# INTERCONNECTS TWG

Contents

## EXECUTIVE SUMMARY

### INTRODUCTION

The Interconnects Technology Working Group (TWG) focuses on the technologies needed for realizing optical connections between integrated photonic components in computer servers and other equipment used in hyperscale data centers and other high-performance data communications applications.

There is significant strategic importance to this market due to replacement of discrete server systems used in millions of businesses and institutions by large, independent, Cloud Data-As-A-Service (DAAS) providers; Hyperscale Data Centers (HDCs) are becoming THE key elements of our future information technology infrastructure. A large infusion of photonics will be required to handle the Zettabytes that will be processed in these data centers and communicated to millions of end users.

Three main application domains can be distinguished:

- Optical communication at the Integrated Circuit (IC) package substrate or printed circuit board (PCB) level, connecting electronic/photonic components in a single package, or connecting bare electronic and photonic IC's (PICs) mounted on a substrate and communicating across the PCB or to/from the IC substrate
- Optical connectors that are used for realizing an optical connection between two optical fibers, or a fiber and a substrate waveguide
- Input/Output (I/O) ports consisting of board-edge or mid-board optics connecting to other servers, switches, and system networks in a data center or in other (remote) applications

About four decades ago optical interconnect technologies were introduced in outside-plant long-haul telecommunications networks. Since then, the use of fiber-optic interconnect technology has grown strongly and moved closer and closer to the end user. Optical fiber is now THE medium for the transfer of large amounts of data both indoors and outdoors, in e.g. local area networks, fiber-to-the home, local loop telecom and data center networks.

In present-day hyperscale data centers, many of the copper-based rack-to-rack datacom links are being replaced by broadband optical communication links, using pluggable optical modules at the edges of the server boards e.g. fiber optic IO connectors and cable links. In the coming years, as data rates continue to increase past 100 and 400Gbps, the broadband board-level and intra-rack copper-based interconnects will need to be replaced by single-mode fiber and package-level optical communication links due to speed/bandwidth and power consumption advantages over copper.

Drive for further expansion of the optical interconnect technology domain results from the needs of novel applications in markets like the sensor, medical, autonomous vehicle and 5G markets, which have started to apply integrated photonic technologies (e.g. Light Detection and Ranging (LIDAR) sensor systems for automotive guidance and collision avoidance). These new markets and developments in data communication systems, require the development of new photonic interconnection technologies for connector and substrate interconnect applications. The development track of these new interconnection technologies is described in this roadmap.

Fiber-optic connector technology is well advanced, having a 30-plus year history of development and manufacturing for a wide range of applications and densities. The current optical (data) paths to/from a board make use of arrays of optical fibers which are connected to photonic transceiver components at the edges of the board. The intra-board communication and the backplane communication is handled by electrical connections that are embedded in the boards/backplanes and accessed via backplane connectors.

The current generation of optical connectors are relatively costly (typically > 10X the cost of copper) and are sensitive to dust and other contaminations, thus requiring time consuming and tedious cleaning and inspection at each mating cycle. To avoid such burdens, and to reduce mating-induced damage by eliminating glass-to-glass contact, a special type of fiber-optic connector -- the expanded-beam (XB) connector -- is being introduced. The XB connector employs a small optical element (e.g. a lens) to eliminate glass-to-glass contact and broaden the optical beam to reduce the effects of contamination and eliminate the labor-intensive practice of having to periodically dismount and clean cable and IO connector end faces.

Unfortunately, the interest in XB connectors, at least for chip interfaces, runs counter to another need: that of higher-density (waveguide cores/area) interconnects. Today's connector technology typically uses 250 micron spacing from fiber core to fiber core. However, channel waveguides in PICs can be fabricated at 10s of micron spacing; spacing them at 250 microns wastes valuable semiconductor real estate. Therefore, approaches to tighter-pitch fiber cores, or off-chip waveguide fan-outs are needed.

To cope with the ever-increasing board-level data rates and the optical interconnection needs of future photonics applications, optical (embedded) interconnections need to be more broadly introduced at the board level. These interconnections will replace: (1) discrete optical-fiber-based connections between (electro/) optical components on a single substrate and (2) copper interconnects over relatively short distances. The driving forces for pursuing these goals are to achieve: higher bandwidth*distance product, reduced power dissipation, reduced noise and crosstalk, improved signal data rate density (Gb/sec/cm$^2$ of board edge area), reduced latency and reduced interconnect form factor.

Optical interconnect advances are needed in the field of board- and rack-level fiber-optic connectors, for overcoming such issues as dust contamination and damaged fiber interfaces, and for reducing or preventing costly cable cleaning and inspection procedures in the field. One potential approach to this is the development and high-volume deployment of Expanded-Beam (EB) Single Mode (SM) fiber optic connectors.

In future equipment applications, optical fiber may be supplanted by channel waveguides fabricated in planar substrates such as PCBs or specialized interposers. These waveguides would allow compact fiber-free optical connections between photonic chips, e.g. to the micro-miniature silicon waveguides comprising silicon photonic ICs. Key issues associated with the use of channel waveguides in PCBs (and in the PICs) include achieving low optical loss for coupling to other fibers/waveguides and dealing with potential polarization dependence of the properties of the waveguides. For realizing photonic connections between optical waveguides in PCBs and interposer/MCMs the use of sockets is currently seen as the most straightforward approach.

To enable wide-scale deployment of substrate-level optical interconnects, a combination of technology gaps and related infrastructure gaps must be addressed for the following topics:

- ergonomic, low-mating-force, environmentally-robust, contamination-resistant, low-optical-loss connectors for SM systems
- low-loss, easily-terminated polarization-maintaining (PM) connectors and chip interfaces, for use in single-polarization and polarization-encoded photonic circuits
- high-volume, low-cost manufacturing approaches for optical connections in packaging of Photonic Integrated Circuit (PIC) devices, with approaches to achieve tighter channel pitch
- solder-reflow-compatible packages and connectors for integrated optics modules
- an optical coupling component "tool kit" or process design kit (PDK) providing low-loss coupling components suited to a wide variety of foundries and PIC designs (for chip edge, surface grating and adiabatic evanescent coupling to PICs)
- low-loss substrate-embedded SM waveguides, compatible with solder reflow, that can eliminate the need to manually install fly-over fiber optic cables on board-mount modules

## NEEDS

Needs < 2025

- > 16 fiber SM expanded-beam connectors, for cables and modules with manual termination
- SM expanded beam connectors, for optical backplane connectors, frontplane, and midplane applications.
- 24 port module optical interface using pigtailed, 1-D waveguide array pitch converter
- Module-board attachment via socket.
- PM-fiber-based connectors, jumpers, and breakouts for polarization sensitive systems

Needs 2025-2030

- > 32 fiber SM expanded-beam connectors, for cables and modules, allowing semi-robotic termination
- Low-loss optical waveguides for integration in PCBs (<0.2 dB/cm)
- Simplified module-to-PCB waveguide coupling, e.g. via evanescent coupling, gratings, or turning mirrors
- Module-board attachment via reflow<128 port module optical interface using multicore fiber, 1D waveguide/microlens array, interposer

Needs 2030-2040

- > 64 fiber SM expanded-beam connectors, for cables and modules, allowing robotic termination
- >128 port module optical interface using multicore fiber, 2D waveguide/microlens array, interposer

## INTRODUCTION

Historically, optical interconnection had its first great success in long-distance communication applications where the bandwidth and cost benefits were compelling, but more recently it has been applied to shorter and shorter distances, thanks to the increase in channel speed that drives the cost of electrical interconnects higher than that of optical interconnects. In some cases of intermediate length, while it is still technically feasible to transmit electrically, the cost of repeaters and signal conditioning becomes too great compared to optics which has a lower cost slope with distance. Furthermore, applications outside traditional communications have begun to proliferate; these include applications in various types of chemical, biological, and physical sensing, as well as specialty applications like optical control systems for phased-array radar, etc.

Fiber optic interconnect has been widely deployed in long-distance communication applications, such as undersea cables and metropolitan networks. In these cases, the low loss and high distance*bandwidth product of fiber, combined with the relatively small number of transceivers per fiber length (and the availability of fiber amplifiers) make fiber the most cost-effective transmission medium.

As data rates in and between large numbers of racks in supercomputers, data centers, and telecommunications switches have increased, optical interconnects have become the low-cost solutions for rack-to-rack interconnects. Now, OEMs also recognize the paradigm shift from Cu to optical circuitry for short-distance interconnect at locations deeper into the rack addition to high bandwidth, now include the potential for lower power dissipation and consumption, improved noise immunity, smaller size and weight, lower latency and easy compatibility with future system speed upgrades.

However, the economic requirements of long- and short-distance applications are entirely different. In long-haul fiber, the overall interconnect system cost is dominated by the installation of the cable, so that the use of expensive connectors and transceivers is possible. In the case of substrate-level connections inside and between racks, this is not the case; cost of higher-density optical connectors and transceivers dominates the cost of the interconnect system. To date, the excessive cost of these components has been a barrier to wide replacement of copper by fiber in short-reach (< 5m) communications.

Penetration of fiber and integrated optics into sensing systems is driven by several potential benefits of optics relative to conventional approaches. These include: high sensitivity, small size, EMI immunity, electrical isolation capability, and the potential for high levels of integration. Market applications include avionics (gyroscopes), structural monitoring (architectural and mobile platforms), medical sensors, food safety, environmental monitoring, and electrical utility and industrial sensing. Some specific examples of current and planned applications of fiber optic sensors include temperature sensing in oil wells, and mechanical stress sensing in structures such as bridges (both using fiber Bragg gratings), current and voltage sensors for power utilities, fiber gyroscopes for aircraft and weapons, and ring resonator chemical and biological sensors for medical diagnostics, food safety, and chem-/bio-warfare applications.

In all these sensing applications, Photonic Integrated Circuits (PICs) are either used directly as the sensing transducer, or are convenient, compact and cost-effective approaches for building the optical system to e.g. interrogate a fiber sensor. PICs are able to perform a variety of optical functions such as: optical emission; modulation (analog or digital) of optical intensity, phase, polarization, and wavelength; detection of optical intensity, phase, polarization, and wavelength; wavelength multiplexing and de-multiplexing, or splitting or combining of optical intensity or polarization.

## SITUATIONAL (INFRASTRUCTURE) ANALYSIS

In the following sections, optical interconnection technologies will be considered for applications ranging across analog and digital communication links in tele/data communication systems, sensor systems, phased-array antennas, medical systems, LIDAR in the automotive industry, and automated manufacturing systems.

For applications outside communications, interconnect systems requirements are generally similar to the communications applications, that is, determined by the distance of optical propagation involved. However, there are areas where the different types of application requirements diverge, such as:

- applications like medical/bioweapons sensing and radar requiring high robustness which may not be subject to the same cost constraints as typical communication applications
- many sensing applications which require the integration of microfluidic elements, generally not required in communications applications
- many sensing and imaging applications may require wavelengths of operation outside the standard tele/data comm 1.3/1.5 μm bands

The situation analysis for interconnects in communication systems is given below. This analysis is also used for describing the situation for all other applications.

Application areas will be classified by the length of the optical link as follows:

1) rack-to-world connections (i.e. LAN and telecom network connections)
2) rack-to-rack connections
3) inter-blade connections between blades in a single rack (potentially through an optical backplane)
4) intra-blade connections between modules on a single blade
5) intra-module connections within a module (e.g. between different chips in a module)

***Application area 1: Rack-to-world connections (lengths > 500m)***

Optical links for distances over 500 meters are common today. They are typically based on pluggable optical modules having various data rate, wavelength and distance specifications; the modules are plugged into sockets at a board edge. These provide a convenient optical connection, for either single Transmitter/Receiver (TX/RX) pairs multimode (MM) or single mode (SM) fiber, and may use multiple parallel groups of fibers or wavelength multiplexing to increase capacity. However, since these pluggable modules are placed at the edges of PCBs, signal degradation results from the copper traces that transfer signals from mid-board electronics to the board-edge transceivers. Thus, there is a drive to place the electrical-to-optical conversion (the transceivers) mid-board near the signal sources (on-board optics) or more likely on the same substrate in a module ("co-packaged optics"). An added benefit of having an optical connector at the board edge rather than a copper connector is the improvement in data connection density (Gb/s per board edge utilized) that can be achieved. This is because copper connectors, to avoid crosstalk and maintain signal integrity, have larger channel spacing than do optical connectors. Furthermore, this benefit can be multiplied by wavelength multiplexing to improve the data capacity of each optical fiber.

There are different benefits and issues with the configurations in which the TX/RX components and the electronic ASICs are packaged in the same or different modules. Co-packaging provides the shortest copper path for the high-speed signals between the different components but requires a high degree of customization dependent on the details of the ASIC. Separate packaging of the TX/RX and ASICs allow a more modular "mix & match" design approach with simpler component and module design, at the cost of longer copper connections.

In future systems, the number of signal channels and the bandwidth per channel will dramatically increase, necessitating the use of high-bandwidth, large-channel-count mid-board transceivers. These transceivers are likely to be integrated into multi-chip electronic/photonic signal processing modules to achieve per bit cost-scaling reductions. Such a mid-board signal processing module with a relatively complex electronic/photonic system will consist of multiple photonic and electronic ICs which will use an interconnect substrate for realizing the optical and electronic interconnections between the ICs. To facilitate optical connections at the edge of the board, embedded high density optical interconnects will be needed.

***Next-generation technology for Application Area 1***

- Pluggable mid-board or co-packaged SM modules, to reduce copper trace length and related impairments (AA1.1)
- SM dust- and damage- tolerant connectors (e.g. expanded-beam) for modules. (AA1.2)
- Low-mating-force, dust resistant, high-density front panel connectors (AA1.3)
- EMI management and agency certification testing (which is more difficult with each increase in speed), and insertion loss standards, including understanding of variations across component suppliers (AA1.4)

**Table 1.** Technology development topics for Application Area 1 (rack-to-world).

| | AA1.1 Electrically-pluggable mid-board SM module | AA1.2 Expanded-beam SM connector for module interface | AA1.3 Low-Loss, dust-resistant, front panel connector | AA1.4 Measurements and Standards for EMI and Loss |
|---|---|---|---|---|
| Manufacturing processes | X | X | X | X |
| Manufacturing equipment | X | X | X | X |
| Materials | | | | X |
| Quality/Reliability | X | X | X | X |
| Environmental technology | | | | X |
| Test, Inspection Measurement | | | | X |
| Attenuation | X | X | X | X |
| Density | X | X | X | |

For rack-to-world applications, single-mode fiber networks are now preferred because of the high bandwidth*distance capability of the fiber. At current data rates (≤100 Gbps per channel), it is possible to mount pluggable transceivers at the edges of PCBs, to make access and replacement easier. However, as channel speeds increase, there will be a drive to move transceivers away from the edge of the boards or to co-package transceivers with other ICs to shorten the copper traces between transceivers and signal sources and thereby reduce signal impairment and loss. These mid-board transceivers will need to be compact (to save valuable PCB area), reliable (to save downtime and replacement costs), dust resistant (for easy installation and maintenance), and offer low cost in terms of $/Gbps (to allow scaling to very high throughput).

The desired increase of the aggregated data rate of such modules, both driven and enabled by the increased capacity of CMOS switches, Field Programmable Gate Arrays (FPGAs), and Multi-Chip Modules (MCMs) , will be addressed in multiple ways. One approach consists of simply increasing the data rate capacity of each fiber (using WDM or advanced modulation formats), another consists of an increase in the overall number of parallel optical channels (Space Division Multiplexing) serving each transceiver. There are even attempts underway to use different optical modes of multimode fibers to carry different signals.

In addition to the use of multi-fiber ribbons, denser optical IOs at the chip interface, can potentially be achieved by using multicore optical fiber, development of which is now progressing rapidly. In such fibers, multiple cores are arranged in a regular geometric pattern, typically a linear array or a hexagon. The cores are spaced by 10s of microns, and thus can be compatible with the pitch of high-index waveguide arrays used in semiconductor PICs. However, there are many practical issues that must be addressed before wide deployment of multi-core fiber. First, maintaining the geometric accuracy of the core locations required for adequate alignment for low-loss coupling to another fiber or waveguide is difficult (such fiber has not reached the high geometric precision of single-core SM fiber yet). Second, the rotational alignment is now critical (more so than in PM connectors). Third, there can be optical crosstalk between the cores. Fourth, for the case of end-fire coupling, 2-D core patterns in the fiber require 3-D fanouts on the PIC, which are difficult to fabricate (in the case of surface coupling, e.g. using gratings, accommodating 2-D core patters with 2-D arrays of surface couplers is more straightforward).

Mid-board pluggable transceivers or co-packaged modules can fit well into the rack-to-world application, since optical interconnect is already accepted and cost sensitivity is moderate so that manual installation of transceivers and/or co-packaged modules (optical TX/RXs and ASICs in a single package) is feasible. However, reduction in module cost is always desirable. Since today's module cost is driven by packaging cost (dominated by fiber alignment, attachment, and testing) a critical area for technology development is in packaging. Today, module manufacturing functions are performed manually or semi-manually, mainly because of the tight mechanical alignment tolerances required for efficient optical coupling. In the future, new processes and equipment for rapid and automated alignment, attachment and testing of fibers or other optical interface components for semiconductor waveguide devices are needed. Of course, for the connector parts used in these modules to provide reliable performance at low cost, new tooling and processes for achieving the required tolerances in molded parts will be required.

A key metric in interconnect evolution is interconnect density, often characterized in Gbps/cm$^2$ of board area. The overall density of the interconnection is determined by the module size, which is in turn often limited by the size of the module optical and electrical connector interfaces (today's array connector ferrules are typically larger, and have larger channel pitch, than the chips to which they interface). Thus interconnects which can perform a pitch transformation from the chip to the connector, and tighter-optical-channel-pitch connectors are both needed enabling technologies.

Module packaging technologies, materials and processes are intimately linked to reliability, since a fundamental element of reliability is the demonstration of low optical loss which is stable across operating and storage conditions, and loss is in turn affected by both the accuracy and the stability of fiber alignment and attachment.

There are several approaches to address future needs for increased optical I/O count, small-form-factor low-profile packages, and manufacturability (e.g. compatibility with solder reflow processes allowing SMT technology to be used to assembly the module on the PCB).

A conservative approach is to rely on moving from legacy approaches that use fiber arrays bonded into v-grooves and actively aligned and butt coupled to devices, to higher I/O count v-grooves holding multicore fiber or reduced-diameter fiber (to permit reduced v-groove pitch).

One developing approach is a waveguide-assisted coupling configuration using a short intermediate waveguide array to connect the PIC to a multifiber connector. This may be a glass or polymer waveguide array, achieving low loss coupling (for example using evanescent coupling between the intermediate waveguide and the PIC waveguide), pitch conversion, and eventually vertical and horizontal as well as lateral redirection of the beam. In this case, a way of achieving self-alignment of the waveguide to the PIC is crucial for cost-effective manufacture.

Another approach proposed to lower system assembly cost and higher reliability is to use expanded-beam connectors at the module and front panel connections. These connectors can relax the mechanical alignment tolerances required, lowering assembly cost. They can also, by virtue of their non-contact, expanded beam coupling, provide reduced sensitivity to dust and damage, therefore providing higher reliability. And they have low mating force so can scale to higher fiber channels than physical connect connectors that rely on Hertzian contact stress to slightly deform the connectors to fully close the gap between the fibers. Unfortunately, low-loss expanded-beam connectors have been difficult to realize, especially for use with single-mode fibers. This is largely due to the difficulty of molding the complex optical polymer ferrules with the fiber holding structures adequately aligned to the beam-expanding optics. For expanded-beam connectors to become practical, new tooling and processes for molding optical materials to higher tolerances must be developed. At the PIC side, this approach will require micro-lens arrays (1-D or 2-D) to be accurately aligned to waveguides to provide an enlarged collimated beam. However, the size of the optical elements, constrained by the size of the beam expansion desired, can impose a limit on minimum pitch for both edge and surface coupling, and on the minimum real estate dedicated to coupling in surface coupling approaches. Ideally, the lens-to-PIC alignment would be done passively. Several techniques can be used, from die-to-wafer assemblies of micro lens arrays, to wafer-scale fabrication of micro lenses directly on top of the PIC.

In a last approach, one or several PICs may be optically coupled to a common larger-dimension photonic interposer (made of glass, Silicon On Insulator (SOI), or organic laminate). The photonic interposer provides optical routing between the PICs via embedded optical waveguides and may also provide pitch conversion and optical coupling to an edge connector or to a motherboard. Electronic integrated circuits (EICs), for example ASICs such as Ethernet switches, can be mounted on the same interposer; such an arrangement of EICs and PICs is also referred to as “co-packaging”. The use of these co-packaged or multi-chip modules allows higher shorter Cu connections between the PICs and EICs, thereby improving signal integrity. Co-packaged optics is the end goal for maximizing the reduction of cost, power and size of mid-board optics modules. The adoption rate and timing of co-packaged modules will depend on the engagement of, and standardization push from, the mega datacenter companies.

One issue with existing multifiber connectors, whether of physical contact or expanded-beam design, is the cost of terminating fibers in the ferrules. Today this process is performed manually; in the future equipment and processes to achieve automated low-cost, high-throughput termination must be developed.

Figures 1a and 1b below show schematic configurations for mid-board modules. In Figure 1a, the optoelectronic transceiver and the ASIC are both mounted on a system PCB. In Figure 1b, the optoelectronic transceiver and the ASIC are mounted on a separate substrate that forms the base of a package. This second configuration has advantages of shorter copper connections, and more options for optimized thermal and electrical characteristics of the substrate.

(A) Pluggable mid-board SM/MM modules, to reduce copper trace length and related impairments
(B) Expanded-beam SM/MM connectors to reduce precision requirements and contamination sensitivity
(C) Low-mating-force, dust resistant, high density front panel connectors

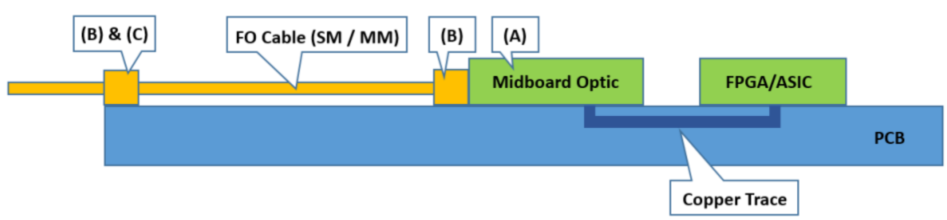


**Figure 1a.** Technology approach for next generation systems in Application Areas 1 and 2.

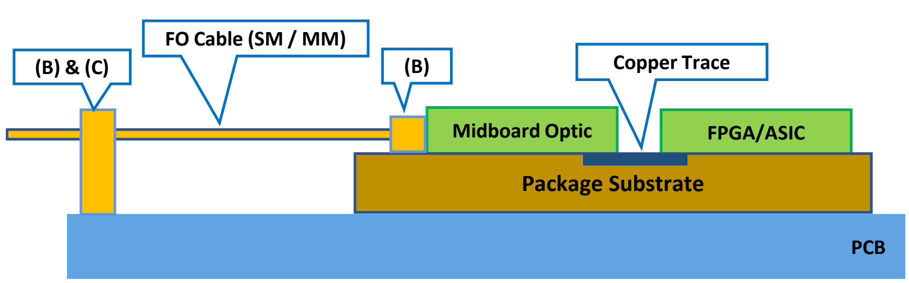

**Figure 1b.** Co-packaged version of Figure 1a, for Application Areas 1 and 2.

### *Application Area 2: Rack-to-rack connections (lengths 500 - 5m)*

These are also common today. As in the case of rack-to-world interconnects, pluggable transceivers are typical implementations at the moment, but Active Optical Cables ("AOCs" which are fiber cables having transceivers permanently attached to each end, thereby easing internal optical component interaction requirements), are also implemented for short run applications where cable routing with transceivers attached is not too cumbersome. However, mid-board modules and interposer-mounted optical modules with transceivers are anticipated in the future.

### *Next-generation technology for Application Area 2*

- Pluggable mid-board SM modules, to reduce copper trace length and related impairments (AA2.1)
- Front panel and blind mating expanded-beam SM connectors to relax contamination and mating damage sensitivity (AA2.2)
- Low-mating-force, dust resistant, high-density front-panel connectors (AA2.3)

**Table 2:** Technology development topics for Application Area 2 (rack-to-rack)

| | AA2.1 Pluggable mid-board SM&MM module | AA2.2 Expanded-beam SM&MM connector | AA2.3 Front panel / back-plane connector |
|---|---|---|---|
| Manufacturing processes | X | X | X |
| Manufacturing equipment | X | X | X |
| Materials | | | |
| Quality/Reliability | X | X | X |
| Environmental technology | | | |
| Test, Inspection Measurement | | | |
| Attenuation | X | X | X |
| Density | X | X | X |

For rack-to-rack connections, the development topics are similar to those for rack-to-world connections, with 3 notable exceptions.

First, since the connection lengths are shorter than rack-to-world connections, multimode fiber used with a VCSEL-based transceiver has adequate bandwidth*distance performance to become a viable option to meet today's requirements. Because of the relative ease of packaging MM VCSELs vs SM integrated photonic transceivers (because of the relatively large, well defined emission area), relative cost benefits are possible. For this reason current rack-to-rack optical connections are dominated by MM VCSEL-based optics. However, low-cost wavelength multiplexing is much more difficult with MM fiber so system bandwidth does not scale as easily as with SM transceivers. This is expected to drive the implementation of SM, WDM-based silicon photonic links in the longer term.

Second, because of the much higher link count in rack-to-rack compared to rack-to-world connections, the cost of the transceiver and connector components is much more important. This amplifies the need for new low-cost manufacturing equipment and processes.

Third, because of the high channel count at switch boards in data centers, the number of fibers leaving the board can be very high, so that the areal interconnect density (fibers per vertical area at the board edge) becomes very important. This is not just a matter of needing room for the connectors, but also due to the need to maintain open area for flow of cooling air.

Considerations of loss and reliability are similar to those mentioned in Application Area 1. However, there is some interest in enhanced-reliability transceivers for use in Application Area 2, because of the very large number of transceivers anticipated in a single system, thereby increasing the probability of there being a transceiver failure somewhere in system.

***Application Area 3: Inter-blade optical connections (length 5 – 0.5m)***

Inter-blade (but intra-rack) optical communication is receiving a tremendous amount of current interest, and is the subject of many development programs, especially in systems which are designed for longer useful lifetimes with several planned "speed bump" upgrade cycles. This is because of the very high port count required (some blades will need over 1,000 optical connections), which can leverage the signal density benefits of optical interconnect. There is therefore a large market opportunity for successful product development. In the short term a density increase in inter-blade interconnections can be facilitated via the application of high-density multi-fiber connectors and /or the use of multicore optical fibers. Longer-term solutions for high-density inter-blade interconnections will require high-density embedded optical waveguides and high-density optical connectors interfacing to an optical backplane.

For rack-to-world, rack-to-rack, and inter-blade optical connections alike, one key issue is the location of the optical module on the PCB. Electrically-pluggable board-edge connections such as AOCs, have the advantages that they are easy to design in, add later for upgrading capacity, and replace when needed (i.e., hot swappable). However, they require copper traces to extend to the board edge and introduce another copper connector in the signal path (the traces and connector both contributing associated signal impairments) and can impede air flow. Moving the optical module to mid-board can reduce the length of copper traces involved, but unless the module is soldered to the PCB, this approach still introduces another connector. Furthermore, making fiber cable connections to a mid-board module can be cumbersome and laborious, so that having waveguides embedded in the PCB to couple optical signals from the mid-board module to the board edge would be highly desirable. At this time, polymer waveguides reported for integration into a PCB do not have low enough loss at 1.3 and/or 1.55 micron wavelengths to be practical, but glass waveguides do. New Multi-Source Agreements (MSA's) such as COBO (the Consortium for On-Board Optics) and the Co-Packaged Optics Collaboration (CPO) are defining standards for low- and high-speed electrical connectors, module footprints, power consumption and interface requirements which will help to develop and accelerate use of mid-board or co-packaged optical interconnects.

One alternative to AOCs, mid-board or co-packaged modules is an optical interposer. This is a small "daughter board" that plugs into a PCB and provides a suitable substrate for the optical modules. Advantages of the interposer implementation include:

- Allows the optical modules, viewed as potentially lower reliability than the electronics, to be easily replaced if they fail
- Separates the module mounting process from the standard reflow of the PCB
- Allows the use of different, more expensive, higher-performance materials for the interposer than used in the PCB, e.g. the interposer can be a piece of silicon wafer
- Provides a shorter optical path between modules and connectors to the outside world, thus allowing the use of higher loss-per-distance waveguide materials (this can relax the waveguide propagation loss requirement from ~ 0.02 dB/cm to cross a blade, to ~ 0.2 dB/cm to cross an interposer).
- The interposer-to-PCB interface socket could be standardized, thus separating PCB design from the details of the optical modules

Disadvantages of the interposer implementation include:

- Localization of the optical modules on the interposer requires longer copper traces on the PCB to reach the interposer socket. While this is still better than the case of a board-edge pluggable module, it is not as good as the module location being unconstrained
- The electronic interface between the interposer and the PCB will introduce additional signal degradation
- There are now two "boards" to fabricate separately: the PCB and the interposer
- Plugging the interposer into the PCB, and potentially connecting the output fibers cable(s), requires more labor than the reflowable integrated-waveguide PCB

***Next-generation technology for Application Area 3***

- Standardized mid-board, co-packaged or interposer-mounting optical modules with fly-over fiber-based media (AA3.1)
- Optical embedded waveguides including optical interfacing to an optical backplane or a front panel (AA3.2)
- Optical backplanes simplifying PCB to PCB optical routing (AA3.3)
  - Low-mating-force, high-density low-loss, low-cost, dirt-resistant expanded-beam multimode and single mode front panel, backplane, and midplane optical connectors (AA3.4)

**Table 3.** Technology development topics for Application Area 3 (blade-to-blade).

| | AA3.1 Pluggable mid-board SM&MM module | AA3.2 Optical embedded waveguide | AA3.3 Optical backplane | AA3.4 Front/ mid/ backplane connector |
|---|---|---|---|---|
| Manufacturing processes | X | X | X | X |
| Manufacturing equipment | X | X | X | X |
| Materials | | X | X | |
| Quality/Reliability | X | X | X | X |
| Environmental technology | | | | |
| Test, Inspection Measurement | | X | X | |
| Attenuation | X | X | X | X |
| Density | X | X | X | X |

From the module standpoint, the technology development needs are the same as Application Areas 1 and 2, apart from even stronger pressure to develop manufacturing equipment and processes that can drive module manufacturing cost down.

Many connector technology requirements are also similar to those mentioned for Application Areas 1 and 2, except that now, for easy routing of high-speed signals between blades, optical backplanes and optical backplane connectors will be required. These components need to function like copper backplanes and connectors, allowing blind mating and being resistant to dust that may accumulate at un-mated connectors in vacant blade locations. One issue with today's backplane connectors for optical fiber is that they typically have very high mating force, and thus are not suitable for very-high-fiber count applications, due to the needs for increased rack and card mechanical load bearing robustness.

Future optical backplanes may be based on optical fibers, optical fibers routed on a flexible substrate, or on embedded optical waveguides: that is, channel waveguides fabricated in a substrate such as polymer or glass. If they are fiber-based, manufacturing technology for automatically routing and terminating them in connectors is needed. If they are channel-waveguide based, new connectors for channel waveguides will need to be developed.

An alternative for fly-over cables are mezzanine card with mezzanine card connectors, which provide higher levels of integration at PCB level which may reduce the complexity of electro-optical packaging.

There are indications that optical cabling or flexible embedded optical circuitry might affect the traditional motherboard/daughtercard backplane domain – particularly when future fiber optics becomes dominant, and board-level electronics is shrunk to module-level. Several potential technologies including the use of laminated polymer or glass optical waveguides embedded into a conventional backplane have been investigated; an example is shown in Figure 2 below [1].

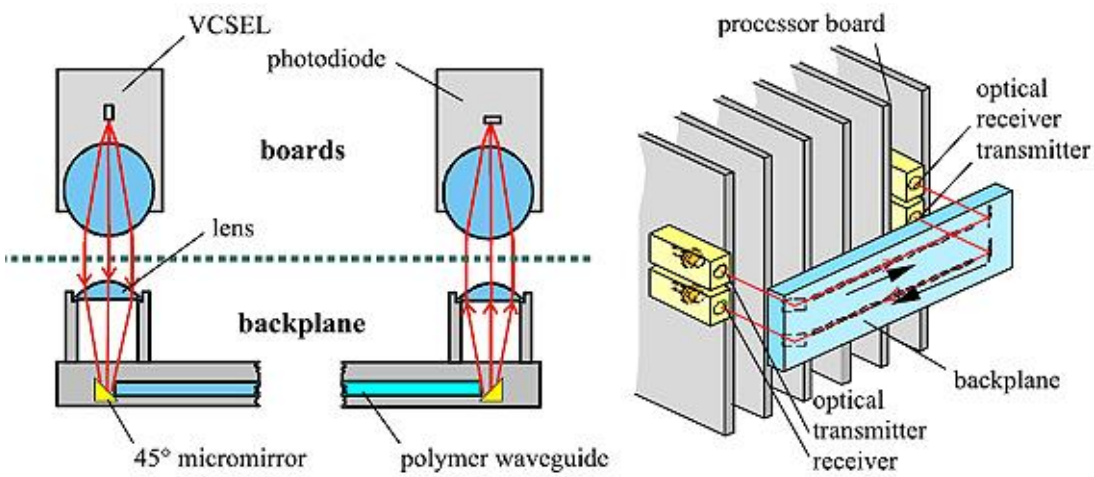


**Figure 2.** Approach for optical backplane connector technology.

The ability to tap the optical layer within a PCB and re-direct an optical signal 90° up into a connector has been a difficult challenge. Recent publications [2] indicate that development work on true optical backplanes is continuing as new technology becomes available.

In systems with many short-range optical connections based on fiber, one of the most significant problems will be the routing and management of larger numbers of fiber cables. For that reason, there has been a long-term drive for the development of waveguides that can be embedded in the blade or backplane, thus eliminating the fiber management problem. Unfortunately, to date, there is no published technology for fabricating embedded waveguides fully satisfying the requirements of low loss over the application lifetime, compatibility with solder reflow (260°C), and having low-loss coupling features for surface mount photonics . For waveguides embedded in PCBs, if communication across a board in a standard 24" x 36" rack is desired, a transmission distance on the order of 100 cm is required. For 2 dB of total propagation loss, 0.02 dB/cm waveguides are required; this value is very challenging (values near 0.2 dB/cm are more typical today). If losses cannot be reduced below 0.2 dB/cm, then the use of embedded waveguides will be limited to small (~10 cm) interposers.

For achieving lowest loss, glass waveguides fabricated either by ion-exchange or laser-writing have shown promise. This is a key area where significant technology advances in materials and manufacturing processes are required.

Note that the implementation of embedded waveguides in blades and backplanes will create new challenges in efficient testing, and will require new levels of reliability, due the high number of connections and the fact that unlike the case of fly-over fiber, a bad optical connection in an embedded-waveguide blade will not be repairable.

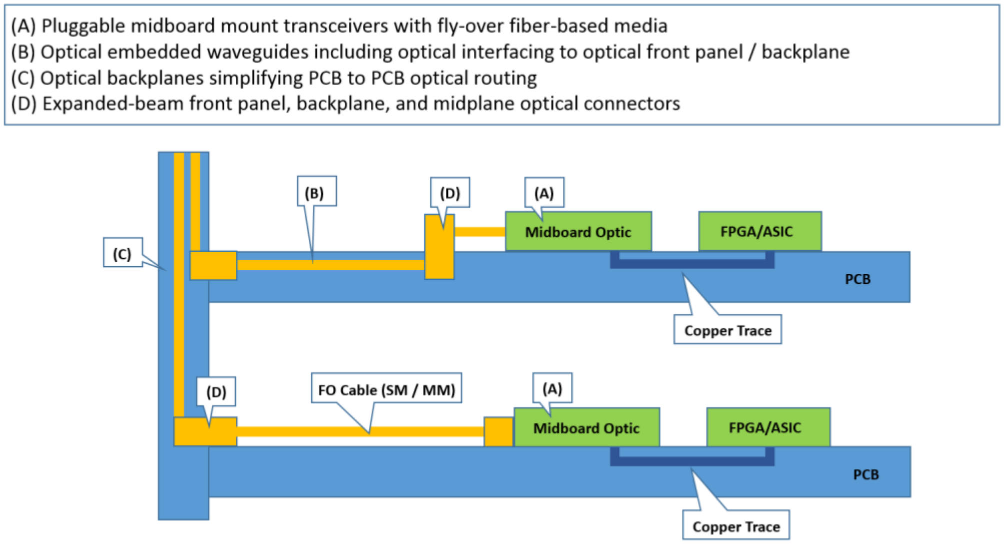


**Figure 3.** Technology approach for next-generation systems in Application Area 3.

### *Application Area 4: Intra-blade optical connections (length 0.5 - 0.05m)*

These are connections across a single blade; they have not been commercially implemented to date because data rates have not yet reached the point where optical communication is required for such short distances. However, in future systems it is expected that multiple electro/optical modules will be placed on a single blade, and that optical channels will provide the densest interconnection medium. Optical interconnections between these modules can be realized via both interposer- and embedded-waveguide-based optical interconnects.

### *Next-generation technology for Application Area 4*

- Reflowable electronic/photonic integrated modules to eliminate manual placement of modules in copper sockets (but still requiring manual coupling of optical connectors) (AA4.1)
- Interposer on PCB to provide electrical and optical traces connecting separate modules on the same interposer, to isolate modules from PCB reflow process (AA4.2)
- Optical embedded waveguides including optical interfacing to optical front panel and backplane (AA4.3)

**Table 4.** Technology development topics for Application Area 4 (intra-blade).

| | AA4.1 Reflowable modules | AA4.2 Interposer on PCB | AA4.3 Optical embedded waveguide |
|---|---|---|---|
| Manufacturing processes | X | X | X |
| Manufacturing equipment | X | X | X |
| Materials | X | X | X |
| Quality/Reliability | X | X | X |
| Environmental technology | | | |
| Test, Inspection Measurement | | X | X |
| Attenuation | X | X | X |
| Density | X | X | X |

Development of practical, cost-effective module-to-module connections across a blade is the "Holy Grail" of optical interconnect, and it presupposes success of the technical developments called for in Application Areas 1-3 above. The primary driving force for optical interconnect between modules over distances shorter than blade dimensions is signal degradation over copper traces at very high data rates/channel, probably in excess of 100 Gbps/channel. For such an approach to be economically feasible, several developments are required. First, labor associated with manual routing and coupling of the optical transmission medium must be eliminated. Second the need to separate the process for electrical coupling from optical coupling during the board assembly process must be eliminated. This means that embedded waveguides, with low-loss module-to-waveguide coupling technology, compatible with standard PCB fabrication technology (including reflow) are essential.

There are multiple technical challenges to the development of practical reflowable modules compatible with embedded waveguides. First, all reflow-incompatible materials in the modules must be eliminated. In the past it was assumed that organic adhesives and molded polymer optical coupling elements would have to be replaced by metal, or glass equivalents. Recently, there has been progress in high-temperature-compatible optical polymers (e.g. Extem™ polyimide and polyetherimide) and hybrid materials (e.g. ormocers) that may provide simpler fabrication options. Second, the structures that provide optical coupling between the embedded waveguides and the module optical interface must be compatible with the positioning tolerances and cleanliness characteristic of the automated module placement and reflow processes.

Related challenges exist for the embedded waveguides that will interconnect the modules. These waveguides may span the entire blade, or may be confined to an interposer smaller than the blade. In either case, new (reflow compatible) materials and processes must be developed to fabricate coupling structures in the waveguides (e.g. gratings or mirrors) that allow low-loss coupling to modules.

In the case of embedded waveguides spanning an entire blade, two dominant types of substrates can be distinguished: Rigid Multilayer PCBs and Flexible PCBs. Both can be "active" or "passive" and all are custom engineered for each application – unlike connectors, which have many standard designs.

Commercial rigid PCB materials include a wide range of organic materials including pre-impregnated epoxy-glass "prepreg" sheets, FR4+ low-electrical-loss laminate materials, copper foil, additive Cu (via chemical processes).To add optical functionality, silicone, glass or other optical materials can be incorporated as outer- or inner-layer optical waveguide layers. The ability to add layers of silicone or other optical polymeric waveguide materials or glass external to PCBs should be relatively within existing technology; but connecting these optical traces to surface-mount components, connectors or fibers will be a major challenge for high-volume manufacturing.

In the case of embedded waveguides spanning an entire blade, new waveguide materials are needed. Current polymer waveguides that can be embedded in PCBs have loss that is too high for practical use, at least at the operating wavelength (near 1310 or 1550 nm) of the anticipated SiPh modules. Materials with loss < ~ 0.02 dB/cm are needed; SM polymer waveguides have loss > 10x higher. Glass waveguides by ion-exchange have reported loss of 0.04-0.05 dB/cm and likely can be reduced to meet the target [3,4]. It has also been embedded within or on top of PCBs [5]. This still is in research though so for now deployed optical interconnect at the board level are essentially 100% done with cables and connectors.

One approach to dealing with the high loss of today's embedded waveguides may be to cluster modules needing optical connections on a common interposer. This provides the benefit of shortening the optical path length to reduce loss. It also allows the use of interposer materials systems which allow fabrication of low-loss waveguides,
but are of limited size due to use of a wafer technology (e.g. SiPh or silica-on-silicon wafers etc.) compared to a panel technology like glass. As in the case of the embedded-waveguide blade, optical coupling and materials challenges remain.

The interposers utilize different layouts depending on the type of interconnections that it needs to provide. In case electronic re-routing or fan-out is required at the interface between PCB and packaged opto-electronic ICs (OEIC) or between two or more packaged opto-electronic ICs, an interposer with up to several thousands of electronic lines will be required. The electrical interface between the OEIC package and the interposer, and between the interposer and the PCB will be realized via a connector or reflow approach. If in addition to this electronic interfacing, optical interfacing will also be required, the interposer will be equipped with optical waveguides.

The use of optical interposers requires optically mating chips and modules to substrate waveguides. The ultimate package interconnect would be Z-axis interconnect, similar to a BGA but with optical interconnect to waveguides on the substrate. This is possible with evanescent or adiabatic coupling and inverse-taper sections for the coupling part of the waveguides. The final package would be a fully-integrated photonic system which will first be heterogeneous and, ideally in the end, monolithic.

Interfacing to and from interposers could require IC socket and PC board type connectors, likely with both optical and high-speed electrical channels. Neither type of connector is yet available and substantial development will be needed before being available for application in systems. An alternate approach is BGA attachment to the PCB and fiber fly-overs to the front panel.

In spite of the early stage of embedded optical waveguides in PCBs and interposers, some efforts to produce standards for such products have begun. For example, the IEC has begun an effort to develop a standard for the geometry and performance of embedded optical waveguides (IEC 62496).

One important infrastructure challenge in development of optical PCBs is the nature of the PCB industry. The PCB market comprises over 1,000 firms worldwide, with organic PCB technology for electronics mature, and with materials and knowhow in the public domain. Firms in this industry are typically neither highly funded nor have sufficient margins to conduct a lot of research. Only a few firms post industry consolidation have strong RDE capability: one or two in the US and in Japan and Taiwan. Therefore, developments in optical PCBs currently depend on government funding and/or university research, or perhaps an unanticipated shake-up in the value chain.

Several issues are associated with the potential Optical PCB (OPCB) supply chain: i) pollution related to PCB manufacturing ii) the aforementioned very few (<5 worldwide) PCB manufacturers exploring Optical PCB) technology, with none of those having actual products; iii) only limited activity to develop flexible polyimide and or polyester PCB technology…which could be key ingredients to a maturing OPCB technology.

When and if SM PCB-embedded waveguide technology does emerge, or systems undergo a radical change to photonic computing in integrated photonic modules, a new breed of interconnect devices will likely be needed. Areas that will need additional development are mass-production-compatible chip edge coupling to external cables or waveguides, optical interposers at the chip/package level, and the PCB-embedded waveguides. The so-called ‘Chicken and Egg’ syndrome impedes some connector developments. The typical connector industry scenario is to develop and make products for a specific customer demand, then for a market – in that order. In some cases, tooling costs are shared between the OEM and connector supplier. Since data centers are typically not OEMs, this introduces a new challenge into the connector manufacturers’ technology development and market coverage.

(A) Reflowable electronic/photonic integrated modules to eliminate manual placement of modules in copper sockets (but still requiring manual coupling of optical connectors)
(B) Interposer on PCB to provide electrical and optical traces connecting separate modules on the same interposer, to isolate modules from PCB reflow process
(C) Optical embedded waveguides including optical interfacing to optical front panel / backplane

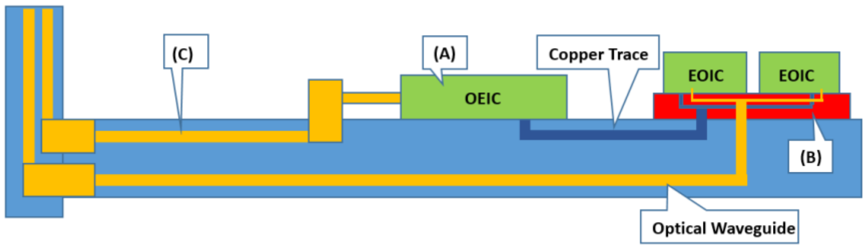


**Figure 4:** Technology approach for next generation systems in Application Area 4.

***Application Area 5: Intra-module optical connections (length < 0.05m)***

These are connections inside a module package. Integrated electronic/photonic modules require high-density low-cost, low-optical-loss assembly technologies that provide an integrated system with adequate reliability and lifetime. Such connections are in wide use today in two types of applications: connections from a PIC to a connector interface at the module wall, and connections between two or more PICs occupying a single package. In the former case, the connection is usually in the form of a short array of fiber stubs or polymer waveguides. In the latter case, the optical connections can be realized by (1) direct coupling using end-fire coupling, relay micro-lenses, fiber stubs or by direct evanescent optical coupling between optical waveguides by placing PICs on top of each other, or (2) via the use of a waveguide interposer to which multiple chips are optically coupled. The most common approaches today use either fiber stubs or relay micro-lenses to couple lasers to waveguide chips. In the case of an interposer the optical (and electrical) interconnections are realized via a submount (e.g. based on SOI or SiN/Si).

In the remainder of this analysis, detailed technology options and their status, needs for new technologies to advance short-range interconnect, infrastructure considerations, and associated roadmap milestones will be discussed in the context of the above application areas.

***Next-generation technology for Application Area 5***

- Optical coupling elements that self-align to PICs, and couple to a package connector interface (AA5.1)
- Module substrates that incorporate waveguides that can couple between multiple PICs in a single module, or to a package connector interface (AA5.2)
- Interposers for low-loss transmission between PICs or PICs and connectors, with metal traces patterned at the wafer scale for low manufacturing costs (AA5.3)

**Table 5.** Technology development topics for Application Area 5 (intra-module).

| | AA5.1 Self-aligning coupling elements to a package connector interface | AA5.2 Module substrates for optical connections within a module | AA5.3 Interposer for low loss transmission |
|---|---|---|---|
| Manufacturing processes | X | X | X |
| Manufacturing Equipment | X | X | X |
| Materials | X | X | X |
| Quality/Reliability | X | X | X |
| Environmental technology | | | |
| Test, Inspection Measurement | X | X | X |
| Attenuation | X | X | X |
| Density | X | X | X |

Optical connections between components within a module are already available. Typically, such connections are between laser sources and silicon photonic or InP PIC chips (e.g. modulators), or between PICs and connector interfaces. Common approaches include free-space relay lenses, short sections of optical fiber, or even polymer waveguides (used in millimeter lengths where their contribution to the total loss is tolerable). These approaches

typically require active alignment steps that are slow and expensive, so new approaches to high-throughput automated assembly of the modules are needed, e.g. self-alignment.
Since the ultimate interest is in modules that are low loss at 1310 nm and 1550 nm, and are reflow compatible, the materials challenges cited above apply here. This means that approaches using current polymer waveguide or organic adhesive technologies are probably not viable long-term solutions.

Within a module, the density of interconnection can be a very important cost driver. Waveguide pitches in high-index-contrast semiconductor waveguides can be small (e.g. less thana few tens of microns), whereas fibers are large (80 or 125 micron diameter), and low-index-contrast waveguides (photorefractive polymers or ion-exchanged glass) require pitches >> 50 microns to avoid cross-coupling. This means that for multi-port devices the spacing of output ports on the PIC, driven by the coupling waveguide medium, must be larger, resulting in larger areas of expensive semiconductor chips being required just for coupling. Therefore cost-effective intra-module interconnect medium needs to have a pitch matching the "native pitch" of the PIC. This is true on the end of the interconnect medium coupling to the PIC, but the interconnect medium may serve as a "pitch transformer" to couple the PIC into a traditional-pitch (250 microns) connector interface.

#### MANUFACTURING EQUIPMENT

- PIC fabrication equipment is well established, and a large infrastructure exists. PICs do not require state-of-the-art lithography, so can be patterned in low-cost, depreciated-capital legacy fabs. However, there may be issues in maintaining equipment in these limited-resolution fabs as volume Si chips (electronics) move to higher resolution
- There is no established manufacturing equipment for high-volume, low-cost fiber/waveguide termination in connectors
- There is no established manufacturing equipment for high-volume, low-cost optical connections to chips

#### MANUFACTURING PROCESSES

- There are no established processes for automated high-volume, low-cost fiber/waveguide termination in connectors
- There are no established processes for automated high-volume, low-cost optical connections to chips
- There are no established processes for automated high-volume, low-cost PCBs with optical embedded waveguides

#### MATERIALS

- There is a need for moldable optically-transparent materials that are CTE-matched to Si and reflow compatible for fabrication of low-cost optical couplers
- There is a need for low-CTE optical adhesives that are compatible with reflow conditions, for use in bonding optical couplers or fibers to PICs

#### QUALITY/RELIABILITY

- New reliability specifications may be required for optical modules meant for use in data center and supercomputer environments, where higher level of cleanliness, but also higher temperatures, may be experienced than addressed in e.g. traditional Telcordia environmental specifications
- New dust test methods and standards need to be developed to address the new SM expanded beam connectors

#### ENVIRONMENTAL TECHNOLOGIES

- Like in regular electronic IC production, the production of PICs involves the use of hazardous fluids and gases (e.g. solvents, etchants and layer deposition gases)

- SiPh processing does not bring much additional environmental contamination hazard (beyond that associated with high volume silicon electronics manufacturing)
- Recently, one of the principal driving forces for the use of optical interconnect over short distances has been the theoretical potential for reducing the overall energy consumption of a data center. In principle, the low attenuation and distortion of high-speed signals in the optical domain can lower the amount of energy required to transmit a logical bit. This could translate into reduced computational power consumption by the center as well as reduced power consumed in air conditioning. The reduced power consumption can reduce not only the operating cost of the facility (the electric bill), but also carbon emission and thermal pollution related to power generation. Unfortunately, at today's transmission speeds with today's optical transmission technologies, the power consumption advantage of optics over copper has not yet been realized at the system level.
- There are known toxicity issues with III-Vs semiconductor elements and compounds thereof, including As, Ga and In
- Indium has limited abundance, is heavily used in transparent conductors for touch screens, etc., and is considered a strategic material with significant future supply risk

### TEST, INSPECTION, MEASUREMENT (TIM)

- TIM approaches for completed modules are well established but require manual mating of connectors to test instruments, so are slow. Also, testing at the module level wastes resources when defective chips are packaged
- Wafer-level testing need to be implemented for low-cost, high-volume manufacturing. This requires a probe system that integrates electrical and optical probes, and for devices that don't emit light, provides light to each device for testing
- Wafer level testing may be challenging for edge coupled devices, since the optical input/output facets may not be accessible prior to dicing the wafer to produce either individual PICs or bars of PICs

## ROADMAP OF QUANTIFIED KEY ATTRIBUTE NEEDS

The development track of substrate/interconnection technology depends strongly on the timing of the transition from Cu signaling to photonics at the chip, package and board level of datacom and computer/server/storage equipment.

In this, four stages are foreseen over the next two decades:

- 2020-25: Heterogeneous photonic solutions with advanced 3D packaging. Embedded waveguides with surface-level interconnects (e.g. grating-based) at the PIC level, Coexistence of MM and SM fiber solutions for rack-to-rack and limited intra-rack interconnect. Introduction of Wavelength Division Multiplexing (WDM). Introduction of SM expanded-beam optical connectors. Persistence of board-edge mounting of transceiver modules but beginning of transition to mid-board transceivers or co-packaging. Pluggable socket for module-board electrical connection
- 2025-30: Silicon photonic and InP PICs will be widely commercialized; monolithic integration will result in single-chip or complex 3D chip solutions. SM optical fiber and waveguides will be used for I/O ports and will dominate to the rack-to-rack level, with rack-to-rack interconnect driving volume manufacturing. SM expanded-beam connectors will be widely implemented. Coarse WDM (CWDM) to 4 wavelengths will be common. Modules may move to mid-board if the value proposition based on shortened copper length (and anticipated reduction in Serializer/De-serializer (SERDES) IC power) is borne out, but still with pluggable electrical sockets
- 2030-35: Use of spatial multiplexing via multicore fiber will begin at scale, probably planar arrays of cores fiber at first, for compatibility with planar interfaces on PICs. Intra-module interconnect via waveguide interposers will begin. Modules compatible with standard reflow processes will become available if

justified by the manufacturing and repair/rework costs relative to modules that are plugged into sockets on the PCB.

- 2035-2040: Spatial multiplexing will be pushed to higher core counts in fiber, and multi-level interfacing schemes (3-D fanouts) will be used in chips. Mid-board modules will be reflow compatible, and optically self-aligning during the assembly process if justified by the manufacturing and repair/rework costs relative to modules that are plugged into sockets on the PCB.

Figure 5 below shows a high-level view of the expected co-evolution of module characteristics. The exact timescale is dependent on many technical and economic factors, but the order of events is expected to be accurate. Note that the final dominance of SM interconnect in the shortest applications may be delayed by packaging cost disadvantages relative to MM interconnect, until both the need for WDM to boost capacity, and the technology for simple, manufacturable SM alignment, are present. However, whenever compatibility with fibers reaching further than several meters is required, SM will win.

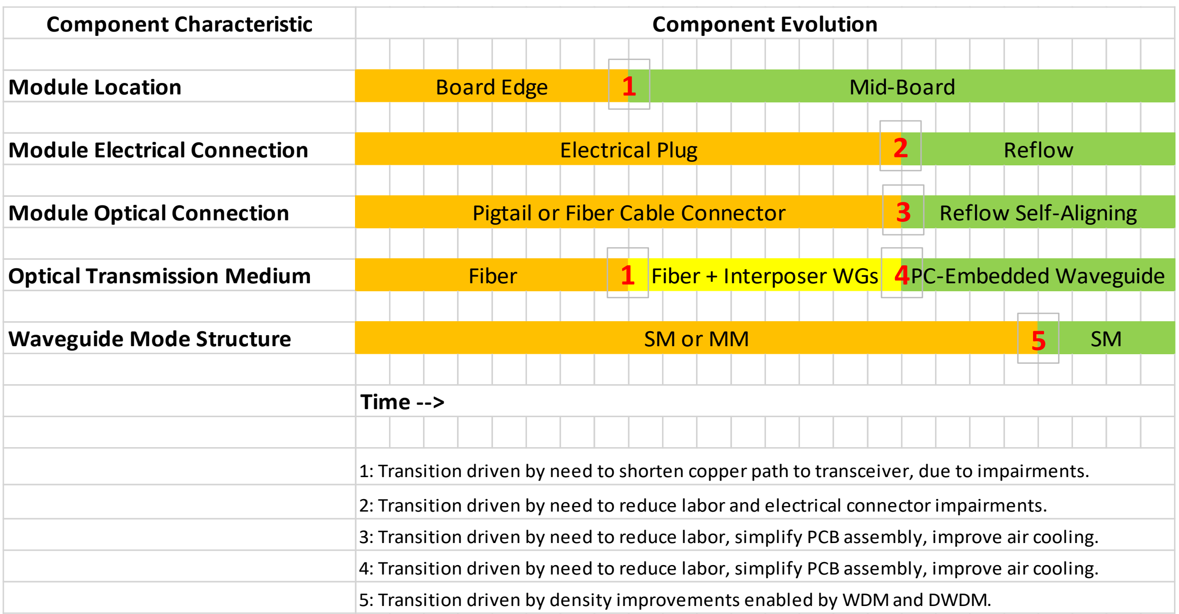


**Figure 5.** Expected general evolutionary trends of optical interconnect technology characteristics.

Further, more detailed information on the evolution of some key attribute needs of optical interconnects is provided in Table 6.

**Table 6.** Evolution of key interconnect application attributes.

| Table 6: Roadmap of Quantified Key Attribute Needs | [unit] | Current | 2025 | 2030 | 2035 | 2040 |
|---|---|---|---|---|---|---|
| **Optical Connector Cost/Bit** | $/Gbps | 0.75 ($30 @ 40G) | 0.0375 ($15 @ 400G) | 0.00344 ($22 @ 6.4T) | 0.00073 ($22 @ 30T) | 0.00030 ($22 @ 72T) |
| **Optical Connector Size (Long x Short)** | mm | 25 x 9.8 (MPO) | 25 x 9.8 (MPO) | 13.6 x 4.8 (XB) | 13.6 x 4.8 (XB) | 13.6 x 4.8 (XB) |
| **Gbps speed per Wavelength Channel** | Gpbs | 100 | 100 | 400 | 400 | 400 |
| **Wavelengths per Fiber Core or Waveguide** | # | 1 | 4 | 4 | 8 | 16 |
| **Cores/Fiber** | # | 1 | 1 | 3, 4 | 7, 16 | 7, 16 |
| **Fiber Pitch** | micron | 250 | 127 | 127 | 84 | 84 |
| **Fibers/ Connector** | # | 1-2-4-8-16-32-64 | 1-2-4-8-16 | 1-2-4-8-16 | 1-2-4-8-16-32 | 1-2-4-8-16-32 |
| **db Loss Budget TX to RX** | dB | 4db | 3db | 3db | 3db | 3db |
| **Key Connector Types** | ---- | LC (MM), MPO, MXC, BP[2] (SM) | Chip-X[3], PCIe (SM) | PCI-X, Chip-X, Co-packaged | Other Co-packaged, TBD | Other Co-packaged, TBD |
| **Sockets/ Interposers** | (status) | In Development | Connector or Chip OEM/OSAT? | Chip-Integrated | Chip-Integrated | Chip-Integrated |
| **Cables** | ---- | 1000s of Racks, Fly-Over | 100s of Racks, | 10s of Racks, WGs[4] | 10s of Racks, WGs | 10s of Racks, WGs |
| **System** | ---- | Discrete Conv. Packaging | 3D-E0-SiPh | Monolithic E0-SoC | Monolithic E0-SoC | Monolithic E0-SoC |
| **Roadblocks** | ---- | SM-WGs, Chip-X | SiPh[5] Integration (SoC) | SiPh/SoC Modules; multicore fiber coupling | SiPh/SoC Modules; multicore fiber coupling; T-stable muxes | SiPh/SoC Modules; multicore fiber coupling; T-stable muxes |
| **Comment** | ---- | Major Issues are WGs & Chip-X | Mid-Board, IO Standards | Modularization < Connectors[6] | | |

**In Production** **Modest Difficulties** **Significant Roadblocks** **Major Technology Challenges**

** Connector cost can be reduced with increased manufacturing volume or via offshoring to low-manufacturing-cost regions. A preferred path is automation and making products regionally where used.*

*1=Experimental 2=Backplane 3=Direct Chip/Package Attach 4=Waveguides 5=Silicon Photonics 6=Miniaturized, semi/monolithic modular circuitry*

In the above chart, neither alignment tolerance nor spectral bandwidth of the components is explicitly included. This is because these parameters are highly dependent on other design parameters which can be combined to achieve the same performance metrics in the chart. For example, alignment tolerance is determined by the optical mode size of the components being coupled. In coupling to standard SM fiber modes, which are on the order of 10 microns in size, lateral alignment to around 1 micron is adequate. When coupling components with smaller waveguide modes, say 2 microns, for a SiPh PIC, lateral alignment to around 0.2 microns is needed for the same level of loss.

Similarly, the spectral bandwidth required for components is dependent not only by the channel speed, encoding scheme and number of channels, but also on the spectral shifts of the multiplexers and sources over the operating temperature range. For current (expensive) wavelength stabilized 100 Gbps telecom channels, 0.8 nm channel spacing is typical. For the case of low-cost, robust systems without temperature control, larger channel separation and broader operating spectra will be required.

In the interconnect evolution the connector developments will follow OEM/EMS requirements. Key areas of development include materials and process technologies, high-speed performance, miniaturization and close attention to system life cycle for optimized reliability vs cost balancing. Mobile system interconnect requirements may drive future micro-scale robotic connector design, plus other dimensional and environmental requirements outside the realm of conventional stamp and form/mold connector processes. However, FO connector developments will be more dependent on telecom/datacom and computer-oriented applications. This may complicate new product development because data center operators are mostly not OEMs. Thus, consortium efforts, with members from the equipment industry, will need to speed up development efforts. They include minimal challenges for existing connectors beyond verification of I/O-midboard and backplane verification. This includes expanded Beam SM MPO and MXC connector designs.

On basis of the optical interconnect technology evolution the following critical, regular and desirable milestones are identified:

Critical Milestones:

CM1 Low-cost packaging approaches for SM PICs
CM2 Low-cost fiber termination technologies for SM fibers
CM3 Higher-density optical fiber and connector interfaces to match PIC waveguide pitch
CM4 Reflow-compatible optical coupling technology for PIC chips and modules
CM5 Broad-wavelength-band optical coupling technologies for PICs, to allow implementation of wavelength multiplexing

Regular Milestones:

RM1 Low-loss expanded-beam connectors for SM fiber, suitable for backplane, midplane, and front-plane use
RM2 Low-profile expanded-beam connectors for PIC chip and module interfaces
RM3 Optical backplanes and mid-planes, providing routing of optical channels between blades in a rack

Desirable Milestones:

DM1 PCB or interposer with embedded low-loss SM waveguides, and in-/out-coupling for modules and connectors. This also includes optical Ball Grid Array (BGA) and Vertical Cavity Surface Emitting Laser (VCSEL) interposer developments with mechanical integrity for advanced Surface Mount Technology (SMT) applications
DM2 Convergence to one or a small number of PIC chip waveguide optical coupling interface designs, to allow development of “generic” packaging technologies, with economy of scale. The optical chip packaging interconnect will include optical IC card edge, Z-axis or waveguide interconnects

**CRITICAL (INFRASTRUCTURE) ISSUES**

Although an extensive infrastructure has been established to support the telecommunications and data communications industry needs to date, and has done so adequately, the penetration of optical interconnect into higher-volume and shorter-distance applications is not adequately supported by this legacy infrastructure. In fact, in some sense, the presence of the legacy infrastructure may be an obstacle to the development of the new infrastructure that is needed. This is because the legacy optical interconnect infrastructure was built around products for long-haul or specialty communications applications where cost was (almost) no object. There was little motivation to drive down manufacturing cost, and so the legacy infrastructure is not compatible with achieving the cost targets essential for wide-scale penetration of optical interconnect into applications of the future. Thus, in the discussion below, it will be apparent that much of the new infrastructure needed is not for performing new functions, but rather for performing familiar ones more efficiently.

A few manufacturers have answered the call of data center applications, e.g. notably Molex in the US. However, most fiber optic connector products are assembled in Chinese and Malaysian factories, so that costs have been driven down to minimums with bench-type assembly using low-cost labor. Now that those costs have risen, notably in China, other assembly options (e.g. Indian and Vietnamese bench assembly) are on the table. Ultimately, sustainable lower costs will depend on high volumes [100,000s to Millions] and automation, which has historically not been the case with these products.

***Equipment for low-cost automated termination of connectors***

Termination of fiber optic cables, defined as the process of installing an optical connector on the end of a cable, is still a manual task (see above). For longer-distance spans of fiber, where the precise final length of the span is unknown until the fiber is laid, connectors are typically field-installed on the ends of pre-laid fiber cable. This can be done either by attaching the connector directly to the fiber cable, or by using mechanical or fusion spicing to splice on a short section of fiber that has been factory terminated with the connector. In this long-span application, the number of connectors that must be installed per length of fiber is small, so the manual process is acceptable.

For shorter-span applications such as rack-to-rack or intra-rack spans in a data center or supercomputer, factory pre-terminated cable assemblies of pre-determined lengths are preferred. These pre-terminated assemblies allow faster installation, use of lower-cost labor, and improved reliability via factory testing for verification of the optical performance before installation.

However, the process for factory production of pre-terminated cable assemblies is currently very similar to what is done in the field. That is, the process is still a high-labor-content manual process involving technicians installing one connector at a time. (This situation has driven the termination business to low-labor-cost regions.) Part of the reason for the manual process is the relatively low volume and moderate price pressure of cable assemblies at present, which has not provided adequate economic motivation to find higher productivity approaches. However, a more fundamental technical reason for the longevity of the manual process is the difficulty of automating the process of terminating fiber with existing connectors. This difficulty stems from a basic design element of most commercial fiber optic connectors: fibers that have been stripped and cleaned to produce a pristine glass outer diameter must be inserted into cylindrical holes having ~ 1 micron clearance in the connector ferrule and bonded in place. Finding an economical path to automating this process is a challenging proposition.

Nevertheless, existence of high-throughput, low-cost factory termination equipment is a key enabling element of infrastructure for manufacturing of future cable assemblies. At the present moment, it appears that this will not be achieved using current ferrule and connector designs. This means that the best path to low-cost high-volume manufacturing of cable assemblies may be to put aside the legacy connector designs and develop new connectors that are specifically designed to enable automated termination with relatively low-cost capital equipment. These new connector designs would be developed in parallel with the new automated termination equipment, to optimize productivity, and the automated termination equipment will need to be widely available.

***Equipment for low-cost manufacturing of packaging, including fiber attachment***

In the process of converting a bare PIC to a finished functional module, the most challenging and cost-intensive step (80% of package manufacturing cost by some estimates) is making the optical connection between the chip and the outside world. This is difficult because it requires precisely locating the optical mode of the PIC and then aligning and permanently attaching a fiber or other optical coupling element to that mode with sub-micron tolerances. Approaches for finding the mode include “active alignment” which involves moving the fiber/coupler relative to the chip to find an optical coupling maximum, “robotic vision alignment” where the fiber/coupler is aligned to fiducial marks on the chip that are designed to be precisely registered to the optical output, or “passive alignment” where there are mechanical interlocking features on the chip and the fiber carrier or coupler to hold the fiber/coupler in alignment with the mode (see section 5.1.6 for more discussion of these technical approaches).

Today, module manufacturers use active or robotic vision alignment to assemble modules. Both of these approaches use expensive micro-positioners to manipulate the fiber/coupler plus either power meters or robot vision systems to provide feedback information for coupling optimization. While suitable alignment systems are commercially available and proprietary alignment systems can be readily developed, these systems are expensive and have relatively low throughput. The low throughput is not intrinsic to the alignment process but is often limited by the set-up time (i.e. the time required to attach fiber input and output cables for active alignment, or by the time required to cure the bonding adhesive.

In the future it would be desirable to eliminate as much of this precision positioning equipment as possible. One option that has been investigated for many years but not yet fully perfected is the use of solder surface tension and etched stops to position elements relative to each other with sub-micron accuracy. This precise positioning technique requires sub-micron dimensional control of fiber/coupler elements as well as location of waveguides in all three dimensions. If successful, this approach could be carried out with pick-and-place equipment and reflow equipment that the module manufacturer would likely already have, thus minimizing investment in new infrastructure. In fact, any technique that piggy backs off the established pick-and-place and reflow infrastructure of microelectronics should have an inherent cost and adoption advantage.

***Foundries for low-cost, high-volume manufacturing of PICs***

Because of the high cost of building a semiconductor foundry capable of PIC production, many companies in the industry are “fabless” and rely on contract foundries to fabricate their PICs. There are already multiple PIC foundries operating, especially for the silicon photonic material system. These include both “pay-for-play” foundries open to any customer, and captive foundries belonging to a company (but sometimes still potentially available for outside contracted work). Currently there is no demand for high wafer throughput (relative to silicon electronics) due to the low level of PIC market consumption; cost scaling at high volume has not been firmly established.

One factor that may interfere with cost reduction with volume increase for PICs is the lack of standardization in PIC processing. Different foundries have different standard (and often proprietary) elements in their design library (their “PDK” for “Process Design Kit”), so it may prove difficult to scale volume by employing multiple foundries with foundry-specific designs.

Furthermore, most of the PIC foundries offer only PIC fabrication up through chip singulation and perform very limited amounts of testing. They do not develop packaging for their components and no standard packaging exists. Therefore, a customer using a foundry to develop a PIC must either develop packaging internally or find another contractor to develop the package. This can be very significant because many of the key performance attributes of the PIC module are critically dependent on the quality of the packaging (e.g. optical properties like insertion loss, return loss, and polarization dependent loss, as well as electrical properties like modulator and detector bandwidth).

Ultimately high-volume low-cost manufacturing will depend on standardization of both PIC process elements and packaging approaches and on co-optimization of these. This may be accelerated by organizations providing PIC fabrication, as well as providers of test, assembly and packaging services; the latter is the objective of the AIM Photonics Foundry and TAP (Test, Assembly, Packaging) facility.

***Supply chain and manufacturing technology for low-loss waveguides embedded in PCBs, with integrated optical coupling mechanisms, like "optical solder bumps"***

As mentioned elsewhere in this document there are many materials and process challenges to overcome to enable PCBs with low-loss embedded optical waveguide interconnect between modules on a board or interposer. or between modules and fiber connectors at the board edge. Furthermore, there is no existing infrastructure suitable for manufacturing PCBs with embedded waveguides. One design constraint is the long wavelengths (1.3 and 1.5 microns) that will be used in these systems. Because of the absorption losses of polymers at these wavelengths, it is likely that inorganic waveguides will be required to achieve adequately low loss. Low loss glass waveguides have been demonstrated but establishing the commercial infrastructure for manufacturing PCBs with sheets of inorganic materials accurately embedded in them remains a challenge. While the process technologies used in PCB fabrication can be used to embed inorganic glass sheets having embedded waveguides with some modifications to standard processes, the real challenge is that there is no current support for the sub-micron alignment tolerances that will be required between embedded optical waveguides and surface optical coupling features.

At present, addressing these infrastructure issues is not economically feasible, since there is no current application to drive volume scaling. The infrastructure for combined electrical and optical PCBs will need to develop after full lab-scale technical solutions emerge. Based on the lessons learned in the evolution of fiber connectorization and fiber attachment in modules, it is important that proposed technical solutions are judged heavily on the difficulty of establishing the infrastructure that will be needed to manufacture them in a cost-effective way. It is estimated that very short reach optical interconnects within boards or racks will need to meet cost targets 4 to 10 times lower than intra-data center interconnects.

***Workforce trained to design, install, and maintain electrical-optical PCBs***

Wide-scale implementation of optical interconnects will require a large workforce of assembly, installation and maintenance technicians that are familiar with the basic concepts of guided-wave optics, optical measurements and precision engineering. Such a workforce does not exist today, and there are sparse existing resources for training one. However, this need has been recognized, and programs to educate certified photonics technicians are beginning to emerge. One example of such a program is the new photonics technician certification program that will be provided by the AIM Photonics Academy, beginning in 2020.

## TECHNOLOGY NEEDS

Tables 7 and 8 below summarize the near and longer-term technology requirements; detailed discussions of each area follow.

| **Table 7. Prioritized Research Milestones (>2025)** | **Relative Priority** |
|---|---|
| Simplified approaches for optical coupling of connectors to PICs, e.g. self alignment | Critical |
| Low-loss optical waveguides for integration in PCBs or interposers | Critical for interposer; Regular for PCBs |
| Low-loss coupling technology from PICs to PCB or interposer waveguides. | Critical for interposer; Regular for PCBs |
| Optical alignment of chips/modules to PCBs or interposers via reflow | Critical for interposer; Regular for PCBs |

| **Table 8. Prioritized Development and Implementation Milestone (≤ 2025)** | **Relative Priority** |
|---|---|
| Low-cost connector termination technology | Critical |
| Low-cost PIC packaging technology (high IO count coupling to fiber connector) | Critical |
| SM expanded beam connectors, for cables and modules | Regular |
| SM expanded beam connectors, for optical backplane, front panel, and midplane applications | Regular |
| Improved cable densities, routing and management technology | Critical |
| Optical interposers for coupling of PICs | Regular |

In addition to the above prioritized needs, the following additional needs have been identified:

- Availability of core competencies like:
  - Electronics/photonics technology and intellectual property
  - Photonics circuit design
  - Computer-aided design for manufacturing and design collaboration
  - Precision injection molding with mold equipment suppliers
  - Materials technology with materials suppliers

- Ferrule technology: designs and methods allowing automated mass production of single mode optical device interfaces and interconnection cables

**Table 9:** Evolution of Technology Elements to Support Interconnection Applications.

| Roadmap of Supporting Interconnection Technologies | [unit] | Current | 2025 | 2030 | 2035 | 2040 |
|---|---|---|---|---|---|---|
| **Expanded Beam MM Connector** | | | | | | |
| Waveguide-to-waveguide loss[1] | dB | 1.5 for 12 fibers | 1.5 for 64 fibers | 1.0 for 64 fibers | 0.5 for 64 fibers | 0.5 for 64 fibers |
| Fiber density | #/mm$^2$ | 0.05 | 0.26 | 1.0 | 1.0 | 1.0 |
| Float for backplane application | mm | +/- 1 | +/- 1 | +/- 1 | +/- 1 | +/- 1 |
| Termination process | --- | manual | Semi-robotic | robotic | passive | passive |
| **Expanded Beam SM Connector** | | | | | | |
| Waveguide-to-waveguide loss[1] | dB | 1 dB for 12 fibers | 0.75 for 16 fibers | 0.5 for 32 fibers | 0.5 for 64 fibers | 0.5 for 64 fibers |
| Fiber density | #/mm$^2$ | 0.05 | 0.26 | 1.0 | 1.0 | 1.0 |
| Float for backplane application | mm | +/- 1 | +/- 1 | +/- 1 | +/- 1 | +/- 1 |
| Termination process | --- | manual | manual | semi-robotic | robotic | passive |
| Reflection loss | dB | -50 | -50 | -50 | -50 | -50 |
| **Optical Transport Media** | | | | | | |
| SM Fiber loss | dB/cm | < $3x10^{-6}$ single core | < $3x10^{-6}$ single core | < $1x10^{-4}$ multicore | < $1x10^{-4}$ multicore | < $1x10^{-4}$ multicore |
| Waveguide interposer loss | dB/cm | < 0.16 | < 0.16 | < 0.16 | < 0.16 | < 0.16 |
| Waveguide interposer PDL | dB | < 0.2 | < 0.2 | < 0.2 | < 0.2 | < 0.2 |
| Embedded waveguide loss | dB/cm | < 0.02 | < 0.02 | < 0.02 | < 0.02 | < 0.02 |
| Embedded waveguide PDL | dB | < 0.2 | < 0.2 | < 0.2 | < 0.2 | < 0.2 |
| **Module** | | | | | | |
| Module optical interface A. Pigtail fiber type and pitch | N.A. | Pigtail, pitch 250 microns | Pigtail, pitch 250, 127 or 84 microns | Less than 84 microns, or Multicore Fiber | Multicore Fiber | Multicore Fiber |
| Module optical interface B. Channel waveguide geometry | N.A. | 1D array waveguide | 1D array waveguide with pitch converter | 1D array waveguide, self-aligned waveguide array | 2-D waveguide array | 2-D waveguide array |
| Module optical interface C. Lens assisted | --- | Actively aligned 1-D micro-lens array | Actively aligned 1-D micro-lens array | Self-aligned 1-D micro-lens array | Self-aligned 2-D micro-lens array | Self-aligned 2-D micro-lens array |
| Module optical interface D. Interposer | --- | N.A. | Glass interposer | Glass or SOI interposer | Glass or SOI interposer | Glass or SOI interposer |
| Number of optical ports | --- | 12 | 24 | 24-128 | >128 | > 256 |
| Chip-to-medium coupling loss | dB | 1.5 | 1 | 1 | 0.7 | 0.5 |
| Chip-to-medium alignment | N.A. | Machine vision/ Active alignment | Machine vision/ Active alignment | Self-aligning/ Machine vision | Self-aligning | Self-aligning |
| Chip-to-medium coupling BW | nm | 40 near 1310 | 60 near 1310 | 100 near 1310 | 100 near 1310, or 1530-1565 | 100 near 1310, or 1530-1565 |
| Module-board attachment | N.A. | socket | socket | Reflow/ socket | Reflow/ socket | Reflow/ socket |
| Maximum assembly temperature | C | 80 | 80 | 260 | 260 | 260 |

| | | | | | | |
|---|---|---|---|---|---|---|
| Module price target | $/Gbps | 3 | 1 | 0.30 | 0.10 | 0.03 |

[1] These are the coupling losses at the interface between 2 fibers, 2 channel waveguides, or a channel waveguide and a fiber.

### *Qualitative Analysis of Main Challenges*

The primary impediment for future wide implementation of optical interconnect occurs for the short-distance, high-channel-count connections within a rack (our Application Areas 3,4, and 5). In these areas, there have been many demonstrations of module and connector technologies that are able to meet basic performance requirements (e.g. the Avago MicroPOD and MT-based multifiber connectors), but these approaches have not been broadly adopted, but rather have been used in specialized or demonstration systems because of excessive applied cost relative to copper. In this case, "cost" is intended to mean all the applied costs associated with use of the optical interconnect technology, including manufacturing cost of chips and modules, assembly cost of boards, yield and failure issues, labor and system maintenance.

The most important components of the applied costs are those which currently do not adequately scale downward with production volume. These include the following:

- Fabrication of SM fiber coupling elements
- Termination of fiber cables in connectors
- Maintenance of connectors
- Environmental stability issues with fiberoptic connectors
- Optical coupling of PICs to the passive optical interconnect media (fibers, connectors, and substrate-embedded channel waveguides)
- Assembly of the chips/modules onto the PCB
- Routing of fibers/waveguides from module-to-module over the PCB, or from blade-to-blade e.g. through a backplane

Each of these applied cost components will be considered separately below.

### *Fabrication of Single-Mode fiber coupling elements*

Fabrication of precision coupling elements, e.g. ferrules, for low-loss coupling of single-mode fibers or channel waveguides is difficult because of the tight mechanical tolerances that must be held. For coupling of conventional single-mode fibers (mode size ~ 9 microns), sub-micron alignments must be held. For coupling of tightly-confined waveguides on PICs, where mode sizes can be < 1 micron, tolerances are hundreds of nanometers. Such tolerances are very difficult to hold in low-cost fabrication processes (e.g. molding of plastics or glass, or casting of ceramics) so post-fabrication "touch-up" machining and/or sorting of parts is common; this results in higher cost. It also drives PIC designs where the waveguide mode field is expanded up to that of single mode field for relatively higher alignment tolerance at this critical interface

Note that the use of expanded-beam connectors, proposed herein for relaxing the alignment tolerances in the expanded beam path, does not circumvent the requirement of high precision, in this case for the alignment of the fiber/waveguide to the beam-expanding optics.

There is a need for new technology for fabricating precision coupling elements with high throughput and low cost, either by refinements of currently-used injection molding processes, or development of new innovative processes. Development of these components requires close coordination with the intended end-use applications and devices, as their initial ability to meet industry performance standards may be limited, until manufacturing and assembly processes mature over time.

The manufacturing cost of fiber optic connectors can be further reduced by:

- Domestic or China automation for fiber optic products, to replace current operator bench assembly
- Leverage of other low-cost labor areas, such as Vietnam
- Lower material costs via global sourcing and new, lower-cost materials
- High-volume automation for industry standard products, e.g. I/O connectors and cables

***Termination of Fiber Cables***

In conventional fiber optic connectors, fiber(s) are mounted in a precision ferrule (e.g. an MT ferrule), then the ferrules are held in mechanical registration by surrounding connector body parts. Termination of the fiber cable is the process by which the ferrule is attached to the fiber cable. This process involves stripping away the cable matrix and fiber buffer, cleaning the glass fiber, inserting the fiber in precision holes in the ferrule, bonding the fiber in place, and finishing the fiber ends for proper ferrule-to-ferrule mating geometries (typically via multi-step polishing to produce an angled, domed interface). This entire precision process is performed largely manually, with significant labor cost content.

In lieu of standard fiber-to-fiber physical contact mating, expanded-beam ferrules such as the expanded-beam multi-mode MT, offer a non-physical-contact (air gap) interconnect which greatly reduces labor and processing requirements, although with a trade-off in optical performance (loss). This tradeoff has diminished over time as the manufacturing process technology matures but expanded-beam performance remains below current physical-contact MT ferrule performance, especially for multi-row versions. Expanded-beam ferrule assembly steps trade precision laser cleaving of fibers for polishing, eliminating the need for polishing machines and film. The benefits of expanded-beam multimode ferrules in early system deployments (e.g. dust resistance) have proven to be attractive. However, their performance and ecosystem maturity do not yet support their use in the broader marketplace. Single-mode version of expanded-beam products are in the early stages of development and have proven to be much more difficult to design and manufacture.

The bottom line is that there is need for ferrule assembly technology that simplifies and automates fiber termination processes to drive down costs via labor reduction, and more importantly, increase inherent manufacturing capacity and reduce lead times. This development is often hampered due to lack of standardization in raw materials, connector types, and end-use configurations.

***Maintenance of current connectors***

Most current connector technologies, especially those for multiple single-mode fiber cables, rely on physical contact between polished fiber ends for optical coupling. The fibers are polished so the region over the fiber core is slightly domed, then adequate contact force between the fibers pushes the cores into intimate contact (excludes air), thereby providing a low-loss interface. It can be difficult to prepare the multiple fiber ends in an array connector to achieve physical contact over all fibers in a ferrule. Furthermore, the presence of any dust or other debris between the fiber ends causes poor contact and excess loss or reflection. This means that the connectors must be carefully protected or cleaned to eliminate particulates. For this reason (and others), there is a desire to use expanded-beam connectors instead of physical contact connectors. Expanded-beam connectors use optical elements (lens or mirror arrays) to expand the fiber/waveguide mode to a collimated beam with much larger diameter than the fiber core, thus relaxing lateral alignment tolerances between the beams, eliminating the need for physical contact, and reducing sensitivity to particulates and associated need for routine cleaning. Unfortunately, the mechanical tolerances for alignment of the fiber/waveguide to the expanding optics are still comparable to those for alignment of the un-expanded beams, and the angular alignment precision between connectors tightens. For this reason, there are currently no commercially-available expanded-beam array connectors for SM fibers.

There is thus a strong need for expanded-beam connectors suitable for SM fiber/waveguide applications, that can be fabricated and terminated in a high-throughput, low-cost, environmentally reliable process.

***Environmental issues with fiberoptic connectors.***

RoHS, or Directive 2002-95-EC, concerns the use of hazardous materials in electronic products. These materials include Pb, Cd, Hg, hexavalent chromium, PBB (Polybrominated Biphenyl Ether) and PBDE (Polybrominated Diphenyl Ether). Connectors have gone through the RoHS/WEEE (Waste Electrical and Electronic Equipment regulation) redesign or materials substitution cycle with significant (multi-million dollar) start-up costs, but without major roadblocks. There have been cost, logistics and supply chain issues. Connectors from well-established, reliable connector manufacturers are, for the most part, RoHS/WEEE compliant, or covered under exemptions. This has required major connector suppliers to focus significant internal resources on meeting these requirements.

Some current examples of exemptions applicable to fiber optics include:

- Pb and Cd used in optical glass: the original RoHS challenge has been met with nearly 100% of eligible production. Compliance cost is estimated at $60-100M. 20-30% of product & technical engineering resources were devoted to meet this challenge in the 2004-2006 timeframe, and ongoing efforts are being conducted in new and substitute materials, documentation, traceability, etc.
- Medical Devices will now be covered in both RoHS Recast Directive 2011/65/EU and REACH.
- Military applications are not within the scope of regulations.

In the United States, the National Electric Code (NEC) does not require halogen-free cable, but does require low smoke cable. It requires both good fire resistance and low smoke density if the cable burns, and it does require that cable be enclosed in conduit in riser cables and other applications. Jacketing such as FEP (fluorinated ethylene propylene) has good fire resistance but generates very toxic combustion gases. Toxicity is not covered in the NEC. European codes such as REACH consider toxicity as a third criterion and are banning the fluorinated polymers, hence, the halogen-free cables are used much more extensively and run at higher volumes resulting in equivalent pricing to the FEP materials.

With respect to halogen-free connectors: the connector industry is cognizant of potentially hazardous materials associated with providing flame retardance. Molex's position is one example of the industry's move toward non-BFR-CFR-PVC materials, and has adopted a conservative definition for this trend – one that meets customers' definitions. A product that has < 900-ppm (0.09%) bromine, < 900-ppm chlorine, and < 1500-ppm (0.15%) of bromine and chlorine combined, meets the requirement. Over time, prices will drop as supply catches up. In the meantime, a premium may be paid for these substitutes, including: polyethylene. fluorinated polymers (FEP, ETFE, PVDF), ethylene propylene diene elastomer (EPDM), polyurethane.

***Optical coupling of PICs to the passive optical interconnect***

As mentioned earlier, the optical mode sizes for the SM waveguides on PICs are typically smaller than fiber modes, and often less than 1 micron. Furthermore, the modes of the PIC waveguides are often not circularly symmetrical, as are fiber modes. Also, many processes for fabricating PIC waveguides result in polarization-dependent properties, so that they must be used with polarizing or polarization-maintaining fibers; these fibers have to be properly rotationally oriented as well as laterally aligned. These issues cause difficulties in alignment for coupling PICs to fibers and can limit the coupling efficiency that can be obtained at best alignment. Coupling can be improved by using mode-expanding "mode transformers" on the PICs or adding optical components (e.g. lenses) in the chip-to-fiber path. Nevertheless, < 1 micron tolerance assembly processes are needed for good coupling. Today, the lowest-loss coupling is achieved by launching light through the PIC-fiber combination, robotically adjusting the relative positions to maximize coupling, then gluing the parts in place (typically with UV cure adhesive, or dual UV + thermal cure). This is a slow process, and it requires skilled labor to initiate optical coupling between the parts to begin the robotic process.

It is emphasized that the above comments are generic and apply to all the various specific approaches to PIC-fiber coupling being pursued, including end-fire, grating coupling, and evanescent coupling configurations.

For the future, there is thus a need for new technology that allows rapid alignment and attachment of fibers/couplers to PICs. This could be based on robotic vision, or, better yet would be completely passive, where alignment would be achieved by precision mechanical interlocking structures, or by solder bump surface tension.

***Assembly of the chips/modules onto the PCB***

Currently, the optical coupling mechanisms used on PIC chips and in modules are not compatible with solder reflow. This is because there are often organic adhesives present, or injection molded polymer components (e.g. lenses or waveguides) that degrade at reflow temperatures. This means that the modules are typically manually placed in electrical sockets on the PCB after the electronic components have been attached in standard reflow. For cases where there are a small number of modules per PCB this may not be a serious drawback. However, for cases where there will many modules per PCB, the labor involved in plugging each module into its socket during the assembly process will be prohibitive. Additionally, the use of a socket instead of reflow introduces additional interfaces in the copper path, this leading to more signal degradation.

There is a strong desire for reflow-compatible PIC chips and modules; the problem is primarily one of achieving mechanically and thermally stable optical coupling. This suggests the need for inorganic precision optical components, as well as inorganic bonding agents (e.g. solder).

Note, however, that if the PIC modules are reflowed onto the PCB, rework becomes more difficult than for a socketed module, and has to be performed on a PCB of very high value.

***Routing of fibers/waveguides***

Even if the chips/modules can be attached to the PCB via conventional reflow, this cannot be done while fiber cables are attached. This is both because the fiber cables cannot tolerate the reflow temperatures, and because the cables would exert forces on the parts that would not be compatible with proper registration after soldering. To date, this issue is addressed by plugging cables onto the PICs/modules after they are mounted on the boards (also done manually today). This cable routing process is labor intensive, and the fly-over cables are reliability risks due to the potential for snagging. Attempts to address the routing and reliability issues have been made by laminating the fibers to carrier films, or by creating rigid, custom-contoured cables. These approaches are partial solutions, since they still require manual mating of the cable assemblies with the PCB components.

There is therefore a strong desire to replace temperature-sensitive, awkward, fly-over media with optical waveguides embedded in the PCB, where coupling of chips/modules to the PCB is automatically accomplished during reflow. Unfortunately, this is a difficult challenge, that involves new materials and fabrication technologies, as well as significant changes in supply chain and manufacturing infrastructure. Key new technology and infrastructure components required to enable the combined electrical-optical PCB include:

- Reflow-compatible, low-loss, PCB-embedded SM (perhaps PM) waveguides
- Self-alignment technology for positioning chips/modules relative to PCB optical ports
- Expanded-beam optical coupling technology for board-to-chip/module connections, potentially incorporating pitch transformers to convert from tight waveguide pitch on the chip to wider waveguide spacing on the PCB
- Board-edge coupling technology for blade-to-backplane connectors
- Design software for combined electrical-optical PCBs
- Manufacturing infrastructure for electrical-optical PCBs

### *Education and Training Needs*

Effective development of new optical interconnect technologies requires coordinated input from across a wide range of traditional disciplines. This is because optical interconnect modules and media present complex and coupled problems that span electrical engineering, semiconductor processing, mechanical engineering (especially the precision engineering specialty), guided-wave and classical optics, chemistry, materials science, polymers, ceramics, adhesives, metallurgy, and robotics. Typically, engineers become experts in interconnect not via study of interconnect as an academic discipline, but after having focused in some relevant discipline, become interconnect experts via long experience in the field. This path to interconnect expertise while currently working, is not a time-efficient way of building a large dedicated interconnect workforce.

At large companies which are sufficiently dedicated to optical interconnect as a core business, large cross-disciplinary teams can be assembled to provide all the expertise necessary. However, at companies with fewer employees, this may not be economically feasible. There is therefore a need for cross-disciplinary training that can allow smaller teams to effectively address interconnect development. This requires education that continues to provide a broad scope as a student advances to the Master and Doctoral degrees, rather than the traditional narrowing of scope found in today's technical education. Perhaps the most important aspect of the education is for the student to be trained to recognize and address tradeoffs between different requirements in the overall interconnect system, thereby contributing efficiently to overall system optimization. It is recommended that academic degree programs in Optical Interconnect Engineering be developed to address this current gap in training.

## GAPS AND SHOWSTOPPERS

For widescale implementation of optical interconnect in high-volume short-distance applications that offer the most growth potential, the most important near-term gaps and showstoppers are those associated with achieving cost-effective displacement of embedded high-performance copper interconnect. From the performance standpoint, optical interconnects have many benefits over copper that have already been discussed; these are widely recognized. Furthermore, there have been many "hero" demonstrations and high-end deployments in which optical interconnects have been successfully implemented in demanding applications such as world-class supercomputers and core routing.

However, such implementation has not taken place on a large scale because copper interconnect, though inferior in performance, for reach less than 3 meters has acceptable performance at a fraction of the cost of optics. Factors driving the high cost of optics include the following:

- Cost of cable termination
- Cost of optical connection to modules
- Lack of widely-available and reliable methods for analysis of overall cost-of-ownership vs copper solutions (parts, assembly, reliability, maintenance, workforce)

These factors will be considered in more detail below.

### *Cost of cable termination*

Today, the termination of optical fiber cables to connectors is primarily a manual process, generally performed by factory technicians. Steps in the process of applying a connector to a cable include most or all of the following manual steps: 1) separating the fiber from the cabling material, 2) stripping the buffer from the fiber and cleaning the fiber, 3) threading the fiber(s) into the tight-fitting cylindrical holes of the connector ferrule and fastening them, 4) cleaving the fibers 5) generating an optical polish on the end of the fiber, 6) assembling the ferrule into a connector body and strain relieving the cable, 7) testing and qualifying the completed connector assembly. None of these steps is currently automated to a significant degree.

One reason for the lack of automation is that the operations that must be performed are delicate, precise, and difficult to automate. Another factor is that the high variety of different product designs forces the need for flexibility of the automated equipment and also long setup times. This means that designing and building automated termination equipment will be expensive; such an investment is not justified by the current size of the cable assembly opportunity. Therefore, this issue has an aspect of the "chicken and egg" paradox: the process won't be inexpensive unless automated and producing in high volume, but high volume must be assured before the investment in automation is justified. Potential approaches for dealing with this showstopper include: 1) companies deciding to risk investing in development of automated termination equipment, based on the confidence that it will ultimately enable market growth and pay off, or 2) development of new connector technologies, specifically designed to enable cost-effective automated termination.

***Cost of optical connection to modules.***

The cost of making optical connections to single-mode optical elements (lasers, PICs, etc.) in optoelectronic modules has long been recognized as a dominant element of the module manufacturing cost. Estimates of the portion of the module manufacturing cost associated with optical coupling (assembly and testing) are as high as 80%; this cost is the result of the difficulty of aligning fibers or channel waveguides to PICs to submicron tolerances needed to optimize optical coupling between them, the need to maintain that alignment during initial curing of the bonding adhesive, and the need to optically test each connection to verify performance.

There are 3 main classes of alignment used in making optical connections to devices:

- Active alignment, where light emanating from the device is coupled into an output fiber or connector interface that is in turn coupled to a power meter, and the components are moved relative to each other using a precision positioner to maximize the detected power. This requires that the device be connected to an electrical or optical input, and the output fiber/connector be connected to a power meter. Although the movement to maximize the coupling may be done automatically, the connections to energize the device, and the connection to the power meter are generally done manually. This technique is used with both edge-emitting and surface-emitting devices.
- Robotic vision alignment, where the device is not energized, but a vision system is used to locate the optical emission area of the device (typically indicated by fiducial marks nearby on the chip), so that the fiber/connector can be placed in registration with the emitting area by a precision positioner. This technique is widely used with surface-emitting devices such as VCSELs and grating-coupled PICs but is difficult to use with edge-coupled devices, where fiducial marks are on the top surface and the waveguide exit is on the chip edge.
- Passive alignment, where via mechanical intermating features (e.g. etched grooves to align fibers to silicon waveguides), solder surface tension, or other effects, the fiber or connector interface can be aligned without the use of precision positioners. In the case of solder surface tension alignment, the components would be placed in rough alignment, then the solder would be reflowed to move the parts into adequate alignment. Such an approach would require no investment in specialized precision alignment equipment. This approach also offers the potential manufacturing advantage of performing many alignments in parallel, in batch processes or potentially even at wafer level. However, this approach has been very difficult to scale up in manufacturing, due to issues of cleanliness and friction at the micro-scale. And of course, there can be little compromise in coupling loss due to tight link budget requirements.

Once the fiber/connector has been aligned with the device interface, it must be attached in a way that is adequately stable under the conditions of device use and storage. Typical approaches today include UV light-cure adhesives, sometimes in conjunction with thermally-cured adhesives applied after alignment is achieved to improve stability, and solders.

Significant technology gaps exist in the 3 approaches above, especially relative to manufacturing cost. Active alignment is both capital and labor intensive, since it requires both high-accuracy robotic positioners, as well as human intervention to make the optical connections to device and output fiber before alignment can begin. Robotic vision alignment is also capital intensive, requiring the addition of robotic vision hardware and software to the high-accuracy robotic positioners. Furthermore, in both these cases the capital is poorly utilized due to the long cycle time associated with curing the adhesive and verifying alignment for each device serially.

Passive alignment offers the promise of the lowest-cost manufacturing, both from the capital, labor and throughput perspectives. However, to date, high-yield self-alignment of single-mode components having the required sub-micron dimensional precision has not been demonstrated.

***Lack of widely available and reliable methods for analysis of overall cost-of-ownership vs copper solutions (parts, assembly, reliability, maintenance, workforce)***

A realistic cost/benefit analysis comparing copper to optical interconnect in important target applications like data centers is very complicated, because there are many interrelated and conflicting requirements. Ultimately, customers want to transfer data between parts of their systems at the least possible cost for the data rates required; they are interested not only in the purchase price of the optical components, but also in the costs of installing, powering, cooling, maintaining and upgrading the entire installation. Quantitative trade-offs must be made between more-highly-paid optical technicians vs lower-paid copper technicians for installation and maintenance. Optical signals may dissipate less power at high data rates, reducing utility power consumption and air conditioning costs, but the initial investment in components is higher. Optical transmission requires the addition of optical sources and detectors, in addition to high-speed electronic laser drivers and amplifiers, whereas copper drive circuits can be integrated into the electronic logic. Optical connections, at least for distances less than several meters, do not require link length compensation transmission impairments, whereas high speed copper links require link-length-dependent drive compensation. High-speed copper connections require bulky cables, reducing the overall functional density of the system, whereas optical fibers are flexible and have a small cross section.

Ultimately, the gap here is the lack of a credible and widely available way of comparing the overall economic impact of the use of optics vs copper, to give system architects the confidence to make the shift from copper to optics.

***Lack of low-loss technology for integrating waveguides and couplers with PCBs***

Beyond the economic barriers, there is at least one major technology gap for longer-term implementation of optical interconnect at the substrate level. This is associated with the “Holy Grail” of optical interconnect, where optical waveguides would be embedded in PCBs much like electrical traces are today, and where optical connections between chips/modules would be made via processes as simple as solder reflow.

One difficulty in the use of today’s optical interconnect in interconnect-dense systems is the complexity and labor cost of installing large numbers of fiber optic cables. The cables must be installed after the boards are fully assembled (since the cables are not compatible with reflow), and as/after the boards are installed in the system rack (to establish the distance between connection endpoints and length of assembly needed). This is true even when backplane connectors are used, since the cables must be installed between mid-board modules and the backplane, and the backplane itself must be populated with cables. The situation is reminiscent of the days when electrical circuits were assembled using wire wrapping.

Ultimately, assembly and maintenance of these blade-in-rack systems would be dramatically improved if the optical connections could be handled like copper connections. That is, instead of fiber cables above the boards, optical signals would be carried by waveguides embedded in the boards. These embedded waveguides would route signals from module to module on the board, and to front panel or backplane connectors to destinations off the board.

Unfortunately, there is currently no practical technology for embedding waveguides in PCBs for transmissions over board- or rack-scale distances. While there have been lab demonstrations of multi-mode polymer waveguides operating near 830nm wavelength in the past, there have been no demonstrations of PCB-embedded single-mode waveguides at wavelengths of interest for silicon photonic or InP PIC transceivers (mostly 1310 nm, but some 1550 nm). This is because all polymers have carbon-hydrogen absorption bands that lead to excessive propagation loss.

One approach that has been suggested is to re-focus research on shorter wavelength emitters (e.g. visible), where polymer loss is adequately low. Unfortunately, such wavelengths are strongly absorbed in silicon, so that silicon photonics technology cannot be used, thus forfeiting the potential for full integration. Furthermore, to achieve multi-functional integration in a new waveguide material, single-mode waveguides would be required, so that the waveguide size and alignment tolerances would be reduced in proportion to the wavelength.

Low-loss glass waveguides fabricated via ion-exchange or laser-writing processes and operating at 1310 nm or 1550 nm are commercially available. A scalable process for fabricating glass waveguides at panel, rather than wafer scale, and then embedding them in PCBs in a way that is compatible with existing manufacturing practices, while not a fundamental gap, still needs to be worked out. A more critical gap is technology to couple a transceiver mounted on the PCB into the embedded waveguide with low loss in a scalable manufacturing process.

### *Other Potential Gaps and Showstoppers*

Beyond the strongly-economics-based interconnect related gaps and showstoppers discussed above, the following additional factors could interfere with further implementation of optical interconnect:

- Sub-Miniaturization Barriers to Conventional Fiber Optic Connector Technology: With connector housings at several millimeters and the optical fibers they encase at 125 microns, but waveguide cores less than 10 microns in size, there appears to be room for miniaturization.
- Sub-Miniaturization Barriers of the Electronic Packaging Platform (e.g. HDI, 3D, Printed Electronics): Requirements for pitch below 200μm require innovative electrical interconnect designs which may also require advances in micro-robotic assembly. In the former OEM-vertically integrated technology model, this would have been more easily possible – but is less likely today in the exploded global supply chain and multiple outsourcing of subsystems and assembly.
- PCB Development and Supply: Mainstream merchant PCB technology is not currently moving strongly in the OPCB direction — and many PCB/board assembly houses have low/no R&D budgets to do so. Flexible circuitry with embedded waveguides may come into play, especially for short-distance connections between PICs and module connectors; this is the boundary where subminiaturized FPC connectors are approaching minimum size limits.
- Barriers to Modularization of PIC photonic circuitry: This will require chip-to-chip optical interconnect, likely in the form of an interposer. At present, an open question is will be whether this will be designed by connector manufacturers, the semiconductor OEM, or the OSAT (Outsourced Assembly and Test) firms that do much of the packaging. Currently, this question is being considered by the IPSR-i Roadmap team, and could result in a new special interest group under the auspices of i-NEMI.
- Need for Manufacturable Optical Socket or Interposer Designs: Reflow-compatible optical-electrical sockets and interposers, preferable with standardized designs, will be needed to support the OPCB industry.
- Raw Material Cost Inflation: This has been a serious issue for connectors and other products; Cu, Ni, Sn, Au and many plastic materials have experienced significant price escalation and deflation cycles. Proprietary efforts by manufacturers have developed minimalist/substitute materials and processes to minimize the impact of these cycles. Still cost fluctuations have been reflected in higher prices and/or thinner margins. This is typically not a supply shortage issue, although there are some shortages that could result from recent globalization into unstable regions in Africa and a questionable China going forward. It is anticipated that solutions will continue to be found, combined with price increases where necessary. Recently there was a commodities deflation cycle which has stabilized with high demand for electronic materials.

- Part Cost: For commodity parts, competitive price pressures continue but are constrained by the maturity of this industry and its already having aggressively squeezed out costs, including via the use of offshore venues. However, these formerly low-cost labor locales are now experiencing inflation and higher labor and logistical costs, so that options for further cost reduction are limited.

### RECOMMENDATIONS ON POTENTIAL ALTERNATIVE TECHNOLOGIES

It is well understood that for distances of more than a couple of meters and at today's data rates, optical interconnect is the only technically viable solution. The combination of high channel rate, spatial multiplexing provided by multifiber cables, wavelength multiplexing, and nearly distance-agnostic signal quality has led to very high levels of fiber deployment.

However, for distances shorter than a couple of meters, the alternative, dominant and firmly-entrenched interconnect technology is high-speed copper. This takes the form of traces on high-performance PCBs, or specialized cables like twinax. Over time, remarkable progress has been made with copper media and associated electronics, such that for over 20 years there have been many predictions of the imminent demise of copper and large-scale adoption of optics…but this has never happened.

Reasons for the persistence of copper are many, but include: existing infrastructure for manufacturing in volume at acceptable cost, lower component cost/Gb/s than optics, familiarity of system designers, confidence in reliable performance, ease of maintenance, and the inevitable fear of change. Ultimately, the distance*bandwidth product has not reached the point where optics becomes the clear winner. 100 Gbps*m has historically been shown to be a transition point driven by economics as much as technology, although this is not an exact boundary between the two technologies.

While copper may not be a long-term alternative, it will likely be perceived as a low-risk alternative in the near term. Approaches to prolonging the dominance of copper could include:

- Modifying system architectures to minimize the length of high-speed paths where possible
- Development of new signal processing schemes to improve copper performance
- Transition from electrical traces across organic PCBs to flyover twinax cables

These possibilities will be considered further below.

#### *Modifying system architectures to minimize the length of high-speed paths where possible*

In current system designs, e.g. in data servers, it is common to disaggregate switching, storage, and routing functions between blades in a rack, or even different rack. This imposes the requirement of high-speed communication over many channels at distances of a few meters. As the level of integration in the chips performing these functions increases, it may be possible to combine these chips on single boards and in single packages, perhaps by using multi-chip module or 3-d chip integration packaging technologies. This approach would reduce the number of high-speed lines traveling more than a few centimeters is reduced. This could dramatically delay or reduce the market for short-range optical interconnects.

#### *Development of new signal processing schemes to improve copper performance*

Driven by the need to transmit more data over a relatively expensive telecommunications fiber infrastructure there has been a lot of recent attention given to more complex signal encoding and processing techniques. Multi-level modulation formats started in long-haul transmission with phase-shift keying (PSK) and are moving to quadrature amplitude modulation (m-QAM) adopted from traditional wireless communication applications now that coherent transmission is being adopted. In the short reaches within data centers, four level pulse amplitude modulation (PAM4) is rapidly being adopted for 100 Gbps to retain the use of simpler 25 GHz drive and receive circuits. In addition to signal encoding, digital signal processing in the receiver has become a powerful tool to recover linear impairments on the signal channel. While it is most powerful in long haul coherent systems where the full complex signal is recovered (real and imaginary parts of the electric field), transmitter and receiver equalization in direct direction systems is also becoming commonplace. Besides equalization, forward error correction has also been widely adopted from traditional wireless communications at the expense of latency and bandwidth overhead. However, much of the optical transmitter and receiver electronics developed for these approaches can equally well be applied to boost the performance of electrical transmission. The price paid for the improved signal rate through encoding and processing techniques is typically higher drive/receive circuit complexity, larger power dissipation, and lower signal-to-noise leading to increased bit error rate and/or lower reach. However, to minimize the number of channels transitioning to what has been more expensive optical interconnects users have been willing to make this trade-off.

***Transition from electrical traces across organic PCBs to twinax cable flyovers***

Today's short-reach interconnects such as within hyperscale data centers have direct attach copper (electrical connector modules and twinax cabling) between the servers and top-of-rack switches, and optical fiber between the TOR and higher switching levels. Within the switch boxes, the electrical signals are run over copper traces on the PCB from the pluggable optical transceivers on the front panel to the Ethernet switch ASIC package mounted on the PCB. As transmission speeds have increased, the signal loss of organic PCB material has become a challenge to maintaining overall signal integrity. Moving the optics from the front panel closer to the switch ASIC, either on-board or co-packaged with the ASIC is one alternative.

A competing alternative is to use lower loss twinax cable from the ASIC edge to the front panel optics. Electrical crosstalk, connector density, and cable management are some of the issues that need to be addressed with the twinax cable flyover alternative.

The evolution of optical interconnects from fiber toward PCB-embedded optical waveguides to manage high-channel-count I/O could be delayed by alternative technologies. One important possibility is:

***Heavy use of DWDM at the module level to reduce the number of fibers needed in a system***

The development of PIC transceivers capable of DWDM means that huge quantities of data can be transported by a single fiber. In telecommunications systems, for example, 64 wavelengths at 25 Gb/s each can provide aggregate data rates of 1.6 Tb/s. Ideally, this means that 1/64 the number of fibers is needed to transport the same amount of data as in a system where each fiber carries a single 25 Gb/s signal. This potentially not only reduces the fiber count, but also the connector count and complexity, and the labor associated with routing and managing fibers in the system. This can potentially extend the longevity of fiber cables as the transmission medium. In fact, long haul systems of ~100 wavelengths using dual polarization and QAM-16 for 400Gbps per wavelength (40 Tbps total) are available and being deployed today.

Of course, a critical difference between a data server and a telecommunications long haul line is that the telecommunication signals have a common (optical) destination, whereas this may not be true for the signals in the server. Thus, the use or DWDM may impose undesirable constraints on the system architecture. One key to making high levels of wavelength multiplexing feasible in short-length applications is devising an architecture where all the data on a high-capacity fiber are directed to the same location. As mentioned above, this becomes simpler as more highly integrated "systems in a package" with high-capacity ASICs and co-packaged transceivers are developed.

Another key element for enabling DWDM modules appropriate for use in a dense interconnect server environment is having cost-effective wavelength multiplexers that are temperature stable enough to be function properly when co-packaged with high-thermal-dissipation ASICs. Such multiplexers are not yet available, thereby driving the focus to simpler, more stable, easier to fabricate coarse WDM (CWDM) for these applications.

Not to be forgotten is that broadening the operational spectrum of the system to allow more WDM channels requires not only the development of the multiplexers, but also achieving low wavelength-dependent loss for all components of the system (including vertical couplers and other PIC waveguide devices) over the operating spectrum.

Looking out over the next decade the following associated alternative technology trends are expected.

- Trend towards more-highly-integrated SiPh and SoC; this will reduce/eliminate the need for many outboard connectors. This scenario may result in disaggregated functional modules connected by SM fiber
- The SiPh package may replace the outboard PCB assemblies with highly-integrated 3D PIC packages

**REFERENCES**

[1] "Linking with Light: High-speed optical Interconnects", N. Savage, IEEE Spectrum, Volume: 39 , Issue 8 , Aug. 2002

[2] J. Jou et al, "400 Gb/s optical transmitter and receiver modules for on-board interconnects using polymer waveguide arrays", OSA Continuum 1, p 658 (2018)

[3] Brusberg, Lars, et al. "Single-Mode Glass Waveguide Substrate for PIC Packaging." *2019 IEEE CPMT Symposium Japan (ICSJ)*, 2019, doi:10.1109/icsj47124.2019.8998659.

[4] M. Neitz, J. Röder-Ali, S. Marx, C. Herbst, C. Frey, H. Schröder, K.-D. Lang, "Insertion loss study for panel-level single-mode glass waveguides,"; Proc. SPIE 10109, Optical Interconnects XVII, 101090J (20 February 2017); doi:10.1117/12.2252802

[5] Fraunhofer/TTM EU funded project(s).

**CONTRIBUTORS**

Peter Maat: ASTRON, Netherlands (TWG Co-Lead), maat@astron.nl - Chair
Terry L. Smith: 3M Company (retired), USA (TWG Co-Lead), terrysmith9898@gmail.com - Chair
John MacWilliams, US Competitors, LLC, USA – Chair
Tom Marrapode: Molex LLC, USA
Alan Evans: Corning (retired), USA
Stephane Bernabe: CEA-Leti, France
Felix Betschon: VarioOptics, Switzerland
San-Liang Lee: NTUST, Taiwan
Patty Stabile: TU/e, Netherlands
Voya Markovich, (IPSR Substrates Chair), Microelectronic Advanced Hardware Consulting, USA
Mustafa Mohammed, DOW Corning, USA
Dana Korf, Multex, USA
Robert Pfahl, iNEMI/IPSR, USA
Marika Immonen, TTM, Finland
IPC 2015 Technology Roadmap

## APPENDIX B - CURRENT STATUS OF INTERCONNECT TECHNOLOGY

In this appendix the current status of optical interconnection technologies will be considered for applications with analog and digital communication links ranging from tele/data communication and sensor systems to phased arrays, medical systems, LIDAR in the automotive industry and automated manufacturing systems.

For applications outside communications, systems requirements are generally similar to the communications applications, that is, determined by the distance of optical propagation involved. However, there are areas where the different types of application requirements diverge:

- applications like medical/bioweapons sensing and radar may not be subject to the same cost constraints as typical communication applications;
- many sensing applications require the integration of microfluidic elements, which are generally not required in communications applications.

Below the situation analysis for interconnects in communication systems is given, this analysis also addresses the key interconnect challenges associated with other applications.

Application areas will be classified by the length of the optical link as follows, from longest to shortest:

1) rack-to-world connections;
2) rack-to-rack connections;
3) inter-blade connections between blades in a single rack (potentially through an optical backplane);
4) intra-blade connections between modules on a single blade;
5) intra-module connections within a module.

***Application area 1: Rack-to-world connections (lengths > 500m)***

For rack-to-world applications, single-mode fiber-based networks have long been dominant because of the high bandwidth*distance capability of the fiber. At current data rates (100 Gbps and below per wavelength channel), it is possible to mount pluggable transceivers at the edges of PCBs, to make access and replacement easier.

*Current product status:*

*Technology description:*

- Cages with electrically-pluggable transceivers incorporating thermal and EMI management.
- Optical transceivers (Tx/Rx) with-single mode fibers utilizing wavelength-division multiplexing (WDM) or multiple parallel single mode fibers.
- Single fiber or multifiber cables.
- Fiber distribution via patch panels and optical distribution frames (ODFs).

*Drawback of current approach in future systems:*

- Inefficient use of board edge space, leading to low density connectivity (fibers/cm$^2$).
- Thermal management due to air flow limitations and heat sink blockage by large, high-wattage transceiver packages.
- Power consumption, link impairments due to copper connectors and traces to board edge requiring repeaters, equalization and error correction.
- EMI/RFI mitigation due to high speed electrical signaling
- High-mating-force and dust sensitivity of multi-channel optical connectors.

***Application Area 2: Rack-to-rack connections (lengths 500 - 5m)***

These are also common today. As in the case of rack-to-world interconnects, pluggable transceivers are typical implementations at the moment, but Active Optical Cables ("AOCs" having transceivers permanently attached to fiber cables eliminating optical connectors and easing internal optical component interaction requirements), are also implemented for short-run applications where cable routing with transceivers attached is not too cumbersome. However, mid-board modules, interposer-mounted optical modules with transceivers and co-packaging of the transceivers with EICs on the same substrate are anticipated in the future.

AOCs provide a convenient copper-to-copper connection, enabling the user to treat the link "as if it were a fast copper link", upgrading easily, and requiring minimal understanding of fiber optics. AOCs and pluggables may be based on multimode (MM) or single-mode (SM) fiber, and may use multiple parallel fiber channels or wavelength multiplexing to increase capacity. Due to the shorter distances, and relaxed power budgets, MM fiber and multiple connectors in the signal path are more suitable, and VCSEL sources can be used. (Nevertheless many hyperscale data centers have converted entirely to SM fiber for both performance and future proofing.) However, because of the higher ratio of connectors and transceivers to fiber, there is more cost pressure on these components. In addition to use for 5 - 500m connections in supercomputers, data servers and telecom switches, links of this length may eventually find wide application in consumer high-definition television applications, as well as data transmission in automotive and avionics platforms. Such applications will be attractive due the reduced size and weight of the optical interconnects compared to copper of the same bandwidth.

<u>*Current product status*</u>

*Technology description:*

- Cages with electrically-pluggable transceivers incorporating thermal and EMI management.
- Optical transceivers (Tx/Rx) with multimode or single mode fibers utilizing WDM or multiple parallel optical fibers.
- Single fiber or multifiber cables.
- Fiber distribution via patch panels or pre-connectorized FO cabling.

*Drawback of current approach in future systems:*

- Inefficient use of board edge space, leading to low density connectivity (fibers/$cm^2$).
- Thermal management due to air flow limitations and heat sink blockage by large, high-wattage transceiver packages and ICs.
- Power consumption, link impairments due to copper connectors and traces to board edge requiring equalization and error correction.
- High-mating-force and dust sensitivity of multi-channel optical connectors.
- EMI management and agency certification testing which is more difficult with each increase in speed including variations across component suppliers.

***Application Area 3: Inter-blade optical connections (length 5 - 0.5m)***

On Board Optics are today only found in specialized high-end telecommunication and experimental systems, e.g. cutting-edge supercomputers. Transceivers typically mount mid-board on a PCB via an electrical socket. The optical transmission medium is typically conventional fiber cable, or fibers bonded to a flexible film (to provide fiber routing, shuffling and management), and may involve an optical backplane, typically configured from fiber. Optical connections between the medium and the transceiver may be via high-density, multi-fiber, remateable connectors (sometimes consisting of multiple ferrules ganged in a connector assembly) and/or fiber pigtails. In the case of fiber, direct runs across/between blades in a rack are possible, especially when there are a limited number of connections. The barrier to wider implementation is the cost of the optical systems relative to copper systems of the same capacity. As data rates, system size and densities increase, the cost difference does decrease, but, to date, not enough to justify optics in most inter-blade applications.

<u>*Current product status:*</u>

*Technology description:*

- On-board or mid-board-mounted optical modules with PCB surface mounted electrical interfacing connected via high density ribbon or small diameter multifiber cables.
- High density front panel or blind mating optical connectors utilizing standard MT ferrules supporting single mode and multimode fibers types or multimode expanded beam MT interfaces.
- EMI/RFI containment at the front panel via metallic or conductive polymer-based connector adapters

*Drawback of current approach in future systems:*

- Difficult routing of cables from PCB to PCB, due to no alternative to fiber media.
- No optical backplane technology for low-force mass connection of optical channels.
- Manual assembly of optical modules, cables on cards.
- Current lack of single-mode high-density expanded-beam ferrule interfaces.

Connections between blades in a rack, as mentioned earlier, are generally copper based at the moment. The dominant barrier to the use of optics is the fact that copper can provide the performance needed at today's channel rates of up to 100 Gbps. As data rates increase in the future, distances for which copper can be used will decrease (due to loss and signal distortion), and the cost/Gbps will increase, due to the need for better mechanical precision and higher-performance materials (e.g. dielectrics).

***Application Area 4: Intra-blade optical connections (length 0.5 - 0.05m)***

These are connections across a single blade. These have not been commercially implemented to date because data rates have not yet reached the point where optical communication is required to address copper interconnection impairments. For future systems it is expected that multiple electro/optical modules will be placed on a single blade. Optical interconnections between these modules can be realized via both interposer- and/or PCB-embedded-waveguide-based optical interconnects.

<u>*Current product status:*</u>

*Technology description:*

- Module-to-module optical interconnects on a PCB are not in general use today, due to the drawbacks below, and the acceptable performance of copper at current data rates..

*Drawback of current approach in future systems:*

- Manual application of fiber/fiber-flex based interconnections
- Manual application of modules to PCB—plug into a pre-mounted socket
- Process incompatibility (e.g. reflow of modules on boards with optical interconnect polymer components, optical fiber)
- Limited edge density of copper
- Bandwidth limitation of copper traces on the PCB
- High attenuation at high frequencies of copper traces on the PCB
- High cost of optical transceivers

***Application Area 5: Intra-module optical connections (length < 0.05m)***

These are connections inside a module package. Integrated electronic/photonic modules require high-density low-cost, low-optical-loss assembly technologies that provide an integrated system with a proper mechanical stability for adequate reliability and lifetime. Such connections are already in use today in two types of applications: connections from a laser or PIC to a connector interface at the module wall, and connections between a 2 or more PICs that are present in a single package. In the former case, the connection is usually in the form of a short array of fiber stubs (the dominant commercial approach) or polymer waveguides. In the latter case, the optical connections can be realized by (1) direct coupling using either end-fire coupling, relay micro-lenses, fiber stubs or by direct evanescent optical coupling between optical waveguides by placing PICs on top of each other, or (2) via the use of an interposer to which multiple chips are optically coupled. The most common approaches today use either fiber stubs or relay micro-lenses to couple lasers to waveguide chips. In the case of an interposer the optical (and electrical) interconnections are realized via a (e.g. Si or SiN) sub-mount.

*Current product status:*

*Technology description:*

- Optical waveguide media to couple chips within a module, or to couple a chip in the module to a connector interface at the module package boundary. Examples include short sections of polymer waveguides that connect PICs to MT connector ferrules, or fiber stubs that connect lasers to waveguide chips such as modulators.

*Drawback of current approach in future systems:*

- Slow alignment and attachment of coupling medium to package connector.
- Slow alignment and attachment of coupling media to PICs.
- Manual application of modules to PCB—plug on.
- High loss per length of polymer waveguides at silicon photonics operating wavelengths.
- High manufacturing cost.

***Roadmap starting point***

In the main body of this analysis, detailed future technology options and their status, needs for new technologies to advance short-range interconnect, infrastructure considerations, and associated roadmap milestones will be discussed in the context of the above application areas.

The development track of substrate/interconnection technology depends strongly on the timing of the transition from Cu signaling to photonics at the chip, package and board level of datacom and computer/server/storage equipment. This transition is dependent upon manufacturing cost reduction for optical solutions as much as on performance relative to copper.

The roadmap starting point for optical interconnection technology can be described by its development stages of the past years:

- 2015-18: Hodge-podge of co-existing proprietary, company-specific and standard interconnect designs, which do fulfill existing applications, if at a high cost. Existing fly-over optical fiber interconnect, mostly MM with no PCB-embedded waveguides.
- 2018-20: Initial attempts at evolution of standards based on an interim hybrid approach to photonic chip packaging, parallel to what exists today in InP transceivers. First use of SM in electronic packaging with discrete hybrid transceivers. First use of embedded waveguides with peripheral interconnect and SM connectors and cables.