CIR
enlightened analysis™

Silicon Photonics and Optical Integration

April 2007

CIR
PO Box 3840
Glen Allen, VA 23058
www.cir-inc.com

Silicon Photonics and Optical Integration

Creating optical chips in a manner similar to electronic chips has been discussed for a couple of decades now and there have been impressive demonstrations. Until very recently, however, the commercial applications for optical integration have been few and far between. Typically, optical integration has involved compound semiconductors that are relatively expensive and hard to work with and early implementations of optical integration have often looked more like cool technology in search of an application than a way forward. In addition to such market problems, there are also implicit technical problems including excessive heat from multiple lasers on a chip and the ability to create precise alignments for optics at the chip level in a mass production environment. These have all been tackled in different ways, but none can really be said to be solved.

At its technical best, integrated optics way overshot network requirements. At the height of the optical boom, the main analyst on this report remembers seeing an NTT Labs researcher who talked about the 1,000-channel AWG that his organization had built and how this could serve as the basis for many new kinds of optical chips with equally impressive multi-channel capabilities. Although a truly impressive achievement, a half dozen years later most WDM systems being installed have channel counts measured in the 10s – at best. Many are in single digits. Paradoxically, optical integration could underachieve too. When equipment vendors actually tried out some of the products of optical integration in the early part of the last decade they discovered that it worked less well than a more conventional solution. The cure, it seemed, was worse than the disease.

As we stress throughout this report networking applications have now evolved to a point where optical integration is far less likely to overshoot needs than a few years ago. There now seems to be a belief in the industry that silicon photonics and optical integration will play a major role in optical component design in the next few years. Almost two thirds of the respondents to the informal Web survey carried out by CIR thought that it would play a major role by 2010.

The Future Role for "Traditional" Integrated Optics

Our survey respondents foresaw an impact on most areas of components business, although several believed that optical integration using – for example, InP – was a much nearer term technology than silicon photonics. In fact, this kind of approach is already fairly standard in the form of integration of lasers and modulators. In recent presentations for the IEEE's Higher-Speed Study Group, CyOptics and others have suggested that this kind of integration might be key to building high performance components for high data rate networks. Some examples of the kind of components that this firm has in mind are provided in Exhibit 3-3. In addition to CyOptics, another firm that deserves a mention in this area is Infinera which has designed its own core chip that is apparently an InP-based muxing array that enables 10 x 10 Gbps of WDM capability.

The fact that Infinera has also demonstrated a similar 40 x 40 Gbps array shows that it potentially controls a very powerful technology for the future, although, of course, the 40 x 40 array is nowhere near commercialization. And, while the objective of this kind of optical integration is on both enhancing performance and lowering costs, our sense of the market (and of what can be accomplished with InP and similar materials) is that the focus is more on the former than the latter. The approaches of Infinera and CyOptics discussed above differ in that one (CyOptics) uses monolithic vertical integration in which several different functionalities are integrated together in the same chip, while Infinera is looking at horizontal integration; that is the creation of arrays.

Another useful distinction is between monolithic integration and (the much less challenging) hybrid integration. For example, IBM has recently announced its development of an integrated 160 Gbps transceiver which is built by bonding a silicon chip for the electronics with an integrated 4 x 4 matrix array of 10 Gbps VCSELs. The argument for a hybrid approach is that it can combine the best of silicon and other semiconductors. However, it is still quite hard to achieve in practice. An appropriate process for creating a hybrid chip in large quantities has yet to be found. One production strategy that has been considered is epitaxial regrowth procedures, but these are complicated and have low yields. Intel and the University of Santa Barbara have recently demonstrated another approach to hybridization which uses a bonding process to essentially glue a silicon chip including waveguides to an InP chip, which would provide gain.

Exhibit 3-3
New Integrated Products for Parallel Solutions

Product	**Markets**	**Issues**
CWDM Products:		
DFB + PIN diode	Up to 10 km over single-mode fiber	Possible short-reach applications for multi-mode fiber, but this could be very challenging. If 5 x 20 Gbps format is used, PIN may be inadequate at longer reaches
DFB + PIN diode + EDC	LRM-like solution for up to 300 meters over multimode fiber	Would probably use 10 x 10 Gbps format and would require availability of advanced low-cost EDC
DFB + APD	Up to 10 km over single-mode fiber	
DWDM Products:		
EML + APD	40 Km single mode	Could use any format

Source: CIR, based on CyOptics presentation

Commercialization Potential of Silicon Photonics

In the blossoming area of silicon photonics, the idea is to leverage the decades of experience with making (electronic) chips from silicon as well as the installed base of (silicon) semiconductor manufacturing capacity to build low-cost optical components for a variety of applications. Performance is still important here, but the focus is on lowering costs and volume production, much as it is in the regular semiconductor industry. This would seem to be very useful for certain areas – notably PONs and Ethernet (and possibly Fibre Channel) where high volumes seem likely to be a market requirement.

Before discussing the unique aspects of silicon photonics, it is somewhat important to define "high volume" in the context of this report. Traditionally, optical components have been created in the thousands or tens of thousands. And the manufacturing of these products have been something of a "craft industry," involving a small number of highly skilled technicians in a clean room. As optical Ethernet, optical Fibre Channel and PONs become a major part of the networking business, volumes of specific optical components will have to rise into the millions. But this is still small potatoes compared with the volumes of chips that must be manufactured for the computer and mobile phone industry. These run into the hundreds of millions annually for certain individual kinds of chips. The bottom line is that while volumes currently being talked about in the optical components business seem huge to veterans of the optical components industry, they may actually seem quite unimpressive to those in the semiconductor industry. So while several foundries and chip makers are currently quite interested in the opportunities presented to them by silicon photonics, nobody really expects to fill up their fabs as a result.

However, even though silicon photonics may never translate into huge numbers of wafers passing through a fab or foundry, optical components created by silicon photonics may soon represent a decent throughput and while these components will certainly carry a premium compared with similar electronic chips created in much higher volumes, significant price/performance improvements are to be expected. There are two reasons why this opportunity is only just beginning to emerge. One is – as we have already stressed – volume opportunities for optical components are only just beginning to appear. The other is that until very recently, the optical properties of silicon were supposedly such that this material did not lend itself to making optical components in commercial quantities, although R&D work on silicon photonics in university labs dates back to the 1980s. Three major challenges have been (1) getting silicon to emit light efficiently compared with InP. Silicon tends to emit energy in the form of heat. (2) the absence of a strong electro-optic effect, which means that silicon is not very good at modulating a laser beam and (3) Silicon is poor at photodetection—converting photons into electrons—at the infrared wavelengths commonly used for optical communications.

All of these issues are facts of nature and aren't going to change, but recent R&D by firms such as IBM, Intel, Kotura, Luxtera and STMicroelectronics, as well as by university labs have made progress in getting round these problems and have therefore made the conventional view of the role of silicon in the optical components business seem quite limited. Apart from commercial firms, labs that have become important centers for silicon photonics include Cornell, MIT, UCLA, the University of Catania (Italy), the University of California at Santa Barbara, the University of Rochester, the University of Trento (Italy), and the University of Surrey (U.K).

Silicon photonics has seen some impressive developments in recent years, especially when one considers that silicon photonics was once supposed to be impossible in a commercial sense. On the passive side of the house, silicon has been shown to be an adequate waveguide. Rib waveguides have been built with silicon that exhibit propagation losses of well under 1 dB/cm. In addition, the tight mode confinement allows sharp bends without excessive bend loses. Silicon waveguides could also play a role in creating the first silicon optical amplifiers. It turns out silicon is a suitable material for Raman amps, because the Raman gain coefficient of silicon is very high and waveguides can be used confine the mode to a very small area. (This is one example, at least, where the physics of silicon works in favor of silicon photonics.) However, a silicon photonics op amp would probably have to be a hybrid affair in which a (non-silicon) pump laser would be used.

Important definitional and market positioning issues are raised by silicon photonics. Does silicon photonics include every attempt to create optical products with silicon; conventional silicon waveguide products such as those offered by JDSU, for example? Or does it just include firms that rely on fully-fledged CMOS processes? The conventional wisdom would seem to be that CMOS should be involved, although this could change, especially if "silicon photonics," increasingly becomes something of a buzz word. Conversely, there are some different opinions about just how far it is worth pushing the silicon photonics paradigm. For example, Intel has built silicon lasers. But IBM believes that actually building these silicon lasers in high quantities may simply not lend itself to existing CMOS processes.

This laser issue is important and today most, if not all, silicon lasers are hybrids in some sense. They may use an external laser pump source, and this approach is apparently relatively easy to achieve. However, this kind of optically pumped solution does not lend itself to mass production. Much closer to the optical integration ideal is the well-publicized version of an electrically stimulated silicon laser chip due to Intel-funded work at the University of California at Santa Barbara. This is a hybrid chip made from both InP and silicon. However, we are not talking about an InP pump laser here. Rather the InP layer generates the initial light that is then lased in a silicon waveguide. According to one press report, the silicon part of the chip contributes 97 percent of the light energy, and it is the silicon that determines the laser's performance and wavelength, not the InP.

In this chip, the bonding of the InP and Si is done at wafer scale. Combining InP and Si to create such hybrid chips is quite hard to achieve in practice. An appropriate process for creating a hybrid chip in large quantities has yet to be found. Epitaxial regrowth procedures have been tried, but these are complex and have low yields. Intel and the University of Santa Barbara apparently used a bonding process to essentially glue a silicon chip including waveguides to an InP chip, which should prove scalable in a commercial setting. The resulting hybrid device operates at 1550 nm, making it suitable for telecom applications, and the general technique could be used in amplifiers as well. In the lab, the researchers got the lasers to produce light of 1.8 milliwatts, which would make it suitable for chip-to-chip optical interconnects, using an input power of 65 milliamps. When Intel announced this breakthrough during 2006, it said that it expected this kind of laser to be mass produced by the early part of the next decade, by which time its performance would be considerably improved. By that time, the other parts of the silicon photonics jigsaw puzzle should be in place.

It should also be noted that in the context of market positioning that while we believe that silicon photonics will be of immense importance for the future of optical communications, not all firms are really focused on this space. Luxtera, for example, does seem to be somewhat focused on providing densely packed transmission subsystems that provide high bandwidth at a modest cost, but is also interested in using silicon photonics in consumer electronics (displays) and in optical interconnection for computers. Not surprisingly, optical interconnection is also of interest to IBM and Intel, which sees silicon photonics as the way forward for chip-to-chip and board-to-board (and possibly ultimately on-chip) interconnection; an application that could require billions of components if widely implemented. This latter use of silicon photonics would solve an emerging problem in which the escalating speeds of processors turn the networks connecting them into a bottleneck. At a much smaller volume level much of the current interest in silicon photonics (and more optical integration more generally) is being generated by interconnection in the high-performance computing sector and in the U.S. there has been government funding for this kind of project.

A Timetable for Silicon Photonics:

The most likely evolution of silicon photonics will occur in four phases.

The first phase will involve silicon wavelength products and is with us now. Kotura licensed Bookham's old silicon photonics IP portfolio back in 2005 and is beginning to turn it in to real products. The Kotura/Bookham technology covers the design and manufacturing of waveguide devices including VOAs and modulators and Kotura has already debuted a portfolio of VOAs and VOA arrays including integrated power monitoring features and a VOA array combined with a tap and detector. To tell the truth, Kotura is doing something very similar to what Bookham was doing with the original version of this technology years ago and the VOA market never really did take off the way that it was expected. But several things are likely to have changed. First, the technology is likely to have improved – Bookham had some real teething problems with its version. Second, the actual (as opposed to the imagined) addressable market has expanded to include ROADMs, a product category that barely existed when Bookham was playing in this space. Third – a point that hardly needs explanation – Kotura has far fewer competitors than it would have done in (say) 1999.

For our purposes, however, the kind of approach to silicon photonics that is represented by Kotura's efforts in this space is important because its actually generating revenues now. Tropic Networks has signed a two-year supply agreement for Kotura's Ultra Fast Variable Optical Attenuator (VOA) Arrays. Supposedly the deciding factor in this deal was the fast power management capability that these arrays offer and which may in part derive from the silicon photonics approach that Kotura has adopted.

The second phase of the evolution of silicon photonics is strongly related to the accelerating shift towards 10 GigE. Somewhat simplifying, all previous generations of Ethernet have begun by providing very high-speed connectivity to large machines in data centers and then spread to servers and local hubs and then on to the desktop. At the moment, 10 GigE is just beginning to reach the high-end of the server market and will need to see serious price reductions if it is ever to get much closer to the customer. There is some controversy about whether 10 Gbps would ever be needed that close to the customer and (separately) about whether very short reach 10 GigE could best be provided with copper. As far as the "10 Gbps is not needed" claim goes, I would simply point out that this kind of thing has been said about almost every generation of Ethernet. Somehow the applications expand to meet the available bandwidth. As far as the copper claim goes, one should remember that 10 Gbps copper interfaces are complex chips themselves and cabling for copper 10 GigE is also very expensive. Although, copper-based 10 Gbps will certainly be part of the 10 GigE wave, there is no reason to suppose it will be the whole of the answer.

This is the part of the market that Luxtera is already tackling. It is already sampling its 10 Gbps silicon photonics transceiver technology, which it says will bring down the cost point for 10 Gbps transponders to a disruptive degree. However, it is worth noting that, while Luxtera can offer a low-cost transceiver on a chip, it is very much a hybrid product. The laser itself is an InP semiconductor laser and is a long way from the kind of silicon lasers that we discussed above. Perhaps these could be added into the mix later. But the important thing is the kind of technology that Luxtera is working with could potentially propel silicon photonics into the mainstream of the telecommunications/data communications and help expand the 10 Gbps market by millions of units.

The third phase of silicon photonics could help swell such volumes by orders of magnitude. Much of Intel's work on fiber lasers is aimed at providing next-generation interconnects for connecting up boards in computers and eventually for connecting up devices on the board or even for interconnecting devices on a chip. IBM is also interested in this possibility for obvious reasons, as is Luxtera to some extent. At present about 99 percent of internal computer interconnection is based on copper and in most cases this seems to do just fine with the task. However, the computer industry has made changes in its interconnection material of choice in the not-too-distant past, the previous paradigm being aluminum. Much of the work that is going on in optical interconnection space at the present time is concentrated on the fairly nichey area of high performance computing, but as multiprocessor architectures become the norm in everyday personal computing, optical interconnects may prove the best way to ensure that interconnection does not become a serious bottleneck impeding computing speeds. And the kinds of low-cost optical transceiver chip that is likely to result from those efforts seem likely to find applications in all forms of optical communications; FTTx seems an obvious beneficiary of this technology.

The last of the phases that we foresee for the evolution of silicon photonics is represented by a class of applications that are too far off to really talk about in any detail or with much certainty. But optical integration in general and silicon photonics in particular would seem to have a potential role in all of these areas, although frankly one of the uncertainties about all of these areas is whether it will be the silicon version of integrated optics that prevails. These more futuristic apps are of two sorts. First there are networks operating at speeds above 10 Gbps. These would seem to cry out for integrated optics in much the same way that 10 Gbps does now – more so in fact. Then there are futuristic applications associated with what has become known as the "physics of information," both of which are based on the quantum mechanical phenomena such as "entanglement." The two applications are quantum encryption and optical computing. At first blush these sound very futuristic indeed. But quantum encryption is already being used by the military and large financial institutions, while primitive optical computers have already built, though they are not yet in commercial use. Both applications will need photon generators and if they eventually turn out to become part of mainstream IT, silicon photonics may have a role in making them so.

For information about other CIR reports, please visit our website at www.cir-inc.com or contact us at (804) 360-2967.