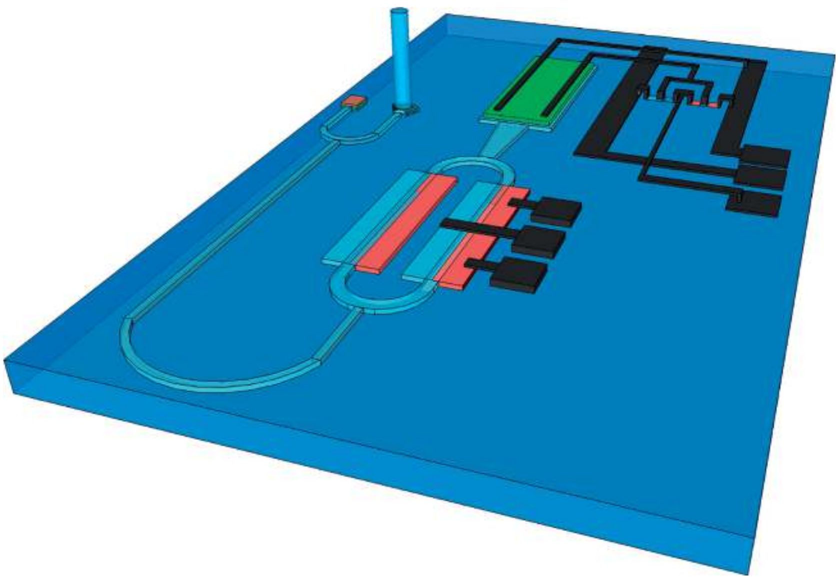

Figure 1.4 Conceptualization of a CMOS/photonic circuit. Light is coupled via on-chip laser or vertical fibre into a grating coupler. The light is then modulated, transduced by a photodetector, and inverted via a CMOS inverting circuit. Silicon photonic-electronic circuits can support systems with hundreds or thousands of such components today. Ref. [10]

processing [20], and many more. Areas of particular promise include biosensing, imaging, LIDAR, inertial sensing, hybrid photonic-RFICs, and signal processing.

1.3.1 Data communication

There are several categories of photonic devices where the silicon photonics components are competitive with best-in-class devices, discussed below. But what we regard as perhaps the most transformative work being done in the optical communications field is concerned with creating integrated platforms with modulators, detectors, waveguides and other components on the same chip, all talking to one another. In some cases, transistors are also included in these platforms, allowing amplifiers, serialization, and feedback to all be integrated onto the same chip. Because of the expense of developing such processes, this effort has largely been led by commercial players, mainly aimed at applications in point-to-point data communications. And because of the expense of developing transistor manufacturing processes, the emerging consensus in the field is that integrating electronics through bonding, either at the wafer or a the die level, makes the most sense for the foreseeable future from both a performance and cost perspective.

There is a lot of obvious value in being able to make chips that can perform computations with electronics, and transmit data optically. The vast majority of the early applications of silicon photonics are in digital data communications. This is driven by the fundamental differences in physics between electrons, which are Fermions, and photons, which are Bosons. Electrons are great for computation, because two of them cannot be in the same place at the same time. That means that they interact very strongly with each other. As a result, it is possible to build massively nonlinear switching devices – transistors – using electronics.

Photons have a different set of properties: many photons can all be in the same place at the same time, and except under very special circumstances they don't interact with each other. That's why it is possible [28] to transmit literally terabits of data per second through a single optical fibre: this is not done by creating a single terabit-bandwidth stream of data. Instead, typical high-bandwidth fibre optic transmission systems today make use of a variety of techniques to take a large number of streams of electrical data, at 10, 28 and 40 Gbits/second speeds, and multiplex them together onto single fibres. One such technique, commonly used in telecommunications, is wavelength-division multiplexing [29], where each colour of light is modulated separately and the various colours are all combined onto a single fibre. Other techniques, where coherence is exploited in order to encode information into phase, amplitude and polarization separately, have become common in long-haul and metro telecommunications as well [30]. In fact, these different approaches can be combined, with each wavelength representing a separate stream of data, and also being modulated in phase, amplitude and polarization, in order to transmit many bits per symbol. Thus it is possible to build systems that transmit terabits of data through a single fibre, without any electronics operating faster than 28 or 40 Gbits/second.

These techniques are critical, because fibre optic data transmission over long distances is limited by the expense of laying the fibres: Obtaining right-of-way and physically laying the conduits is expensive on land, and frightfully expensive across oceans [31]. In such systems, even if it requires millions of dollars of equipment at the endpoints to make efficient use of existing fibres, that is economically sensible, since the cost of laying more fibre is prohibitive. What we've seen over the past 50 years is that, because of the very low losses of fibres and the availability of erbium-doped fibre amplifiers [32] (which amplify all of the different optical wavelengths in a given window, around 1.5 μm) fibre optics first became the dominant technology for transmitting data over long distances, between cities. Over time, it came to dominate metro links, and is now dominating the relatively short distances within the data center. In many parts of the world, fiber-to-the-home is the dominant access paradigm, though this has not proven to be true in the United States, where it competes with DSL and other technologies. With the constant demand for increasing bandwidth [33], demand for ever-more-efficient ways of pushing data through fibres has grown steadily.

The broad trend in the data communications market is that as distances get shorter, the price per part drops precipitously, while the volumes go up. Unsurprisingly, silicon photonics commercialization efforts have focused a lot of effort onto the higher-volume, shorter-reach applications, aiming at data centres and high performance computing. Future applications will include board-to-board, short reach connections on the scale of USB, and perhaps eventually CPU core-to-core communications [34], though the case for on-chip core-to-core applications remains quite speculative.

Though not yet on the scale of the CMOS industry, silicon photonics is beginning to be a significant industry in its own right. The first commercial products integrating chip-scale electronics and photonics have recently hit the market [35], and Intel has announced its intention to standardize a format for optical data communication for personal computing [REF-INTEL]. Luxtera announced the sale of their millionth silicon photonic data channel [37], and they are now selling a 100 Gbits/second-class optically active cable (4x28G) [38], fabricated using the Freescale CMOS foundry in Texas [39]. Numerous startups and established companies (Kotura, Luxtera, Oracle, Genalyte, Lightwire/Cisco, APIC, Skorpios, TeraXion, and many others) are actively developing silicon photonic products. Many of the top American defense companies have programs in this area. Moreover, many of the major semiconductor players now have active programs in silicon photonics. Intel, Samsung, IBM, ST and many others have publicly announced activity in this area. Although there have been some over-blown predictions by market research firms [40] claiming that the field will be generating $2B/year in revenue by 2015, it does seem likely that revenues around $1B will be achieved before 2020, as predicted by the authors [11].

1.4 Technical challenges and the state of the art

1.4.1 Waveguides and passive components

There are a wide variety of waveguide geometries that have been developed in silicon-compatible systems; nearly any transparent material with a higher refractive index than glass can be deposited on top of an oxidized silicon substrate and turned into a waveguide. For purposes of CMOS process compatibility, however, the community has converged on a few classes of geometries. The most common are high-confinement waveguides made out of the active device layer of an SOI wafer, etched either fully to the bottom oxide layer or partially etched with a timed process stop [41, 42] (Figure 1.2 and 3.4). It took several years of work to reduce the losses of these sub-micron waveguides to acceptable levels, since the strong interaction of the optical fields with the sidewalls can lead to substantial losses, driven by roughness [43]. Propagation loss can be reduced either by process optimization to smooth the sidewall [43] or by waveguide geometry optimization to reduce modal field strength at the sidewall [44]. Typical losses for high-confinement guides are in the 2 dB/cm range today for cutting-edge processes [45]. Low loss multimode straight waveguides in combination with tight single mode waveguide bends turn out to be an optimal choice for routing, achieving 0.026 dB/cm [46]. Other key passive components such as grating couplers [47] (Figures 1.2, 1.5, and 1.6), distributed Bragg gratings [48], waveguide crossings [49], and arrayed waveguide gratings (AWG) [50] have all been demonstrated, in each case with very low losses. More recently, CMOS compatible waveguides that can be formed into the dielectric back-end process have become available, made out of silicon nitride. With dedicated processing, the losses of these waveguides are extraordinarily low (< 0.1 dB/m), though the compatibility of such processes with front-end active devices is an open issue, given the requirement for high temperature growth [51]. It should be noted that considerable work has been done on low-confinement silicon waveguides [52, 53].

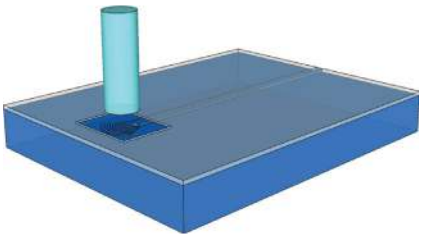

(a) Vertical coupling of optical fibre to on-chip grating coupler

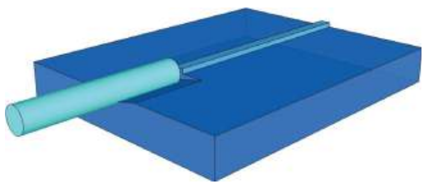

(b) V-groove coupling of optical fibre to on-chip waveguide

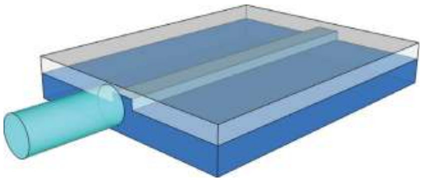

(c) Edge-coupled optical fibre to on-chip waveguide

Figure 1.5 Light coupling techniques. With permission [10].

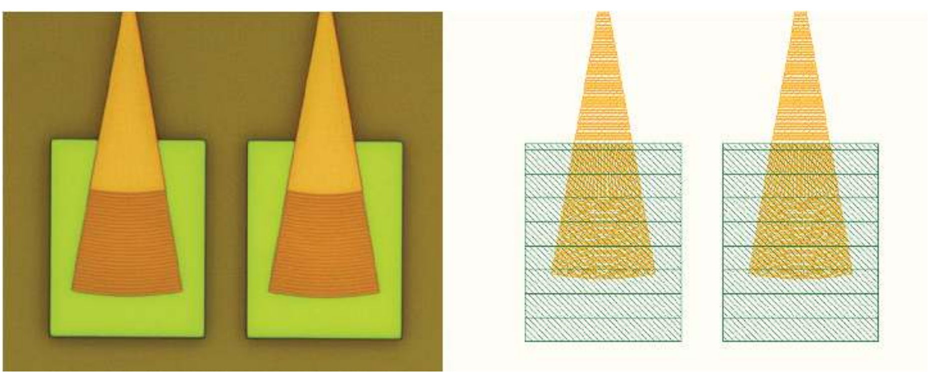

Figure 1.6 Optical micrograph and layout of an optical grating coupler fabricated at IME A*STAR. With permission [10].

There is considerable difficulty in creating compatible high-speed modulators and detectors in these platforms, making them less promising for large-scale integration.

One of the challenges with silicon photonics is coupling light between the chip and optical fibres, and achieving this is in a packaging method that is cost effective [54]. This is typically done using edge or grating couplers, as in Figure 1.5, and described in Chapter 5. Both approaches have demonstrated below 1 dB loss per interface [54–56]. Dealing with polarization is also a challenge, since silicon photonic waveguides are, by default, highly birefringent, namely the optical propagation constants in the waveguides are different for the two different polarizations. The common approach is to build circuits using a single polarization, and to duplicate them when both polarizations are