DesignCon 2013

Advances in Onboard Optical Interconnects: A New Generation of Miniature Optical Engines

Jean-Marc Verdiell, Samtec, Inc.
marc.verdiell@samtec.com; 408/524-3167

Abstract

A new generation of miniature optical engines is enabling novel on-board optical interconnect solutions. In this paper, we review the technology and capabilities of these new optical devices. We address many of the practical design aspects of integrating these devices on-board or directly in an IC package including choice of fiber, engine technology, connectors and mounting solutions, heat-sinking, fiber management, compatibility with existing protocols, testing and interoperability.

Author Biography

Marc Verdiell, Business Development Manager, Samtec Optical Group

Marc Verdiell has a Ph.D. from the University of Paris and is the author of more than 100 peer reviewed papers in optoelectronics. Formerly a Bell Labs researcher and an Intel Fellow, he was the founder and CEO of AlpenIO which was acquired by Samtec in 2011. He is currently the Business Development Leader of the Samtec Optical Group.

1 Introduction

As bit rates and I/O density continue to rise according to Moore's law, traditional Copper interconnect solutions are being challenged. Progress in optical technology and miniaturized packaging are enabling a new generation of miniature optical engines that are small enough to be put directly on board next to, or even within, an IC package. These multi-lane devices are currently being designed for speeds up to 28 Gbps per lane making them a strong contender for the upcoming generation of chip-to-chip, chip-to-backplane or chip-to-optical cable ultra-high speed interconnects. In this introductory paper, we explore new offerings, present the main technology choices and address the practical aspect of on-board and on-IC optical interconnect implementations.

2 Technology

2.1 Optical Engine Overview

Optical Engines are conceptually simple devices, as depicted in Fig. *1*. They are essentially multi-lane analog optical-to-electrical converters packaged into small modules. Such devices can be implemented as mono-directional transmitters / receivers or as bidirectional transceivers. The single directional approach has been the standard for larger lane counts (typically 12 channel devices), whereas the bidirectional approach has been the standard for lower lane counts (typically 4 bidirectional channels). Recently, bidirectional 8 and 12 channel devices have been sampling as pre-commercial products. Larger devices with channel counts in the 100's have been demonstrated in research.

Most optical engines include an electronic driver circuitry that will reshape and amplify the electrical input signal so it can correctly drive the optical element, which is typically a semiconductor laser. The laser is simply modulated on and off by its drive current. Such a modulation scheme is often referred to as OOK, On-Off Keying. In practical implementations, the driver circuitry includes numerous refinements to OOK including temperature dependent laser bias and modulation control, as well as equalization and pre-distortion for driving the laser. At higher bit rates, it also provides equalization on the electrical side. In addition, the capability to turn off a channel and monitor laser health might also be included in the driver. One popular type of laser used is called a VCSEL, Vertical Cavity Surface Emitting Laser, whose basic structure is derived from that of an LED. The laser is of course much faster than an LED as it can be modulated into several 10's of GHz modulation regime. It also outputs much more power and has much narrower optical spectral characteristics. These are all necessary elements for high data bit rate fiber transmission. The light emitted by the laser is then captured by an optical system and coupled to the core of an optical fiber.

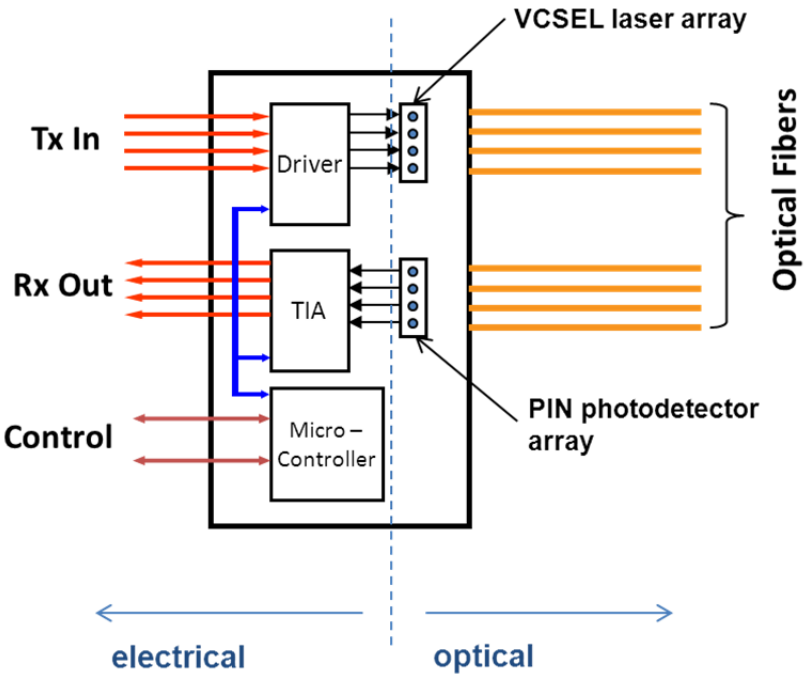

Fig. 1: Typical Optical Engine Block (source: Samtec)

The receiver side is also conceptually straightforward. The light coming out of the fiber is directed via an optical system to a photodetector array. The photodetector, typically a PIN photodiode (named after its P-doped, Intrinsic, and N-doped junction structure) is in turn coupled to an ultra-low noise, very high gain trans-impedance amplifier (TIA) which turns the received photodiode current into an electrically compatible differential voltage output. The TIA output typically incorporates a limiting amplifier (LA) stage and equalization circuitry such as pre/de-emphasis. Advanced functionality such as loss of optical signal detection (LOS), received optical power and squelch might also be implemented.

Optical engines might incorporate a microcontroller to perform internal controls. The microcontroller often interfaces to the system via I2C and enables control of the various programmable transceiver settings as well as reporting temperature, loss of signal and other electrical, temperature or optical alarm conditions (generally referred to as optical digital diagnostics). I2C management is by no means necessary. While telecommunication applications traditionally need to use the optical diagnostics for network test and control features, backplane and embedded applications often treat the transceivers as non-managed units.

Today, most short reach (<300 m) optical engines up to 14 Gbps do not use any digital retiming. This was not always the case. The first 10 Gbps XFP transceivers did incorporate a mandatory clock and data recovery (CDR) circuitry for retiming. As the lasers and high speed circuitry got faster and better, these were subsequently eliminated in the successor SFP+ specification. The same problem, in much tougher guise, is reappearing today at 28 Gbps. The use of mandatory CDRs and retiming in the modules is being debated at the IEEE 802.3 Ethernet forum. Although everyone would like to avoid it due to the significant cost and power penalty incurred, it might need to be included in such open standard links that must cope with a high degree of implementation variability.

However, for backplane and embedded use where the operating environment and engineering parameters are well controlled, the use of CDRs might be overkill. For this reason, Samtec currently plans to introduce 28 Gbps devices with a CDR section that can be optionally switched off.

2.2 Optical Fibers

Limiting ourselves to high bandwidth glass fibers, there are two main fiber types: single-mode or multimode. These fibers differ by the size of their core, which is the tiny region in the center of the fiber that guides the light as shown in Fig. 2.

Single-mode fibers have tiny cores, generally less than 10 μm in diameter. This dimension is so small that only one "optical mode" will transmit through the fiber, which means that the light can only take one single path following the center of the fiber to reach its destination. For this reason, the dispersion, or "smearing" of the pulses of light, is minimized, and the longest transmission distances can be achieved. At 10 Gbps, single-mode fibers are usually needed past 300 m link distance. However, optical transmitters that work with single-mode fibers are significantly more expensive because they use higher cost, long wavelength laser sources emitting around 1500 nm and because the high precision alignment of the laser to the single-mode core of the fiber is more difficult to achieve. Single-mode optical connectors are also more expensive than their multimode counterparts, harder to terminate in the field and much more sensitive to dust and contamination.

By contrast, modern multimode fibers have a large core of 50 μm in diameter. Unlike their single-mode counterparts, the alignment of the lasers is much more relaxed in multimode fibers, and the transmitters and connectors are easier to manufacture. Also, multimode fiber can be coupled to low cost VCSELs emitting at a wavelength of 850 nm. VCSELs have become extremely cost effective because they are easier to build and test, and they are used in extremely large volume consumer applications such as computer mice and DVD players. The disadvantage of multimode fiber is that several modes of light can travel in the core at slightly different speeds, which will eventually smear out a single pulse of light. For this reason, current OM3-grade multimode fibers are usually limited to 300 m transmission at 10 Gbps, a distance generally known as "short reach".

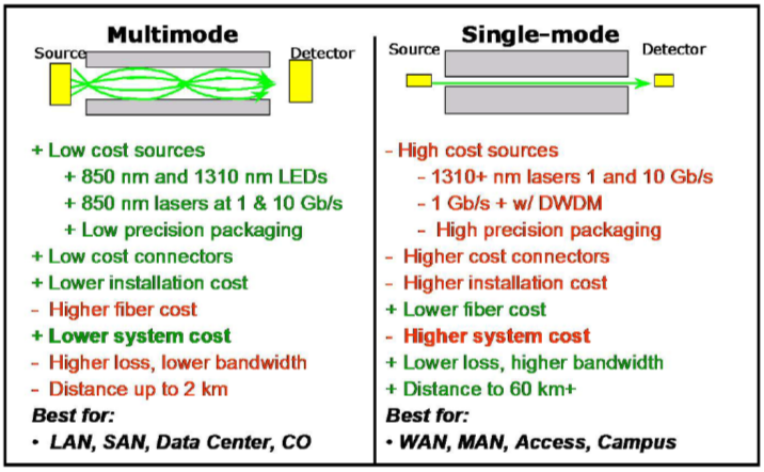

Fig. 2: Multimode vs. Single-Mode Fiber (source: Furukawa)

Counter-intuitively, the cost of multimode fiber itself is higher than that of single-mode, but for relatively short runs (backplanes and rack-to-rack in particular), this does not have any significant economic impact.

Multimode fiber has been the standard fiber deployed for 300 m or less at 10 Gbps. With the exception of Silicon Photonics devices which we will describe below, it is the fiber for which short reach interconnect optical engines have been designed.

Note that there are several versions of multimode fibers. Except for the 62 µm legacy fiber, they differ only by their dispersion characteristics, which determines the maximum reach one can achieve at a given speed:

- 62 µm core multimode fiber. This legacy fiber is a low bandwidth distance product and is not recommended for any new applications.

- OM2 multimode fiber. This 50 µm core fiber can be used for 10 Gbps, shorter length runs. For example, Samtec's current engines can support 50 m of OM2 fiber at 10 Gbps. It is a cost effective alternative to OM3 and might also be prevalent in older cabling installations.

- OM3 multimode fiber. The current standard, also 50 µm core diameter, recommended for all applications up to 14 Gbps. This is the default fiber Samtec ships with current engines and supports up to 100 m distances at 10 Gbps and up to 300 m distances at 5 Gbps. Please note that other vendors' engines might be specified for different maximum transmission length on the same fiber.

- OM4 multimode fiber. The newest 50 µm core fiber standard recommended for upcoming 28 Gbps applications.

2.3 Optical Engine Technology

Fig. *3* shows several examples of optical engine cores from different companies. Note that I use the word *core* on purpose here: this is only part of the story. A complete, practically useable optical engine solution comprises of an electrical connector, an optical connector, a mechanical retention system, fiber stress relief and a heat-sinking solution. In practice, these extraneous elements are often much more important than the optical technology used in the core engine itself. Note that all these cores, despite using different technologies, look relatively similar. Differences will become more apparent when we consider the complete solutions in the next section.

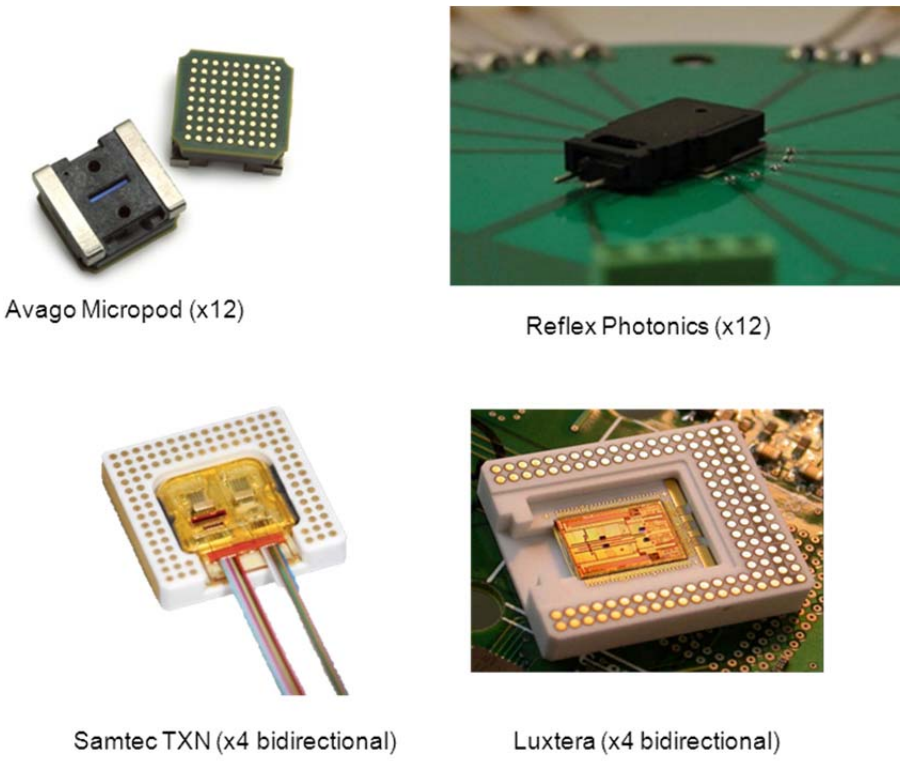

Fig. 3: Optical Engine Cores

Because the core optical technology is often used as a marketing differentiator, comparing vendor claims can get confusing. Technologically speaking, the biggest difference between engine technologies is whether they use a standard VCSEL based approach or a Silicon Photonics approach. VCSEL based approaches are rather straightforward, and a block diagram of such a device has already been given in Fig. *1*. The electrical part is "classic" chip-on-board, while the optical part is usually a small hybrid assembly where a micro-lensing and fiber holding system is used to capture the light coming out of a VCSEL array. The use of precision plastic molded optical blocks, a technology used for high volume cellphone lens camera production, combined with automated machine-vision driven alignment has enabled size and cost reduction of VCSEL-based engines. Such a plastic-optics approach is depicted in Fig. *4*.

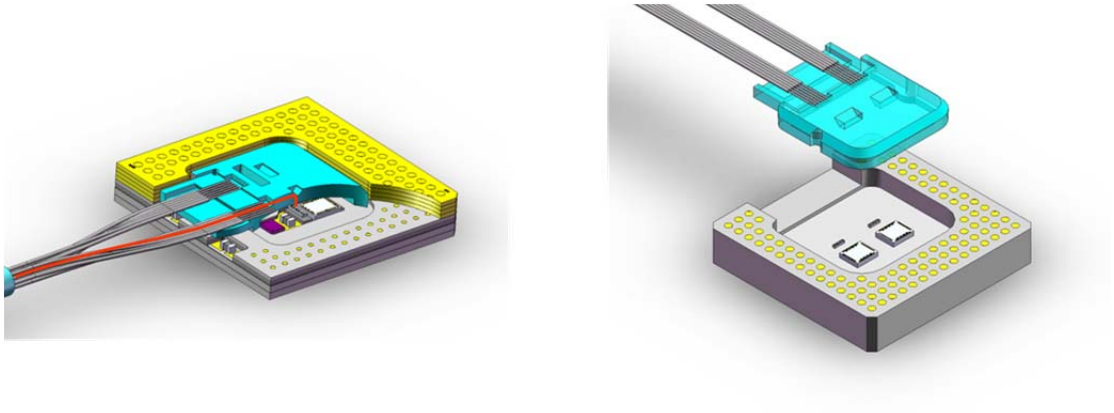

Fig. 4: VCSEL-based Engine with Plastic Optical Block Technology
(source: Samtec)

Silicon Photonics uses a different approach of integration on a Silicon substrate. Obviously, electrical circuits can be manufactured on silicon, but also photo detectors, modulators and optical waveguides are needed to interconnect them optically. The promise is that all the elements can be monolithically integrated into one chip, economically produced in large volume and eventually integrated within processor chips. A picture of such a photonics chip is shown in Fig. 5. The Luxtera engine shown in Fig. 3 above is also realized using Silicon Photonics technology. Note that both of these pictures can be somewhat misleading: unlike the other VCSEL core engines, the modules are shown incomplete as they are not showing the necessary single-mode fiber optical coupling system which would be much larger than the chip itself.

The benefits of Silicon Photonics optical integration have proven difficult to realize in practice because of several practical obstacles. One important element that cannot yet be integrated on silicon is a laser source. Therefore, an external laser source needs to be hybridized and/or coupled to the silicon photonics chip waveguides. Another practical impediment is that silicon is only transparent in the 1500 nm long wavelength band. This means that one is limited to using more expensive long wavelength edge emitting lasers. More importantly, efficient silicon waveguides can only be made for single-mode operation. This means that a single-mode, high precision alignment is needed to couple the laser into the chip and also to couple the light out into the fiber, often negating the cost and manufacturability advantage of the integrated approach. This also makes the silicon-based engines incompatible with traditional short reach, multimode fiber, 850 nm based installations. Finally, even when integrated, the optical devices are very large and very low resolution compared to their electrical counterparts. It might not make economic sense to put the two processes on the same piece of silicon, lest one is willing to incur the high cost of small-feature size processes for the large feature size optical part.

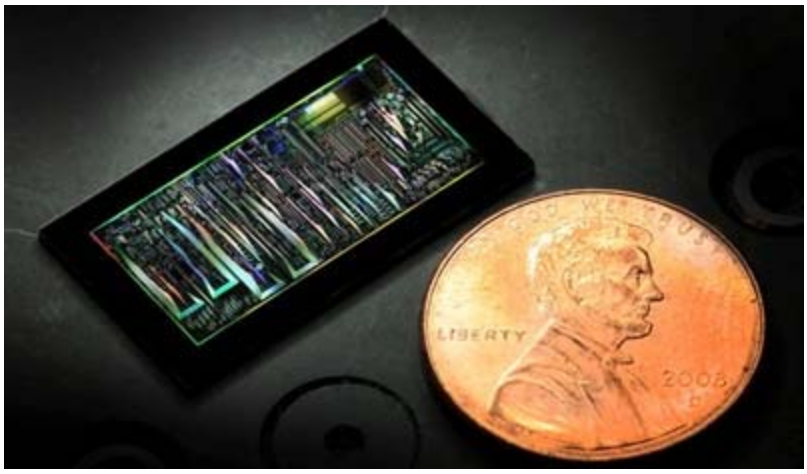

Fig. 5: Silicon Photonics Chip (source: OpSIS)

A lively debate has ensued between visionary proponents of Silicon Photonics technology and commercial makers of traditional VCSEL-based transceivers. The former point to the long term advantages of silicon monolithic integration, while the latter argue that there is nothing inherently wrong, large or expensive with the VCSEL-based approaches. They are also quick to point out the practical shortcomings mentioned above. The debate is likely to continue for a while because arguments of both sides are valid. It will only be settled by the successful commercial introduction of competitively featured and priced Silicon Photonics devices. Several vendors including Luxtera, Kotura, Lightwire and newcomers Aurrion and Skorpios are making continuous progress towards that goal.

3 On-board Optical Engine Types

3.1 Form Factors and Electrical Connector Scheme

As we alluded above, while most engine cores look strikingly i, final packaged products are far more differentiated. One of the first commercial optical transceivers packaged for on-board mounting was the SNAP12, a form factor depicted in Fig. 6. It took advantage of the relatively recent VCSEL technology. The other advances that enabled this form factor were submicron-precision plastic molded MT fiber array connectors and the MEG-Array high speed electrical connector depicted in Fig. 7. It originally provided 12 lanes at 1 Gbps and was one of the first devices to provide short reach links at aggregate speeds of 10 Gbps. This compact form factor has endured. Several of the current engine cores shown in Fig. 3 are available packaged in this form factor updated to speeds of 12x10 Gbps. However, the cost and relatively large size of this first generation solution have mostly limited its use to higher end telecommunication, supercomputing or industrial applications.

Fig. 6: SNAP12 Modules (source: Picolight/Finisar)

Fig. 7: MEG-Array Connector (left, source: FCI) and MTP Fiber Array Connectors (right, source: US Conec)

One of the obvious size limiters of the SNAP12 package is the large electrical connector. One might be tempted to eliminate the electrical connector entirely, but there are many good reasons not to do so. Optics are usually not compatible with reflow soldering temperatures and associated contamination. Often, fiber handling needs to be avoided until completion of the board. Finally, an electrical connector allows replacement of the optical transceiver and to test the board electrically before populating the optical transceivers.

As an alternative to the MEG-Array connector, land grid array schemes have been used more recently to produce miniature engines. One example of such a package is Avago's Minipod, depicted in Fig. *8*, which uses a land grid array connector on the underside as well as a miniaturized Prizm optical connector on the top side.

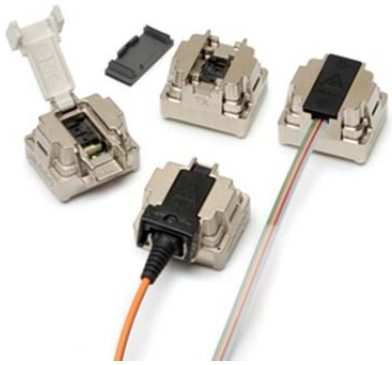

Fig. 8: Minipod Optical Transceiver (source: Avago)

Such smaller form factors have enabled recent demonstrations where the engine is mounted directly on an IC package such as Altera's optically-enabled IC demonstration which uses the Minipod transceiver shown in Fig. 9 below.

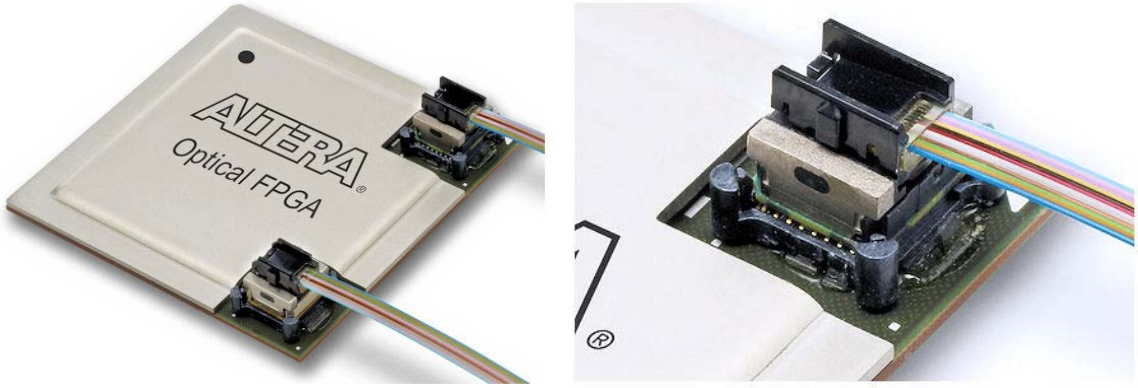

Fig. 9: Optical FPGA Demonstration (source: Altera)

However, land grid arrays were preliminary designed for ICs and have several drawbacks when used for optics. Chief among them is that some sort of compression mechanism or retaining hardware is necessary to hold the device in place. Another difficulty is how to deal with the necessary heatsink, and its retention on the board, while accommodating a fiber output port. Indeed, Fig. 9 only shows part of the story as neither the compression mechanism nor the removable heatsink arrangement is shown. Fig. *10* below illustrates the issue more clearly: Samtec's first generation optical engine, also based on an LGA connector, looks nice and small on its own as shown on the left, but one has to consider the full system including the heatsink and retaining solutions as shown in the center graphic. The resulting density advantage when mounting such engines on a real world board, in this case an x16 optical PCIe interface card using 4 engines, is not fully realized as shown in the graphic on the right.

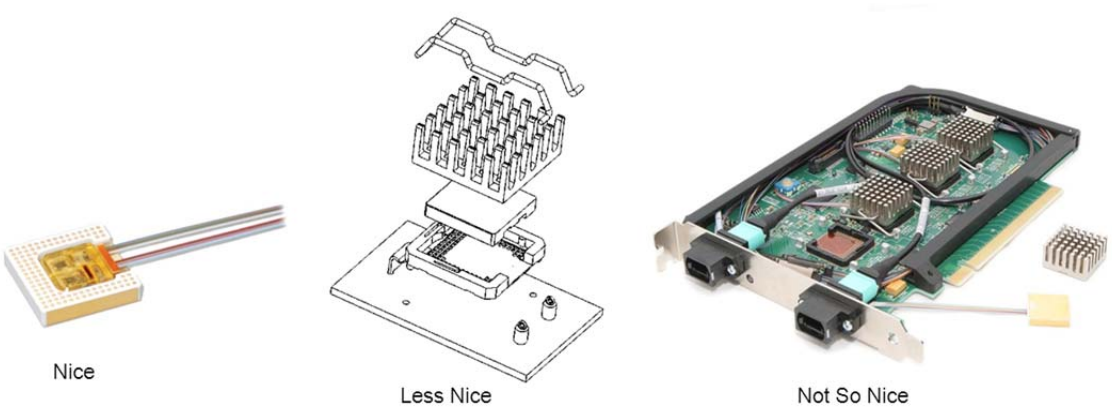

Fig. 10: Bare Engine, Engine with Retaining Hardware, Engine on Board
(source: Samtec/Aprius)

New mounting concepts specifically designed for miniature optical engines are being introduced to address this issue. Fig. *11* below depicts a recently introduced example. This engine inserts with near zero insertion force into a miniature edge-type connector, very much like a miniature memory card. A front connector piece secures the optical engine with a simple latch integrated into the engine heatsink. This connector piece also handles the power and low speed control connections. All high speed data connections are handled by the back piece of the connector which is rated for 28 Gbps.

This arrangement offers multiple advantages: low insertion force, positive latching, small dimensions (an x12 engine connector is 11 mm wide and 3mm thick), simplified high speed board design and lower cost. The connector system is designed with a high density IC package application in mind. Typical size advantage vs. a traditional LGA solution is apparent in Fig. *12*. As shown in Fig. *13*, the connector is narrow enough to allow four 12-channel engines to reside side-to-side along the edge of a 45 mm square BGA package. Another advantage of the edge-type connector system is compatibility with traditional micro-coax Copper cabling as shown in Fig. *14*. The compatible Copper cabling solution can be used instead of the optics for shorter links or for testing purposes before the optics are installed.

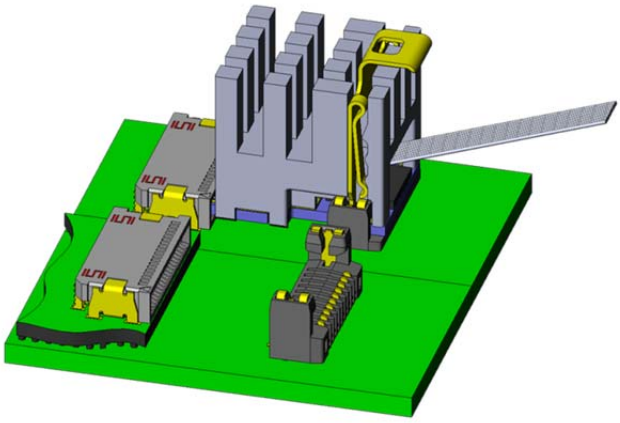

Fig. 11: Miniature Patent Pending On-board Optical FireFly™ Micro
Flyover System (source: Samtec)

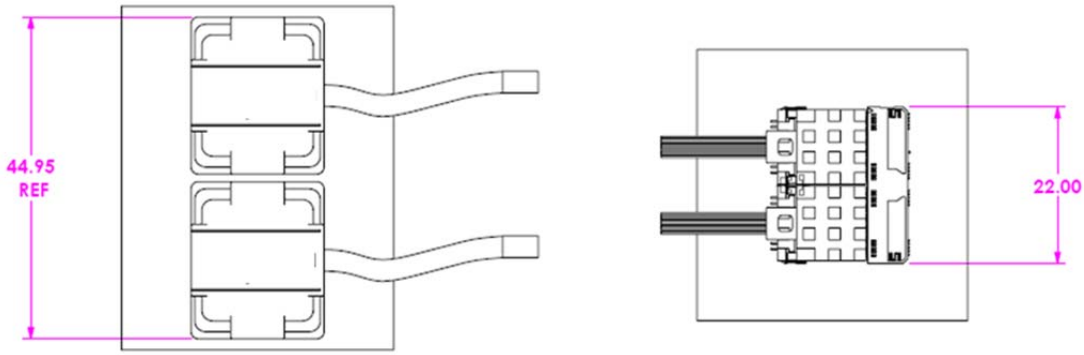

Fig. 12: LGA-based Connector system (left) vs. Miniature Edge Connector System (right)

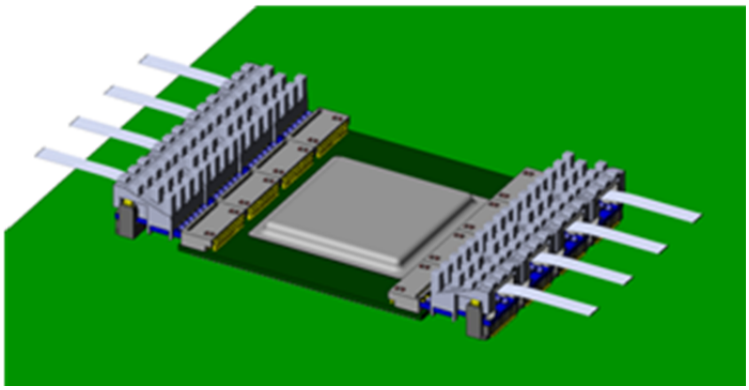

Fig. 13: High Density Optical Connection to an IC BGA Package (source: Samtec)

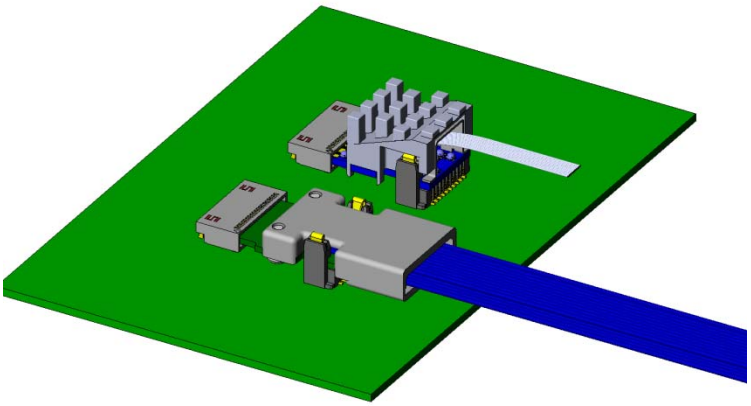

Fig. 14: Plug-compatible Optical and Copper Micro-coax Cabling (source: Samtec)

3.2 Fiber Pigtails

We will talk in more detail about the fiber optic connector options in the next section; however, it is worth mentioning first that optical engines fall into two distinct categories: some have a detachable fiber connector right at the engine, while others have a permanently attached fiber "pigtail". There are pros and cons to both approaches. The removable fiber enables you to exchange the fiber pigtail at will or pair different pigtail lengths to the same engine. However, the downside is a size, cost and performance penalty for that extra connector. The size penalty is particularly noticeable if the engine connector is designed for repeated insertions, calling for a bulky MTP connector like in aforementioned SNAP12. However, it can be minimized if only a few insertions are needed during manufacturing and setup using the Prizm miniature connector solution for example. Note that in many implementations, the optical connector is tucked beneath the heatsink and is therefore not readily accessible. The fiber needs to be attached to the engine before it is placed on the board and the heatsink is installed, negating some of the flexibility advantages. Finally, the cost of the optical "jumper" with connectors at both sides can be significant compared to the cost of the engine.

If the length of the fiber section to the optical connector is fixed at design, the permanent pigtail option might be a preferable and more cost effective solution. Such design usually results in a smaller engine, easier fiber exit handling at the engine and better thermal connection to the heatsink because no fiber connector is in the way. Some vendors, including Samtec, stock tested engines and attach the fiber and heatsink on the engine at the last minute before shipping it, according to customer specifications. This can be more convenient and cost effective than ordering an engine and a custom detachable length jumper separately and assembling the two.

A third interesting option is to use no optical connectors at all! Some vendors offer cables with two miniature engines permanently attached to a length of fiber, thus forming an Active Optical Cable, or AOC. These cables have only electrical connections at both ends. They are therefore robust, insensitive to contamination and are the most cost effective optical link option. They are very simple to connect and deploy, just as one would use a regular Copper cable of fixed length.

3.3 Heatsinking

Heatsinking solutions are another area to be aware of where offerings differ widely. Dissipated power is usually moderate; a full-duplex x4 engine is typically below 1W, while a half-duplex x12 engine is typically less than 2W. However, the area is small, the fiber exit and engine mounting scheme impose some important mechanical restrictions and rated maximum engine case temperature is typically much lower than that of ICs, usually 70 to 75°C maximum. The availability of an integrated, standard or pre-engineered customizable heatsinking solution can be a real asset in shortening the design time. If these are not available, scrutiny should be applied to the mechanical restrictions imposed on the heatsink design, which can make a large difference in the ease of heatsink design and integration.

3.4 AC Coupling

Finally, another often overlooked point to be aware of is the integration of data line AC coupling capacitors in the engine. Engines that do not include these will typically require the customer to put them on their board or IC package, two per differential line, significantly increasing the total space needed to implement a complete solution

4 Optical Connectors

4.1 MT Ferrule

Because parallel optical connectivity solutions have been widely deployed in telecommunication installations for a number of years, the optical connectivity solutions (fiber arrays, jumpers, connectors, etc.) are quite mature.

The development of a cost effective fiber array connector was an important catalyst to the practical deployment of parallel optics in general and optical engines in particular. The plastic MT ferrule developed by US Conec is today's de-facto standard for such connectors. Ferrules are the inside-most part of an optical connector. They hold the fiber with high mechanical precision and perform the critical fiber alignment. The MT ferrule is a submicron-precision, plastic molded part as shown in Fig. *15* on the left and used in MPO connectors on the right. The male and female ferrules are identical with the exception of the presence of two alignment pins on the male ferrule, while the two centering holes are left empty on the female ferrule as shown in Fig. *16*. Fibers are inserted in a line of perfectly dimensioned holes in between the two centering holes/pins. The fiber extremities are polished flush to the front face of the connector during connector fabrication, and the fibers come into actual physical contact when the two ferrules are properly mated. This is referred to as a "physical contact" ferrule system. Such systems are simple, low loss and cost effective, but sensitive to dust. Therefore, careful attention to cleanliness is critical when using such connectors. They are available in both multimode and single-mode versions.

Fig. 15: MT ferrules and MTP® connector
(source: US Coec)

By far the most common ferrule is the x12 single row, but recently x24 ferrules (two rows of 12) have become more commonly available. x48 (4 rows of x12) ferrules have been used in special applications, and up to 6 rows of x12 have been demonstrated in R&D. Because of the wide acceptance of the x12 format, x12 fibber ribbons and x12 jumper cables, even systems with 8 fibers (such as the x4 bidirectional engines) will most often use x12 connectors and x12 fiber arrays; however, 4 of the fibers are left unused and are so called "dark fibers" (see Fig. *19* for an example of a typical x4 arrangement).

The smallest and simplest fiber connection is achieved by using the MT ferrule system with a small metal clip as seen in Fig. *16*. Such a system is for factory installations only and is not meant for user installation or repeated connections. Note that the ferrule itself is not "keyed". In this simplified clip system, ferrules can be mated in two orientations – up-to-up or down-to-up. The up side of the ferrule is only distinguished by the glue dispensing window which can be seen in the right ferrule in Fig. *16*. Therefore, care has to be applied when mating such miniature systems.

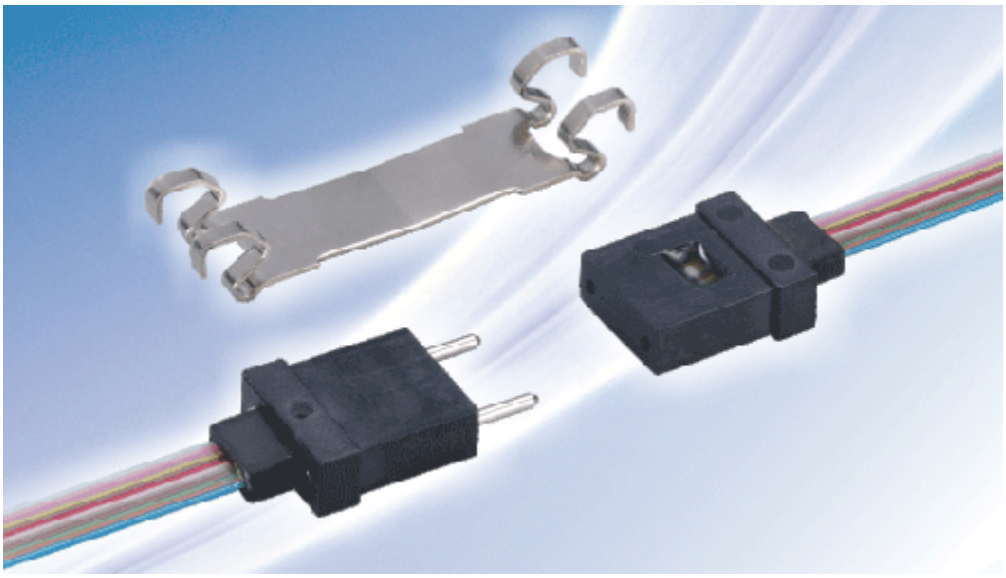

Fig. 16: MT Ferrule Optical Fiber Array Connection System
(source: OE-Tek)

4.2 MPO and MTP Connectors

Fig. *17* shows the Multi-Fiber Push-On, or MPO, connector system based on the MT ferrule. This is an end-user connector meant for repeated hand insertion with keyed polarity. The MPO connector includes fiber strain relief, single-hand operation latching and keying. The connector plugs either directly into the transceiver (for the larger transceiver formats such as SNAP12 or QSFP) or are mated to each other using a bulkhead adapter piece as shown in Fig. *18*. The bulkhead insures proper keying. There are two kinds of bulkhead connectors: a key-aligned and a key-opposed version. For completeness, a higher grade version of the MPO connector with lower insertion loss is available from US Conec under the brand name MTP®.

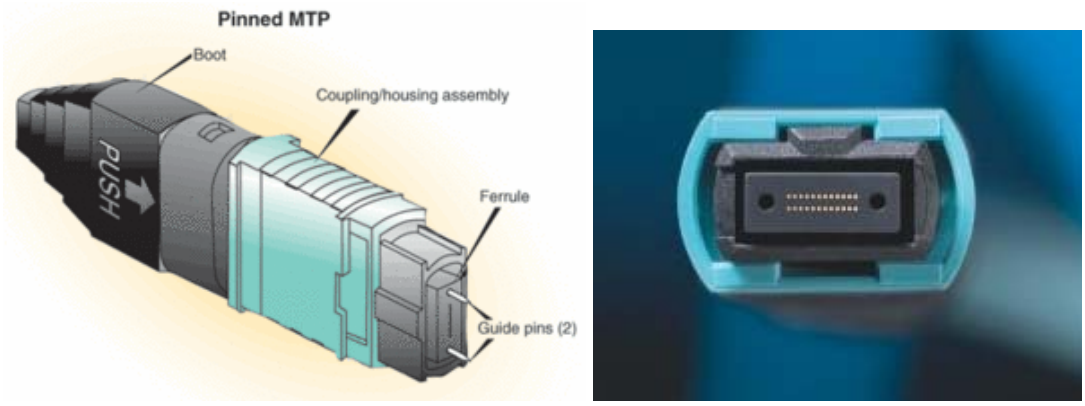

Fig. 17: The MTP® Connector. The MT ferrule and top key are clearly visible on the right image (source: OE-Tek)

Fig. 18 Male and Female MTP Connectors with Bulkhead Connector

4.3 Polarity Considerations

Respecting proper polarity via correct keying is important for proper link functionality. For example, many x4 full-duplex engines will follow the Ethernet QSFP+ standard with receivers (Rx) on one side, four dark fibers in the middle and transmitters (Tx) on the other side as shown in Fig. *19*. These will connect properly with a simple straight-through fiber jumper and a key-aligned bulkhead as shown in Fig. *20*. However, up-to-down key bulkheads and cross jumpers are fairly frequent occurences in data center optical wiring and will prevent the link from operating. It is therefore recommended to check polarity of all connectors and cables, particularly when connecting through a pre-existing fiber plant.

Fiber Number	Function
1	Rx1
2	Rx2
3	Rx3
4	Rx4
5	Dark
6	Dark
7	Dark
8	Dark
9	Tx4
10	Tx3
11	Tx2
12	Tx1

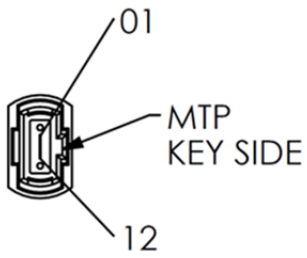

Fig. 19: Typical x4 Full Duplex QSFP+ Optical Connector Arrangement

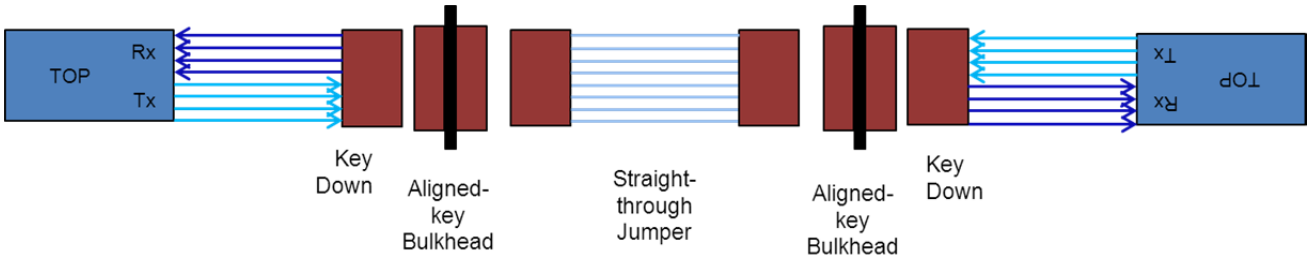

Fig. 20: Properly Keyed Full-Duplex x4 Fiber Link

4.4 Backplane Solutions

Backplane connectivity solutions are also widely available. Several companies carry blind mate connectors based on the MT ferrule. Some can be ganged together for extremely high I/O density like in Fig. *21*, and an example backplane implementation is seen demonstrated in Fig. *22*. Fibers can be routed by laying them onto flex circuits, allowing convenient backplane solutions (Fig. *23*). Several companies are also actively developing waveguide on PCB solutions for even further integration.

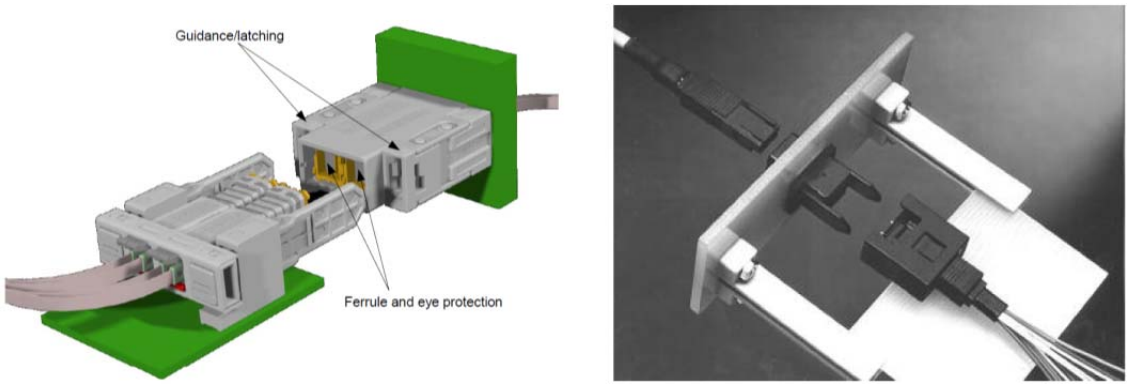

Fig. 21: Blind Mate Optical Connector Solutions (source: Amphenol - left, TE Connectivity - right)

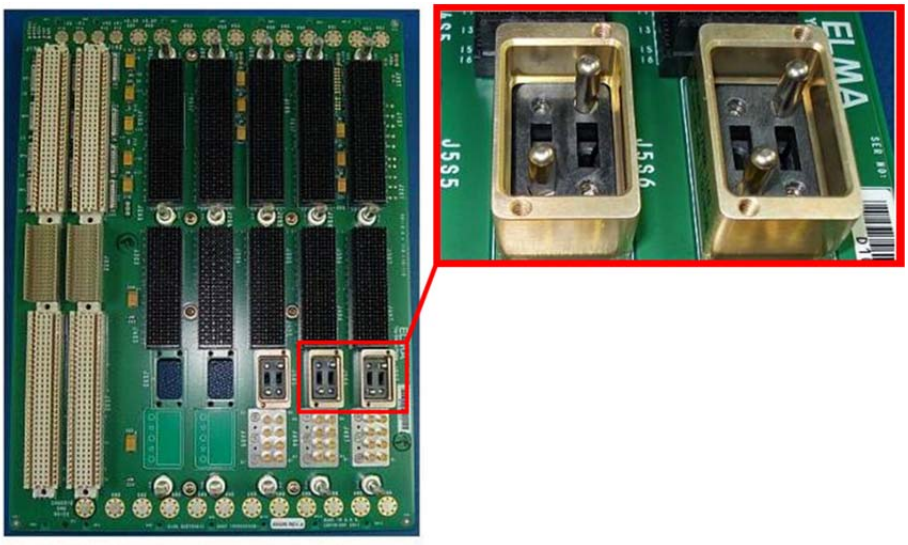

Fig. 22: Optical Backplane Implementation (source: Elma Bustronic)

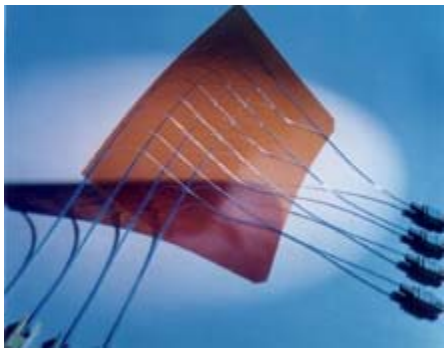

Fig. 23: Optical Backplane on Flex Circuit (source: Molex)

One challenge in achieving this higher level of integration is fiber alignment tolerance control. Consider that rack and panel backplane configurations can have mechanical offsets on the order of 1 mm, yet the optical fibers require offsets of less than +/-0.01 mm. One demonstrated solution relies on a multi-stage alignment system. The first stage (rough alignment) funnels a transceiver into a backplane alignment module shown in Fig. 24. The second stage requires the transceiver to float. This alignment brings the system within approximately +/- 0.2 mm in all directions. As the line card continues to be inserted, the third stage of alignment features engage and bring the system within approximately +/- 0.005 mm. The mechanical motion (engagement) is designed so that there is over travel which allows the line card to optical PCB insertion depth to be off by approximately 1.5 mm as shown in Fig. 25 below.

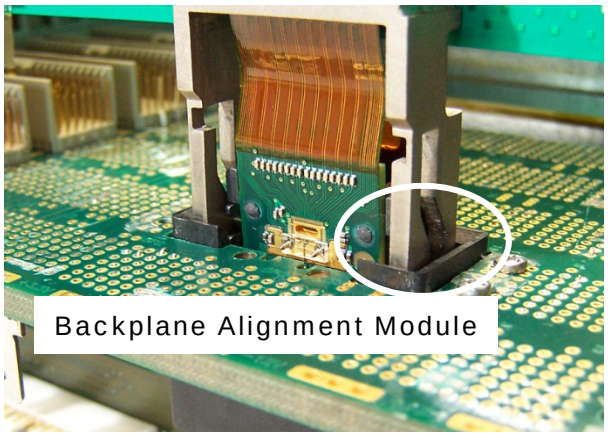

Fig. 24: Line Card Partially Inserted into the Optical PCB (source: Samtec)

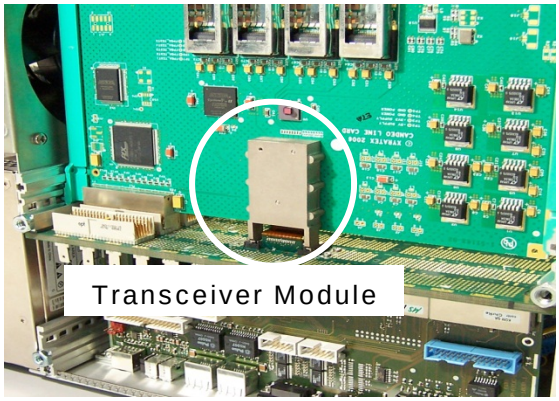

Fig. 25: Line Card Fully Inserted into the Optical PCB (source: Samtec)

5 Protocols and Interoperability

Optical engines below 14 Gbps lane speed are mostly rate independent, analog repeaters. They do not incorporate rate dependent Clock Data Recovery (CDR) circuitry and can be used in theory for transporting many different protocols at any rate – so called rate and protocol agnostic. However, there are many practical considerations that limit what electrical protocols can be successfully transported on an optical link.

5.1 Standard Optical Protocols

The simplest case is that of a protocol with a defined optical transport standard. Engine makers will list which protocols they support, and engines can be engineered and tested against the standard specification guaranteeing a successful link. The most prevalent cases are Ethernet, Infiniband and Fibre Channel. Since a complete specification is available for the QSFP+ x4 bidirectional form factor (defined in the SFF-8436 standard document), itself derived from the SFP+ 10 Gbps single channel standard (as defined in the SFF-8431 standard document), several engine vendors are in compliance with these specifications. Such engines are guaranteed to operate with standard IC chips designed for Ethernet, Infiniband and Fibre Channel with output signals that conform to the standard specification.

Another major advantage of a standard compliant engine is interoperability with traditional optical transceivers. It is therefore possible to interoperate an on-board optical engine with a remote short reach SFP+ or QSFP+ optical module. Indeed, some optical engines can be ordered with a breakout connectors to dual LC which is the standard used for 10 Gbps SFP+ transceivers (Fig. 26). One x4 engine can therefore interoperate with 4 standard SFP+, or one QSFP+.

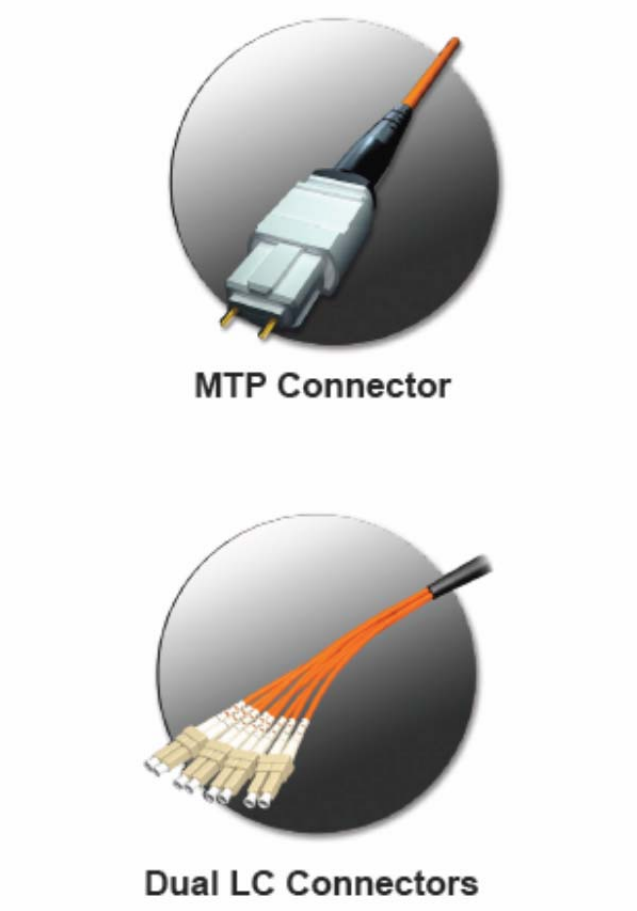

Fig. 26: MTP and Dual LC Breakout Cable Option for x4 engines (source: Samtec)

One notable exception to current Short Reach optical standard compliance is engines based on Silicon Photonics. Since they operate outside of the 850 nm wavelength band and multimode fiber type mandated by the standard, such engines can only be connected to another engine of the same make and type via the recommended single-mode fiber and subject to the length limitations set forth by the vendor. However, since the long wavelength and single-mode fibers are mandated for the Long Reach optical standards, such Silicon Photonics based engines could in theory be interoperable with standard long reach optics.

5.2 Specifically Supported Non-Optical Protocols: the PCIe Case

Some popular protocols are specifically supported by vendors ahead or in the absence of an optical standard. PCI Express (PCIe) is a case in point, and in many ways, typical of what has to be done to support a protocol with many electrical "exotics".

PCIe data signals are relatively straightforward to transport over optics: Gen3 signals are at a 8 Gbps rate, well within the current bandwidth of optical transceivers, and the data stream is balanced and scrambled, a necessary requirement for optical transmission. However, before transmitting data, PCIe links go through an elaborate link training phase that makes heavy use of Copper exotics, not easily emulated over standard optics:

- Link termination sensing
- Idle states (no signal on the Copper lines)
- Sidebands
- Clock forwarding
- Automatic link equalization schemes that assume a linear Copper link

In a closed or embedded system where a known optical link is always present, one viable approach is to disable or not use any of the exotic functionality. Forcing cable detection, disabling idle states, using remote clocking schemes, disregarding sidebands and pre-setting equalization to fixed values might enable the use of generic engines, if such options are available in the PCIe chipset on either end of the link.

But it might be the user choice to use Copper or optical, an existing design might need to be upgraded with minimal engineering change to an optical link or the chipset / design might not provide means to turn off the Copper exotics. In such cases, an engine that emulates all of the Copper exotics is desirable. In the case of PCIe, dedicated hardware resources for the termination sensing, sidebands, clock and idle support are added to the engine as shown in Fig. 27. The optical link then emulates its Copper counterpart, without the need for a dedicated optical standard, and will provide a transparent emulation of the Copper bus.

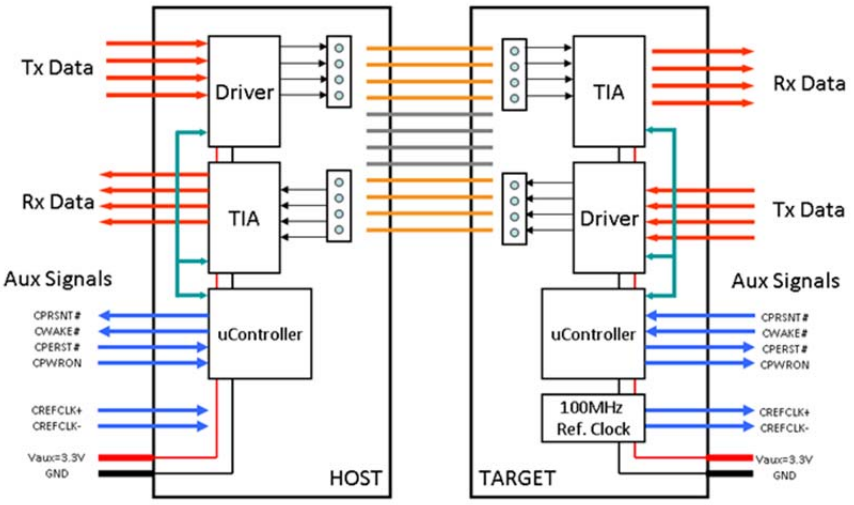

Fig. 27: Optical Engine with Dedicated PCIe Support (source: Samtec)

5.3 Generic Protocol Transport

Despite the data rate agnostic nature of most optical engines below 14 Gbps, not all protocols might be compatible for optical transport. Like in the PCIe example above, some might require additional support or special settings on the IC generating the signals to accommodate the optical nature of the link.

In order to transport a Copper protocol not specifically supported by an engine, one must usually meet a certain set of conditions:

1. Adequate bit rate and spectral content. Obviously, the bit rate must be less than the maximum supported by the engine. However and less obviously, the minimum frequency content must also be above the minimum cutoff supported by the engine. Optical links are AC coupled and have a minimum cutoff, say in the 100's of MHz. Slow protocols with long strings of zeros might have a frequency content below that limit, which will result in occasional bit errors.

2. The protocol must be DC balanced. Standard protocols transported over optics usually have a scrambling scheme that guarantees conditions 1 and 2 just described. 8/10B encoding or 63/64B encoding is such a scrambling that satisfies these conditions, and such encoded protocols are excellent candidates for optical transport.

3. The protocol must be free of Copper exotics, unless specific hardware support is available for them in the engine. We already discussed most of the troublesome ones in the PCIe section:

a. **Sidebands.** Sidebands require extra optical channels to be accommodated, or need to be encoded in the existing fibers (see PCIe example above) or be transported using extra Copper wires ("hybrid" cables). Bidirectional sidebands signals are tricky to support optically (although not impossible), as typical optical channels are mono-directional.

b. **Idle states.** No idle states (no modulated signal, electrical input to DC) should be used unless explicitly supported by the optical engine. This is because optical receivers usually contain a very high gain limiting amplifier. When no modulation is present, the limiting amplifier will go to maximum gain and basically amplify noise to the voltage rails. This will give random data at the receiver instead of the expected quiet idle signal. However, engines that support idle will instead detect the idle input state and "squelch" the receiver, yielding the correct behavior. Note that detecting the idle state and squelching the receiver takes some time, and therefore, the optical link timing will not behave exactly as its Copper counterpart.

c. **No electrical remote load sensing.** Load information is lost in the optical cable.

d. **Compensation techniques.** Pre-set or dynamic equalization meant for compensating losses of long Copper links expect a linear Copper channel. However, most optical links are highly non-linear, in particular due to the limiting amplifier in the receiver. Also, optical links will not reopen pre-compensated eyes, and aggressive pre-compensation should be avoided. The best setup is to reduce compensation settings at the transmitter so a clean eye is delivered to the transmitting optical engine. Unlike a Copper cable, an optical engine will add extremely little distance dependent deformation to a good eye – very little added jitter and eye closure will be observed between a back-to-back link and a 100 m link as shown in Fig. 28. Clean eye in pretty much guarantees a clean eye out. Treat your link as if it was only covering the distance from your source to the transmitter optics.

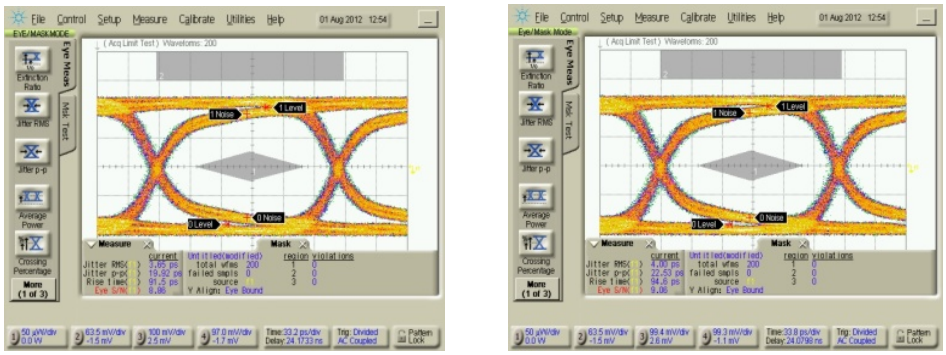

(a) (b)

Fig. 28: Typical Electrical Eye Diagram at the end of a (a) 10 m and (b) 100 m PCIe Gen 2 Active Optical Cable (source: Samtec)

Although this will probably change over the next few years, electrical end-to-end models of optical links for optical engines are not widely available. Optical engines have traditionally been specified for interoperability in terms of the standards mentioned in Section 5.1 Standard Optical Protocols. Therefore, links are mostly specified in terms of such eye mask standards as in the SFF-8431 document mentioned below. If the input signal conforms to the standard input eye mask within the required jitter limits, then the optical link will guarantee that the output passes the output eye mask, with a maximum specified degradation such as added jitter and eye closing. In general, the eye degradation through the link is extremely small compared to copper links. Additionally, some engines are bit error tested according to the standards, generally guaranteed better than 10^{-12} BER for open links (with multiple optical connector in the link) and sometimes up to 10^{-15} BER for closed links (Active Optical Cables). Of course, other parameters such as S-Parameters, input and output amplitude, receiver sensitivity, optical parameters, receiver normal and stressed sensitivity are also part of a typical specification. An example of the type of eye diagrams specified in SFF-8431 is given in Fig. 29. Mostly though, if the protocol input is within the boundary of the input mask, the output eye can be guaranteed. Running the engine at lower than maximum speed significantly increases the link and eye margins.

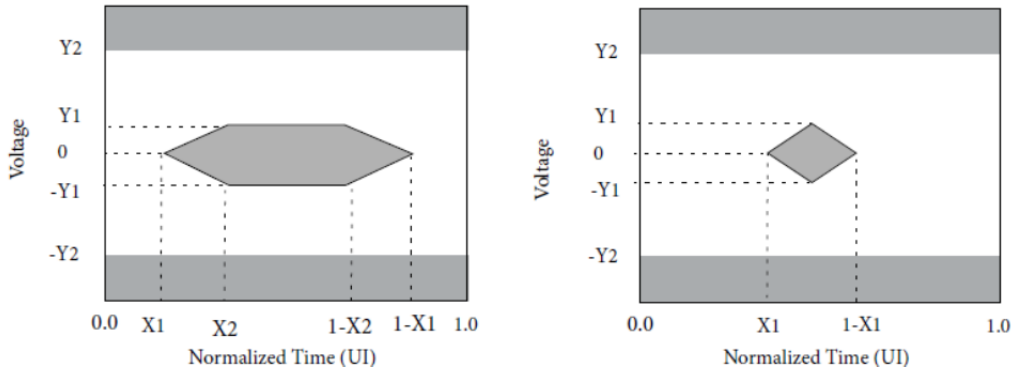

Fig. 29: Engines are Typically Defined by Input (left) and Output (right) Eye Masks (source: SFF-8431)

5.4 Interconnecting Engines with Traditional Optical Transceivers

As we have already alluded to in Section 5.1, the ability to interoperate an optical engine with a traditional standard optical transceiver (mostly short reach SFP+, QSFP+ and SNAP12) might be necessary in some applications. All engines do not support this feature equally. First of all, engines that use Silicon Photonics technology are incompatible with short reach standards. That's because they use 1550 nm, so-called long wavelength laser technology, and single-mode fibers. Therefore, such engines can only be connected to a similar single-mode based, long wavelength engine.

VCSEL based engines may or may not be optically compatible with an open standard. The ones that are compatible will not only have their electrical characteristics follow the standards, but also their optical characteristics: wavelength, power, fiber type, optical spectral width, optical eye, optical modulation, receiver stressed sensitivity, etc. With such engines, the optical output can be transported out of the box and connected straight into a remote QSFP through a regular multimode fiber plant.

For interconnections between a standard compliant engine and multiple single channel SFP modules, a breakout cable such as the one presented in Fig: 26 is a convenient solution.

6 Power and Control

Most engines require dual power pins that are individually filtered, often at 3.3V. One power pin powers the transmitter section while the other powers the receiver section. Individual filtering eliminates crosstalk between the high signal transmitter section and the sensitive low signal receiver section. For example, the SFF-8436 specification recommends the filtering circuitry shown in Fig. 30. In this particular case, a third power supply pin for the digital control section is also provided, although in many engines this is not present, nor necessary.

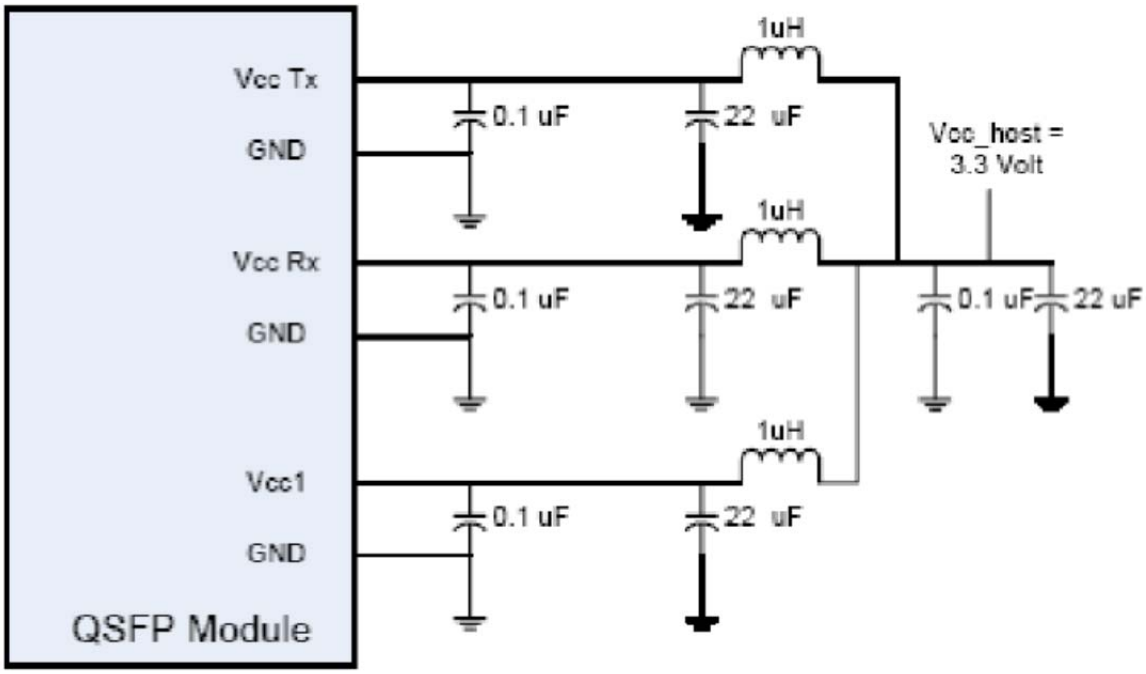

Fig. 30: Example of Recommended Power Supply Filtering
(from SFF-8436)

Some engines also incorporate elaborate digital control and monitoring via an on-board microcontroller. Usually communicating via I2C, information such as engine temperature, voltage, transmitted optical power, received optical power and loss of signal can be queried from such optical engines. Engine parameters such as output pre-emphasis and voltage can also be altered via this interface. Enabling or disabling laser transmission per channel, as well as squelching the receivers per channel, is also implemented in some engines. The SFF-8436 "memory map" standard which has defined the address and value of many registers for this purpose is often used as a template for I2C control for quad channel devices.

7 Upcoming Technology Standards

The march of technology and climb to higher bit rates never seems to cease, and standards are continually evolving to meet the next technology challenge. The old standards get updated for higher speeds, while many formerly Copper-only protocols are seriously considering the addition of an optical version to their standard.

Below is a partial list of standards being discussed as of this writing:

1. The Infiniband standard has just completed its FDR revision, bringing the rate up to 14 Gbps. One important addition besides the speed increase is a standard way of changing output voltage ranges via I2C control.

2. IEEE 802.3bm is in the process of establishing the successor to QSFP+ for 100 GbE. 28 Gbps x4 is being proposed as one of the many options to transport 100 GbE optically over short reaches. An important point of discussion is the re-introduction of CDR circuitry (and its associated cost and power consumption) to support such high data rates.

3. SAS has adopted an option for optical transport and created a variant of the OOB sideband processing that is compatible with optics.

4. Thunderbolt has introduced an optical variant of their consumer high speed 10 Gbps link.

5. PCIe is discussing the incorporation of an optical standard in their OCuLink (read "Optical Copper Link") next generation PCI Express cabling standard.

I am eagerly expecting many more announcements of protocols to be standardized over optics as bit rates continue to go up!