

Chapter 1

Silicon Photonics

Silicon photonics can be defined as the utilization of silicon-based materials for the generation (electrical-to optical conversion), guide, control, and detect (optical-to electrical conversion) of light to communicate information over distance. The most advanced extension of this concept is to have a comprehensive set of optical and electronic functions available to the designer as monolithically integrated building blocks upon a single silicon substrate.

1.1 Why Silicon Photonics?

Fiber-optic communication is the process of transporting data at high speeds on a glass fiber using light. Fiber optic communication is well established today due to the great capacity and reliability it provides. However, the technology has suffered from a reputation as an expensive solution. This view is based in large part on the high cost of the hardware components. These components are typically fabricated using exotic materials that are expensive to manufacture. In addition, these components tend to be specialized and require complex steps to assemble and package. These limitations prompted Intel to research the construction of fiber-optic components from other materials, such as silicon. The vision of silicon photonics arose from the research performed in this area. Its overarching goal is to develop high-volume, low-cost optical components using standard CMOS processing – the same manufacturing process used for microprocessors and semiconductor devices. Silicon presents a unique

material for this research because the techniques for processing it are well understood and it demonstrates certain desirable behaviors. For example, while silicon is opaque in the visible spectrum, it is transparent at the infrared wavelengths used in optical transmission, hence it can guide light. Moreover, manufacturing silicon components in high volume to the specifications needed by optical communications is comparatively inexpensive. Silicon's key drawback is that it can not emit laser light, and so the lasers that drive optical communications have been made of more exotic materials such as indium phosphide. However, silicon can be used to manipulate the light emitted by inexpensive lasers so as to provide light that has characteristics similar to more-expensive devices. This is just one way in which silicon can lower the cost of photonics.

1.2 Silicon Photonics Research

With the goal of developing photonic components that are factory-compatible with silicon microelectronic integrated circuits and optical integrated circuits, silicon photonics has been the subject of intense research activity in both industry and academia. Silicon is an excellent material for confining and manipulating light at the sub micrometer scale, and possesses the added advantage of leveraging the enormous manufacturing infrastructure developed by the silicon microelectronics industry. Silicon optoelectronic integrated devices have the potential to be miniaturized and mass-produced at affordable cost for many applications and markets, including telecommunications, optical interconnects, medical screening, and biological and chemical sensing. Recent developments in diverse areas, such as light sources, modulators, switches, detectors, photonic crystals,

waveguide structures, resonators, sensors, and various subsystems, indicate that Si photonics is an extremely active, and now, firmly established research field.

The aim of this seminar is to reveal some of the remarkable recent progress in silicon photonics from academic and industrial viewpoints and thereby point to future trends in this rapidly evolving field.

1.2.1. Six Building Blocks : Intel's SP Research

Intel's silicon photonics research is an end-to-end effort to build integrated photonic devices in silicon for communication and other applications. In order to “siliconize” photonics, there are six main areas or building blocks for research and investigation.

- An inexpensive light source.
- Devices that route, split, and direct light on the silicon chip.
- A modulator to encode or modulate data into the optical signal.
- A photodetector to convert the optical signal back into electrical bits.
- Low-cost, high-volume assembly methods.
- Supporting electronics for intelligence and photonics control.

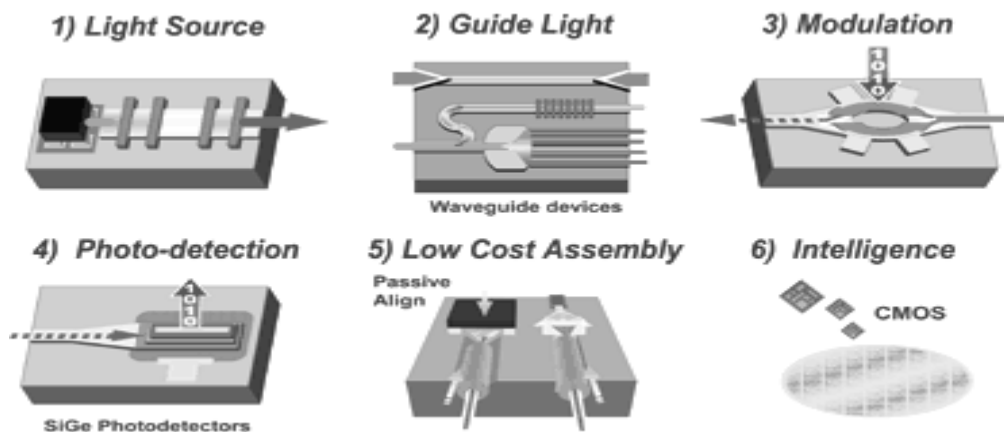


Fig 1.1: Building Blocks of Silicon Photonics

Chapter 2

Light Source In Silicon

Fact: *Silicon is an inefficient light emitter because of a fundamental limitation called an indirect band-gap. An indirect band-gap prevents the atoms in silicon from emitting photons in large numbers when an electrical charge is applied. Instead, silicon emits heat.*

2.1 The Silicon Laser Challenge

A key challenge facing the silicon photonics research is a fundamental physical limitation of silicon: namely, silicon cannot efficiently emit light. While it is capable of routing, modulating, and detecting light, silicon has needed an external light source to provide the initial light.

These external light sources are generally discrete lasers and require careful alignment to the silicon waveguides. The problem is that accurate alignment is difficult and expensive to achieve. Even submicron misalignment of the laser to the silicon waveguide can render the resulting photonic device useless.

A long-standing quest in silicon photonics has been the creation of a laser source that can be manufactured directly on the silicon photonic chip, in high volume, and whose emitted light is automatically aligned with the silicon waveguide.

2.2 The Raman Effect

The term “Laser” is an acronym for *Light Amplification through Stimulated Emission of Radiation*. The stimulated emission is created by changing the state of electrons – the subatomic particles that make up electricity. As their state changes, they release a photon, which is the particle that composes light. This generation of photons can be stimulated in many materials, but not silicon due to its material properties. However, an alternate process called the Raman Effect can be used to amplify light in silicon and other materials, such as glass fiber. Intel has achieved a research breakthrough by creating an optical device based on the Raman Effect, enabling silicon to be used for the first time to amplify signals and create continuous beams of laser light. This breakthrough opens up new possibilities for making optical devices in silicon.

The Raman Effect is widely used today to make amplifiers and lasers in glass fiber. These devices are built by directing a laser beam – known as the pump beam – into a fiber. As the light enters, the photons collide with vibrating atoms in the material and, through the Raman Effect energy is transferred to photons of longer wavelengths. If a data beam is applied at the appropriate wavelength, it will pick up additional photons. After traveling several kilometers in the fiber, the beam acquires enough energy to cause a significant amplification of the data signal (Figure 2.1a). By reflecting light back and forth through the fiber, the repeated action of the Raman Effect can produce a pure laser beam (see sidebar on lasers). However, fiber-based devices using the Raman Effect are limited because they require kilometers of fiber to provide sufficient amplification.




Fig 2.1: The Raman Effect allows energy from pump beam to amplify data at longer wavelengths in glass fiber (a). This could now be done in silicon as well and at less effort (b).

The Raman Effect is more than 10,000 times stronger in silicon than in glass optical fiber, making silicon an advantageous material. Instead of kilometers of fiber, only centimeters of silicon are required (Figure 2.1b). By using the Raman Effect and an optical pump beam, silicon can now be used to make useful amplifiers and lasers.

2.3 Two-Photon Absorption

Usually, silicon is transparent to infrared light, meaning atoms do not absorb photons as they pass through the silicon because the infrared light does not have enough energy to excite an electron. Occasionally, however, two photons arrive at the atom at the same time in such a way that the combined energy is enough to free an electron from an atom. Usually, this is a very rare occurrence. However, the higher the pump power, the more likely it is to happen.

Eventually, these free electrons recombine with the crystal lattice and pose no further problem. However, at high power densities, the rate at which the free electrons are created exceeds the rate of recombination and they build up in the waveguide. Unfortunately, these free electrons begin absorbing the light passing through the silicon waveguide and diminish the power of these signals. The end result is a loss significant enough to cancel out the benefit of Raman amplification.

2.4 First Continuous Silicon Laser

The development of the first continuous wave all-silicon laser using a physical property called the Raman Effect is disclosed on Feb 2005 by Intel's Researchers. They built the experimental device using Intel's existing standard CMOS high-volume manufacturing processes.

The breakthrough device could lead to such practical applications as optical amplifiers, lasers, wavelength converters, and new kinds of lossless optical devices. A low-cost silicon Raman laser could also inspire innovation in the development of new medical, sensor, and spectroscopy devices.




Fig2.2: By inserting a diode-like PIN device in the waveguide, Intel removed the electrons generated by two-photon absorption and produce continuous amplification.

The solution is to change the design of the waveguide so that it contains a semiconductor structure, technically called a PIN (P-type – Intrinsic – N-type) device. When a voltage is applied to this device, it acts like a vacuum and removes the electrons from the path of the light. Prior to this breakthrough, the two photon absorption problem would draw away so many photons as to not allow net amplification. Hence, maintaining a continuous laser beam would be impossible. Breakthrough makes the use of the PIN to make the amplification continuous.




Fig 2.3: The breakthrough silicon laser used a PIN device and the Raman Effect to amplify light as it bounced between two mirrors coated on the waveguide ends, producing a continuous laser beam at a new wavelength.

Figure 2.2 is a schematic of the PIN device. The PIN is represented by the p- and n- doped regions as well as the intrinsic (undoped) silicon in between. This silicon device can direct the flow of current in much the same way as diodes and other semiconductor devices do today in common electronics. Hence, the manufacture of this device relies on established manufacturing technologies and it reinforces the basic goal of silicon photonics: inexpensive, high-performance optical components.

To create the Raman laser, they coated the ends of the PIN waveguide with mirrors to form a laser cavity (Figure 2.3). After applying a voltage and a pump beam to the silicon, researchers observed a steady beam of laser light of a different wavelength exiting the cavity – the first continuous silicon laser.

2.5 First Cascaded Silicon Raman Laser

Since the invention of the laser in the early 60s, countless scientific, industry, military, medical and commercial laser applications have been developed. Each of these applications relies on one or more of the special properties of the laser, such as high coherency, high monochromaticity, or the ability to reach extremely high powers. While most of the lasers are based on light amplification by stimulated emission of radiation predicted by Einstein, a special category of laser is based on stimulated Raman scattering and, therefore, is referred to as Raman laser.

A major attraction of Raman lasers is their ability to use light from an optical “pump” to generate coherent light emission in wavelength regions that are hard to reach with other conventional types of lasers. In addition, Raman lasers can be made from materials such as silicon that do not possess suitable energy band structures to produce laser light by stimulated emission. Four years ago, researchers from Intel’s photonics technology lab demonstrated the first continuous-wave (CW) silicon Raman laser.

Since then, they have made several major improvements. Instead of applying optical coatings on the chip facets to form a laser cavity, we designed a “mirror-less” ring cavity laser which is monolithically integrated on a chip and fabricated entirely in a CMOS pad. We have achieved 10 times reduction in lasing threshold to below 20 mW and 5 times increase in output power to more than 50 mW. This cavity design is scalable to enable size reduction and can be integrated with other silicon photonic components on the same chip. Another significant characteristic of the silicon Raman laser is its extraordinary spectral purity. This is very important for many laser applications such as high-resolution spectroscopy.

These improvements led to the current breakthrough in achieving cascaded Raman lasing in silicon. Cascaded lasing is a unique attribute of Raman lasers. A Raman laser's output beam (which is always at a longer wavelength than the pump) can itself act as a pump to generate Raman lasing at an even longer wavelength. This so-called cascaded process can continue to generate longer wavelengths, as illustrated here. A cascaded Raman laser can be realized by properly designing the laser cavity to take advantage of this unique property. Cascaded lasing enables Raman laser to have greater wavelength coverage, and one can use a low-cost and efficient laser as a pump to generate coherent light in longer wavelength region where it is difficult to achieve lasing by other methods.




Fig 2.4: Cascaded Silicon Raman Laser.

Silicon is particularly suitable as Raman laser material for the near and mid-infrared (IR) regions due to its high Raman gain and optical

transparency in these regions. Through cascaded Raman lasing in silicon, one can convert pump wavelengths in the near IR region for which sources are well-developed and widely available, to wavelengths in the mid-IR region, providing low-cost, compact, and high performance room temperature lasers. Such laser sources are highly desirable for many applications ranging from trace-gas sensing, environmental monitoring, and biomedical analysis, to industrial process control, and free-space communications. Other advantages of using silicon as cascaded Raman laser material include its unique material properties such as high thermal conductivity and optical damage threshold, as well as its extraordinary material purity and great natural abundance.

We show that when using a pump laser of 1550 nm which is well developed for optical communications, we achieved stable, single-mode 1st and 2nd order CW lasing at 1686 nm and 1848 nm, respectively. The Raman laser output power exceeded 5 mW and its wavelength could be continuously tuned over a 25 GHz range. The realization of 2nd order silicon Raman lasing paves the way toward higher order cascaded Raman lasing, and opens a new path to producing low-cost, compact, room temperature, high-performance mid-IR lasers.

Chapter 3

Hybrid Silicon Wave Guide

Figure 3.1 is a cross-section of the hybrid silicon laser, showing the indium phosphide-based gain material (orange) that generates the laser light bonded on top of a silicon waveguide (gray). The silicon substrate, which is marked in gray at the bottom of Figure 1, is the base upon which the other items are placed. On this substrate rests the silicon waveguide. Both the substrate and the waveguide are manufactured using standard silicon fabrication processes.




Fig 3.1: Cross section schematic of a Hybrid Silicon Laser

Both the silicon wafer and the indium phosphide-based wafer are then exposed to the oxygen plasma, which leaves a thin coating of oxide on each of the two surfaces that acts as a glue layer. The oxide layer is only 25 atoms thick, yet it is strong enough to bond the two materials together into a single component.

The oxygen plasma that is used for this layer is similar in concept to the plasma used in fluorescent light bulbs and modern high definition plasma TV screens. Plasma is a gas that has been electrically charged. While

fluorescent bulbs are based on plasma that derives from neon or argon gases, the hybrid laser relies on oxygen plasma to coat the components and make them bond.

When the silicon and the indium phosphide-based material are heated and pressed together, the two oxide layers fuse them together. Electrical contacts, shown in yellow in Figure 3.1, are then patterned onto the device.




Fig 3.2: When voltage is applied to the contacts, current flows, and the electrons (-) and holes (+) recombine in the center and generate light.

As shown in Figure 3.2, when the voltage is applied, the light generated in the indium phosphide based material passes directly through the glue layer into the silicon waveguide below, which acts like the laser cavity to create the hybrid silicon laser. The design of the individual silicon waveguides is critical to determining the performance of the hybrid silicon laser, and will allow future versions to be built that generate specific wavelengths.

Chapter 4

High Speed Modulation In Silicon

4.1 First GHz Silicon Modulator

Optical modulators are used to encode a high-quality data signal onto an optical beam, effectively by turning the beam on and off rapidly to create ones and zeros. Before the year 2004, no one had built an optical modulator from silicon that was faster than about 20 MHz. In February of 2004, the first gigahertz silicon optical modulator was disclosed. By integrating a novel transistor-like device, researchers were able to create a modulator that scaled much faster than previous attempts. In 2005, researchers further demonstrated that this silicon modulator is capable of transmitting data up to 10 gigabits per second (Gbps).

4.2 How Intel's Silicon Modulator Works

To understand how the modulator functions, it is important to touch briefly on the nature of light. Light is an electromagnetic wave that occurs at specific frequencies, some of which are visible and some, like ultraviolet and infrared, that are invisible to the naked eye. When light is emitted it travels in the pattern of a sine wave. (See the top row of Figure 8.) The total distance reached by the peaks and troughs of this sine wave is known as amplitude. When the sine wave is nearly flat, the light is at its dimmest and has low amplitude. When the peaks and troughs are very high and deep, the light shines brightly and has greater amplitude.




Fig 4.1: The effect of adding two light waves together.

When two beams are combined, the resulting sine wave is the sum of the two constituent sine waves. For example, if two sine waves are perfectly in sync and added together (left column of Figure 4.1), the resulting sine wave has twice the amplitude of the individual waves. In contrast, when two waves are completely out of sync (right column of Figure 4.1), the resulting wave has no amplitude. In Figure 4.1, for example, see how the peak of the top wave on the right has a value of +1. When it is aligned with the trough of the lower wave, which has a value of -1, the net result is 0: The peak and trough offset each other exactly. In this case, the two light waves cancel each

other out and the resulting light wave is off for the duration of these two sine waves.

The degree to which two beams of equal wavelength are in sync is called phase by optical engineers. When two waves are in phase, they are aligned so that their peaks line up, as do their troughs. When they are out of phase, their peaks and troughs offset each other and the light dims. These changes in amplitude (the strength of the light) are used on the receiving end by the photodetector to recognize 0s and 1s. Because the amplitude is being modulated to encode the data, this technique is referred to as amplitude modulation (AM). This AM is the same technique used in AM radios. Radio transmitters use changes in amplitude to encode variations in pitch on a radio wave. The AM radio receives the wave and from the changes in amplitude re-creates the sound that was encoded.




Fig 4.2: Intel's modulator splits the incoming light into two beams. Those beams are then phase shifted and recombined to encode the data.

Intel's modulator (see Figure 4.2) takes the incoming light beam and splits it into two beams. When, the beams are phase shifted relative to each other. The two beams are recombined, and the phase shifting changes the amplitude of the resulting beam so that it goes bright and dark, thereby encoding the data.

The significance of this particular advance is that the Intel phase shifter can perform this modulation at speeds in excess of 1 GHz. In so doing, Intel has raised the previous performance ceiling in silicon by a factor of over 50x.

Intel expects to achieve even greater bandwidth by multiplexing these data streams. This approach could bring silicon photonics into an age where 10 Gbps and 40 Gbps bandwidth are not uncommon. At these capacities, silicon photonics could present a compelling, inexpensive choice for commercial uses, especially as backbones and wiring for corporate campuses.

4.3 World's First 40G Silicon Laser Modulator

A Recent breakthrough in Silicon Photonics research at Photonics Technology Lab of Intel reveals a laser modulator that encodes optical data at 40 billion bits per second. We have finally reached the goal of data transmission at 40 Gbps speed, matching the fastest devices deployed today using other materials.




Fig 4.3: 40Gbps Silicon Laser Modulator.

The Intel modulator is based on a Mach-Zehnder interferometer with a reverse-biased p-n junction in each of the arms (Figure 4.3.1a). When a reverse voltage is applied to the junction, free carriers – electrons and holes resulting from the n- and p-dopants – are pulled out of the junction, changing its refractive index via the free-carrier effect. The intensity of the light transmitted through the Mach-Zehnder interferometer is modulated by modulating the phase difference between the interferometer's two arms. This modulation can be very fast, because free carriers can be swept out of the junction with a time of approximately 7ps. The modulator speed is thus limited by the parasitic effects such as RC time constant limit.

To minimize the RC constant limitation, Intel researchers adopted a traveling-wave drive scheme allowing electrical and optical signal co-propagation along the waveguide. The traveling-wave electrode which is based on a coplanar waveguide was designed to match the velocity for both optical and electrical signals, while keeping the RF attenuation small. To operate the traveling-wave modulator, the RF signal is fed into the transmission line using a commercially available driver from the optical input side and the transmission line is terminated with an external resistor (see Fig. 4.3.1a). After packaging the modulator on a printed circuit board,

the researchers demonstrated that the modulator has a 3 dB bandwidth of ~30 GHz and data transmission capability up to 40 Gbps.

The high-speed silicon modulator could find use in various future applications. For example, a highly integrated silicon photonic circuit may provide a cost effective solution for the future optical interconnects within computers and other devices. With the demonstration of the 40 Gbps silicon modulator and the electrically pumped hybrid silicon laser, it will become possible to integrate multiple devices on a single chip (Fig. 4.4) that can transmit terabits of aggregate data per second in the near future – truly enabling tera-scale computing.




Fig 4.4: Multiplexed 40Gbps Modulator enabling tera-scale computing.

Chapter 5

Avalanche Photodetector

Intel has been doing research in this area for more than 5 years and has already reported on silicon modulators, silicon Raman lasers, and hybrid InP-Si lasers. Last year we also published on a photodetector made from germanium and silicon that had a bandwidth of 31 GHz. The use of Ge is important because, unlike Si, it can efficiently detect light in the near infrared which is the standard for communications. The drawback is that so much stress is developed in pure Ge films deposited on Si that defects are introduced near the Ge/Si interface. Careful design and processing is needed to minimize the impact of these defects on the electrical performance of the device.




Fig 5.1: Silicon Germanium waveguide based photodetector.

This is now a different type of Ge/Si photodetector that has built-in amplification, which makes it much more useful in instances where very

little light falls on the detector. It is called an avalanche photodetector because an avalanche process occurs inside the device. First, a negative and a positive charge (electrons and holes in semiconductor terminology) are created when the light strikes the detector. The electron is accelerated by an electric field until it attains a high enough energy to slam into a silicon atom and create another pair of positive and negative charges. Each time this happens the number of total electrons doubles, until this “avalanche” of charges are collected by the detection electronics. This amplification effect (called gain) is the key to the device, and it serves as the motivation for why anyone would try to do this in silicon and not just continue to use traditional InP-based APDs. The materials properties of silicon inherently led to lower noise and better performance in this avalanche process. Another reason relates to this bit economic trivia; an individual 10 Gb/s InP APD can sell for more than \$200 currently and has a semiconductor area of roughly $400 \times 400 \mu\text{m}^2$. Even the much cheaper 1-2 Gb/s APDs used in fiber to the home (FTTH) still sell for \$3-5.

It has often been assumed, however, that while silicon photonics might be lower in cost than InP-based devices, its performance would be inferior. While this is true in many cases, one of the exceptions is the area of APDs, where silicon’s material properties allow for higher gain with less excess noise than InP-based APDs and a theoretical sensitivity improvement of 3-5dB. Sensitivity is the gold standard of detector benchmarks and is defined as the smallest amount of optical power falling on the detector that can still maintain a desired (low) bit error rate. We have recently achieved a monolithically grown Ge/Si APD with a sensitivity of -28dBm at 10Gb/s and a gain-bandwidth product of 340GHz. This sensitivity is equivalent to

mature, commercially-available InP APDs and the gain-bandwidth product (GBP) is the highest reported for any APD, as shown in the graph below. The GBP is important because it describes over what frequency range that the gain of the device is available. InP-based devices typically have a GBP of ~ 100 GHz which means that they would have a gain of 10 at 10 GHz. However, at bandwidths high enough to support 40 Gb/s, the gain falls to about 3 which is not enough to justify the cost. Our device would have a still have gain of 10 at that same point. In order to realize the full performance potential from this material system though, we need to further reduce the dark current that is coming from the defects at the Ge/Si interface, and stop the inter diffusion of Ge and Si that occurs during annealing. This intermixing is problematic since the Ge causes higher noise than if the silicon alone was in the multiplication region. If we are successful, this work will pave the way for developing low cost, CMOS-based Ge/Si APD operating at data rates of 40Gb/s or higher in the future.




Fig 5.2: Gain Bandwidth Product

There are two other directions that we are planning to go with derivatives of this technology. The first is to move to a waveguide-based APD. This will improve the absorption at wavelengths up to about 1600 nm because the effective absorption depth can be much greater in that type of device. It also allows for integration with other optical devices, as such demultiplexers and attenuators. Secondly, we would like to reduce the operational voltage from the industry standard of ~30V bias to something more common in consumer electronics to open up a much broader user base.

Chapter 6

6.1 Benefits of Silicon Photonics

The principal benefit of the hybrid silicon laser is that silicon photonics components no longer need to rely on aligning and attaching discrete lasers to generate light into a silicon photonic chip. In addition, dozens and maybe even hundreds of lasers can be created with a single bonding step. This has several advantages:

- The laser is compact so it allows many lasers to be integrated on a single chip. The first demonstration hybrid silicon laser is only ~800 microns long. Future generations will be significantly smaller.
- Each of these lasers can have a different output wavelength by simply modifying the silicon waveguide properties without having to modify the Indium phosphide based material.
- The materials are bonded with no alignment and are manufactured using high volume, low cost manufacturing processes.

6.2 The Optical Future

As Moore's Law continues to push microprocessor performance, and as increasing volumes of data are sent across the Internet, the demands placed on network infrastructure will increase significantly. Optical communications and silicon photonic technology will allow enterprises to scale bandwidth availability to meet this demand.

In addition, due to the low cost of silicon solutions, servers and high-end PCs might one day come standard with an optical port for high-bandwidth communication. Likewise, other devices will be able to share in the bandwidth explosion provided by the optical building blocks of silicon photonics.

By creating the PIN device to sweep away free electrons in silicon waveguides, Intel delivered a significant breakthrough: a silicon component that can create continuous-beam Raman lasers and optical amplifiers.

Research into silicon photonics is an end-to-end program that pushes Moore's Law into new areas. It brings the benefits of CMOS and Intel's volume manufacturing expertise to fiber-optic communications. The goal is not only achieving high performance in silicon photonics, but doing so at a price point that makes the technology a natural fit – even an automatic feature – for all devices that consume bandwidth. Intel's breakthrough

continuous silicon Raman laser will undoubtedly contribute to the reality of this vision.

6.3 References

- Google Search Engine.
- ieeexplore.com
- research.intel.com
- IBM.com
- Book on Silicon Photonics – Mario Paniccia.