Teaching silicon new tricks

BAHRAM JALALI

is at the University of California, Los Angeles, Department of Electrical Engineering, Los Angeles, California 90095-1594, USA.
e-mail: Jalali@ucla.edu

Arguably the most important element for the electronics industry, silicon is now being given a new lease of life in the world of photonics.

Empowering silicon with optical functions, such as the ability to emit, guide and modulate light, could be the key to creating short-distance ultrafast optical interconnects that overcome one of the most formidable hurdles in scaling the speed of computing[1]. An ultimate aim is the realization of silicon chips that communicate internally, or with other chips, using photons and optical waveguides and thus overcome the bandwidth limitations imposed by metallic interconnects. Yet another opportunity lies in optical sensors with on-chip communication circuitry that can form nodes of intelligent sensor networks used for environmental and health monitoring.

Guided by such visions and propelled by pioneering research conducted in the 1980s and 1990s, silicon photonics has enjoyed spectacular progress in the past five years. The critical size of photonic devices has been scaled to the 100-nm regime, commensurate with that of electronic devices and at the limit imposed by the optical wavelength. Optical amplification and lasing, once considered forbidden in silicon, have been achieved[2]. High-speed and efficient electrical-to-optical and optical-to-electrical conversion is being performed by production-worthy silicon devices. Silicon's nonlinear optical properties, enhanced by tight optical confinement in Si–SiO$_2$ structures, are producing wavelength generation and conversion — central functions in multiwavelength communications and signal processing. Many of these developments are documented in recently published review articles[3,4]. Not meant as a comprehensive review, the present commentary provides a sample of recent achievements and discusses future challenges for this exciting field.

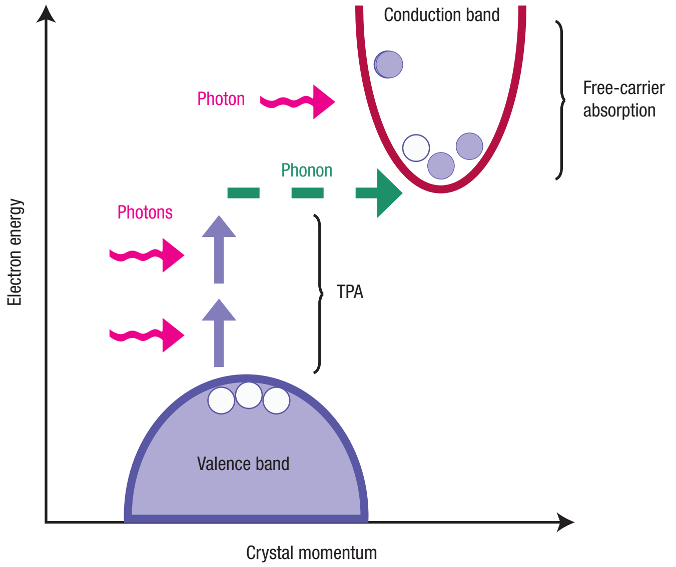

Figure 1 Two-photon absorption generates free carriers that in turn cause a significant amount of absorption. The loss of photons is the main problem facing a new class of silicon devices that perform optical amplification, lasing and wavelength conversion.

POTENTIAL IMPACT

The greatest impact of silicon photonics will most probably be in data-communication applications. Consequently, much research has focused on producing a source of photons, an electro–optic modulator and a photodetector. Optical amplifiers and wavelength-manipulation devices (filters, switches and converters) are also needed to compensate for path losses and to take full advantage of the optical bandwidth available.

It has been known since the early 1990s that one can create high-quality (low-loss and single-mode) optical waveguides on silicon-on-insulator (SOI) wafers[5], the preferred platform for CMOS and the substrate used for the IBM PowerPC and the AMD Opteron microprocessors.

However, coaxing silicon to emit and amplify light, and to change its wavelength, is far more difficult because of hard limits imposed by nature. Silicon is an inefficient light emitter because of its indirect bandgap — the momentum of the conduction and valence bands are displaced (Fig. 1). The vast majority (but not all) of electron–hole pairs, created by

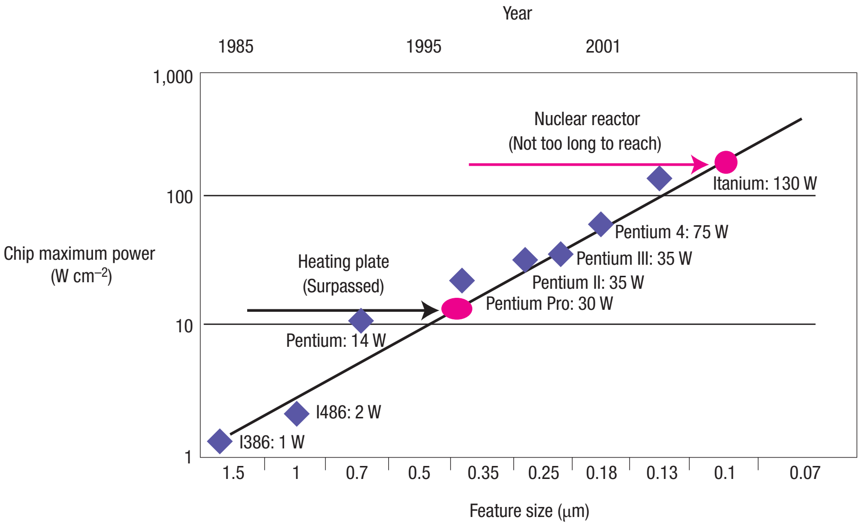

Figure 2 Owing to the rapid increase in transistor count, in accordance with Moore's Law, the power density on a silicon chip has already exceeded 100 W cm^{-2}. For photonic devices to find their way onto a silicon VLSI chip, they must be low power, and more importantly, they must be able to operate reliably at elevated temperatures as high as 90 °C. Figure courtesy of Fred Pollack.

electrical or optical injection, lose their energy to heat before they can create a photon. To make matters worse, injected free carriers (electrons and holes) will also absorb photons (Fig. 1). Any light that is created, or enters from outside, will not be amplified because the rate of stimulated emission is far below the rate of absorption by free carriers[4]. Amplification is the prerequisite to lasing, which is why it was so difficult to make silicon lase using conventional approaches. A second fundamental problem with silicon is its lack of a static dipole moment, a consequence of its centrosymmetric crystal structure. This means that the linear electro–optic (Pockel) effect — that wonderful phenomenon that makes $LiNbO_3$ and III–V semiconductors good electro–optic materials — is absent in silicon.

Today, the only practical way to encode data onto light is by modulating its absorption by carrier (electrons and holes) injection or depletion[6]. Silicon is also unable to detect signals at the communication wavelengths of 1,300 nm and 1,550 nm because such photon energies are less than its bandgap. This problem has been overcome by taking advantage of the small bandgap of germanium that has been grown on silicon[7]. The approach has been very successful with these detectors representing a mature segment of silicon photonics.

RECENT PROGRESS

It is not surprising that much of the current activity is focused on finding ways around the physical limitations of silicon. A conclusive demonstration of optical amplification and the first silicon laser were demonstrated in 2003–2004, by exploiting the Raman scattering, a phenomenon that describes interaction of light with atomic vibrations of the crystal. Today, the question is no longer whether silicon can perform these functions, but rather how well. A limitation of Raman amplification is the narrow gain bandwidth. Unless a broadband or multiwavelength pump is used, the intrinsic 100-GHz Raman gain spectrum is only sufficient for amplifying two 10-Gbit s^{-1} data channels spaced by the industry standard 50 GHz optical frequency spacing. Motivated to solve this problem, Alex Gaeta, Michal Lipson and co-workers at Cornell University recently demonstrated broadband optical amplification — a pulsed gain of 10 dB over an optical bandwidth of nearly 3 THz — by using the Kerr effect and a specially designed waveguide instead[8]. Although second-order nonlinearities do not exist, third-order processes, such as the Kerr effect, are alive and well in silicon.

To avoid the generation of electrons and prevent free-carrier absorption, Raman and Kerr devices use infrared light (1,500 nm), so that the photon energy is less than the bandgap. However, because of the high intensities involved (greater than 10 MW cm^{-2}), two-photon absorption (TPA) occurs (Fig. 1). The resulting free-carrier absorption is the main problem in such devices. To avoid them, devices typically operate in the pulsed mode where electron accumulation does not occur. Because the pulse repetition period must then be larger than the minority carrier lifetime (1–10 ns in SOI waveguides), the utility of this technique is limited to low data rates. Hence, dealing with the TPA-induced free-carrier absorption remains a central focus of silicon-photonics research. Recently, it was demonstrated that by taking advantage of the two-photon photovoltaic effect inherent in these devices, carrier sweep-out can be achieved with negative electrical-power dissipation[4]. In another promising direction, researchers at the Chinese University of Hong Kong (CUHK) and at NTT in Japan have independently shown that ion bombardment can be used to reduce the accumulation of generated carriers by lowering the minority-carrier lifetime[9,10].

In addition, Brinkmeyer, Renner and Strauss at the Technical University of Hamburg have shown advanced optical amplifier designs in which the deleterious effect of TPA can be minimized by using a double-cladding waveguide[11]. By uniformly distributing the pump power throughout the amplifier length, they avoid excessive carrier generation in the input section. However, it should be noted that TPA is not always detrimental and it can, in fact, be exploited to perform useful functions. For example, a TPA photodetector can be used for monitoring the optical power in an integrated multiwavelength channel equalizer[9]. It can also be used as a pulse compressor in an external-cavity mode-locked laser[12].

As for silicon light sources, significant excitement was created in 2000 by the report of internal optical gain from silicon nanocrystals embedded in a silicon oxide host. Although obtaining net gain has been frustrated by carrier absorption and scattering loss caused by nanocrystals, and the approach has not yet produced a laser, electroluminescent diodes that produce spontaneous emission continue to improve[13]. Among recent ideas are the cavity enhancement of emission, and the use of silicon nitride as the host for nanocrystals[14] (instead of silicon oxide). By lowering the tunnelling barrier and hence, the operating voltage, the latter may enhance the device reliability and solve a lingering issue with such devices.

Japanese researchers at Hitachi have also recently reported an electroluminescent device with a 9-nm SOI active region[15] and a CMOS compatible structure. Dubbed a light-emitting transistor (LET), it can emit and detect light. Progress is also being made in hybrid silicon/III–V technologies with the preferred approach being the bonding of an indium phosphide gain region onto a passive silicon microcavity. Using this approach, a joint Intel–UCSB team reported an electrically pumped hybrid laser in 2006 (ref. 16).

FUTURE TRENDS AND CHALLENGES

For photonics to be truly compatible with very-large-scale-integrated (VLSI) chips, it must be compliant with the two basic requirements of the semiconductor industry: power efficiency and real-estate efficiency. Power dissipation is the most pressing problem in the semiconductor industry, to the extent that it finally forced the microprocessor manufacturers to abandon higher clock speeds in favour of multicore architectures[17]. Described by Moore's Law, transistor count has been increasing at an astonishing pace, with the latest microprocessors introduced in 2006 (for example, the Intel Itanium II) boasting 1.7 billion transistors. The unintended consequence is that power density now exceeds 100 W cm^{-2}, an order of magnitude higher than a typical hot plate (Fig. 2). This poses severe challenges for photonic devices, for they must be able to operate reliably at substrate temperatures as high as 90 °C. Diode lasers, such as those used in fibre-optic communication, are packaged with a thermoelectric cooler because their performance severely degrades at high temperatures. It remains to be seen whether these types of lasers, regardless of what material they are made of (silicon, III–V or hybrid), can efficiently and reliably operate on a VLSI chip. Today's use of an off-chip 'optical power supply' is compatible with the thermal requirements of VLSI chips and is, therefore, an intelligent solution. Other photonic devices, such as the optical modulator, detector and amplifier, must also be able to operate at elevated temperatures. High-temperature compliance notwithstanding, photonic devices cannot dissipate large amounts of power on their own without exacerbating chip-heating problems. The time has come for researchers to begin studying the thermal behaviour of their devices in order to determine their true CMOS compatibility. Power dissipation also plays a pivotal role in the trade off between optical and copper-based interconnects. To be sure, the bandwidth of copper interconnects continues to improve by using equalization techniques, albeit, at the price of higher power consumption.

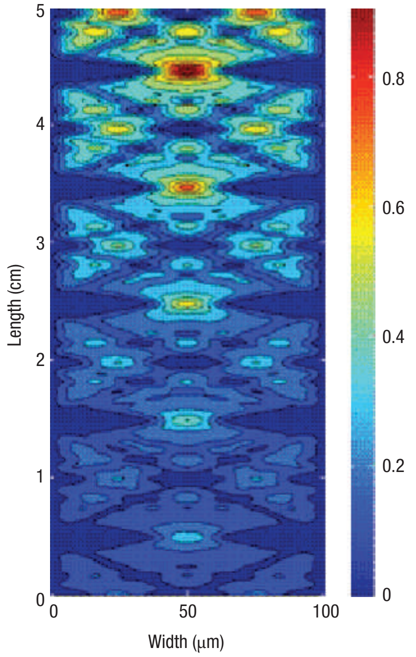

Figure 3 A silicon image amplifier as described in ref. 2. The device combines Raman amplification and self imaging in a multimode waveguide. Used as an image pre-amplifier, its intended application is to enhance the sensitivity of optoelectronic image sensors by elevating the signal above the sensor's thermal-noise limit.

Moore's Law is based on the observation that as transistors get smaller, they also become cheaper and faster[17]. They become cheaper because smaller transistors occupy less chip area. Photonic devices must be real-estate efficient if they are to be economically compatible with silicon manufacturing. Great progress has been made in reducing the size of silicon photonic devices; however, the success is exposing new challenges. For example, Graham Reed and colleagues from the University of Surrey in the UK have correctly argued that reductions in the cross-section of a silicon waveguide come at the expense of polarization sensitivity[18]. Apparently, polarization-insensitive waveguide devices, such as the input/output waveguides needed to interface with optical fibres, have a lower transverse-dimensional limit of approximately a micrometre — giant structures in the nanoscale world of CMOS. For internal interconnects, submicrometre waveguides are still desired to ensure efficient use of the chip area. Another challenge exposed by scaling is the increase in waveguide loss caused by scattering of light at the surfaces, a phenomenon that diminishes the power efficiency of optical inteconnects.

Data communication and interconnects aside, there are many other applications for silicon photonics. Silicon is already used in mirrors for mid-infrared lasers and has good nonlinear optical properties at these wavelengths[4]. These, plus its large thermal conductivity and high optical-damage threshold renders silicon an attractive platform for mid-infrared photonics, with applications in medicine, biochemical sensing and defence. An example is the recently proposed silicon image amplifier that takes advantage of Raman amplification in the presence of the Talbot self-imaging effect (Fig. 3; ref. 2). Such a device, which is being developed by a UCLA–Northrop Grumman team, is designed to improve the sensitivity of laser-based remote-imaging systems. There are many more unconventional applications.

In summary, silicon photonics is making stunning progress and stands to have a bright future, as long as the research community recognizes the real challenges that remain and maintains an open mind with respect to its applications.

References

1. Gibbs, W. W. *Sci. Am.* 81–87 (November 2004).
2. Jalali, B. *Sci. Am.* 58–65 (February 2007).
3. Soref, R. A. *J. Sel. Top. Quant. Electron.* **12,** 1678–1687 (2006).
4. Jalali, B. & Fathpour, S. *J. Lightwave Technol.* **24,** 4600–4615 (2006).
5. Soref, R. A., Schmidtchen, J. & Petermann, K. *J. Quant. Electron.* **27,** 1971–1974 (1991).
6. Soref, R. A. & Lorenzo, J. *J. Quant. Electron.* **22,** 873–879 (1986).
7. Huang, Z. *et al. J. Sel. Top. Quant. Electron.* **12,** 1450–1454 (2006).
8. Foster, M. A. *et al. Nature* **441,** 960–963 (2006).
9. Liu, Y., Chow, C. W., Cheung, W. Y. & Tsang, H. K. *Photon. Technol. Lett.* **18,** 1882–1884 (2006).
10. Tanabe, T. *et al. Appl. Phys. Lett.* **90,** 031115 (2007).
11. Krause, M. *et al.* in *Third Int. Conf. on Group IV Photonics, Ottawa, Canada* 61–63 (2006).
12. Tien, E.-K., Yuksek, N. S., Qian, F. & Boyraz, O. *Opt. Express* (in the press).
13. Iacona, F. *et al. J. Sel. Top. Quant. Electron.* **12,** 1596–1606 (2006).
14. Sung, G. Y. *et al. J. Sel. Top. Quant. Electron.* **12,** 1545–1555 (2006).
15. Saito, S. *et al. Appl. Phys. Lett.* **89,** 163504 (2006).
16. Fang, A. W. *et al. Opt. Express* **14,** 9203–9210 (2006).
17. Muller, D. A. *Nature Mater.* **4,** 645–647 (2004).
18. Reed, G. T. *et al. J. Sel. Top. Quant. Electron.* **12,** 11335–11344 (2006).

Acknowledgements

The author would like to thank DARPA for supporting his work on silicon photonics. He is also indebted to Sasan Fathpour and Prakash Koonath for their review of the manuscript.