TRANSCEIVERS AND INTERCONNECTS

EXECUTIVE SUMMARY

The data explosion generated by the increasing use of social networks, digital entertainment, cloud-computing, and the internet of things (IoT) is radically driving the growth of data centers and the need for high bandwidth, low-latency, low power consumption, optical fiber-based communications. Since 2022, this growth has accelerated to unprecedented levels due to the rapid adoption of artificial intelligence (AI) in general and generative AI in particular. The dynamics of the market have changed. The datacenter (DC) industry, represented by companies such as Google, Facebook, Microsoft, Amazon etc., not telecommunication companies are driving next-generation technology platforms and metrics. These forces are transforming data center architectures to a higher level of integration of photonics components with fiber interconnects that are becoming more intelligent and higher performing.

The biggest change over the past 3 years in optical interconnects is the commercial introduction of co-package compatible Ethernet switch ASICs and associated optical engines (i.e., transceiver photonic integrated circuits (PICs) without their mechanical envelops). This significant event will spur a move to innovative, lower cost detachable connectors from legacy technology - fiber array in glass V-grooves, known as fiber array units or FAUs) that are actively aligned by optical power-peaking and then bonded to the PIC with UV/thermally cured epoxy. Total cost of ownership dominates the decision process for technology adoption and fiber attach is one of the highest categories of the cost pareto. In addition, FAUs, as a 30-year-old technology, is well down its cost reduction curve. The expected interconnect cost reduction from new, lower-cost, co-packaged PIC-to-fiber solutions should motivate front panel, pluggable transceivers are not going away and have a path to 1.6 Tbps and perhaps 3.2 Tbps. The recently announced OSFP-XD multi-source agreement (MSA) follows the historic path of increased parallelism (more channels) but existing and limited faceplate area will drive innovations in small form factor connectors. Likewise co-packaged connectors (both PIC and front panel) will need to be small, and cross-fertilization of space-constrained solutions are anticipated. Another important trend will be automation, which will be critical to achieve the rate of manufacturing predicted by transceiver sales volume. And new technologies such as glass waveguides and expanded beam optics will provide next generation photonic solutions leading transparency of outsourced assembly & test (OSAT) that finally fulfills the promise of silicon integrated photonics, not only tapping into the scale possible with semiconductor wafer-scale fabrication and the leveraging of the silicon integrated circuit ecosystem, but with back-end high speed assembly and packaging. “At this point, high-volume manufacturing abilities are critical to meet the demand for data center optics.” [1]

INTRODUCTION

Three major segments exist in the datacom and telecom market: Telecom core/metro, telecom access, and datacom. While this report will continue to use the phrase “datacom”, it is important to point out interchangeably using the phrase “datacenter” would be acceptable given the dominance of this sub-sector in driving revenue and innovation in this segment. The growth of the telecom core/metro and datacom segments are expected to experience very strong growth over the next decade. Within that growth is the trend to photonics components integrated in PICs which is beginning to accelerate driven by applications that require smaller photonic component solutions, lower power consumption, higher data rates, hybrid packaging, longer interconnect lengths, and scaled economics in terms of $/Gbps. In particular, PIC transceivers are expected to grow quickly over the next decade from the rapid ramping in deployment of 800Gbps systems using 60 to 80GHz bandwidth photonic devices. Within 5 years, 1600Gbps and 3200Gbps are expected.

The obvious yet most difficult next move is to increase the optoelectronic modulation and detection speed from 60-80 GHz to beyond 100 GHz. Recall that 40 GHz analogue bandwidth roughly corresponds to approximately 50 Gbaud NRZ using standing coding techniques and 100Gbps PAM4[1] while 70 GHz analogue bandwidth corresponds to ~100 Gbaud NRZ, and 200Gbps PAM4. Subsequently, 100GHz analogue bandwidth corresponds to ~150 Gbaud NRZ, and 300Gbps PAM4. At the same time, these new optoelectronic devices must be very small, and operate with very low voltage to keep power consumption low.

Optical interconnects focus on optical connections between photonic integrated circuits (PICs) in either computer servers or Ethernet ASIC switches and other equipment used in cloud and hyperscale data centers and other high-performance data communications applications such as AI and machine learning (ML). This market has significant strategic importance due to the transition from discrete server systems used in millions of businesses and institutions (enterprise DCs) to large, independent, Cloud Data-As-A-Service (DAAS) providers. Hyperscale data centers (HDCs) are becoming THE key elements of our future information technology infrastructure. The emphasis is on DCs since this application was the first and remains the dominant one for PIC optical interconnects.

For applications outside communications, interconnect systems requirements are generally similar to 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:

- 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, may require wavelengths of operation outside the standard tele/datacom 1.3/1.5 µm bands, and have potentially lower cost and more relaxed loss targets
- Quantum applications where loss targets are significantly lower than DC
- Neuromorphic computing based on integrated photonic accelerators requiring high speed and highly parallel interconnects between the accelerator and CPU/memory to overcome the electronic memory and pin constraints; photonic links could provide the necessary scale. [2]

In the DC application space, silicon PIC (SiP) transceivers have taken over the 100-to-500-meter space due to their cost competitiveness and ability to meet link requirements. Much of the focus is on SiP but also applies to other material systems (i.e., indium phosphide, InP, and gallium arsenide, GaAs). In fact, the manufacturing scale section parses the three material systems.

PURPLE BRICK WALL FOR TRANSCEIVERS

The transceiver/modules roadmap that displays the purple brick walls is shown below in Figure 1 and Figure 2. Figure 1 shows the individual purple brick walls for each technology segment. Figure 2 shows the purple brick wall challenges that require focus to achieve the metrics laid out in the transceiver/module roadmap.

[1] https://www.edn.com/the-fundamentals-of-pam4/

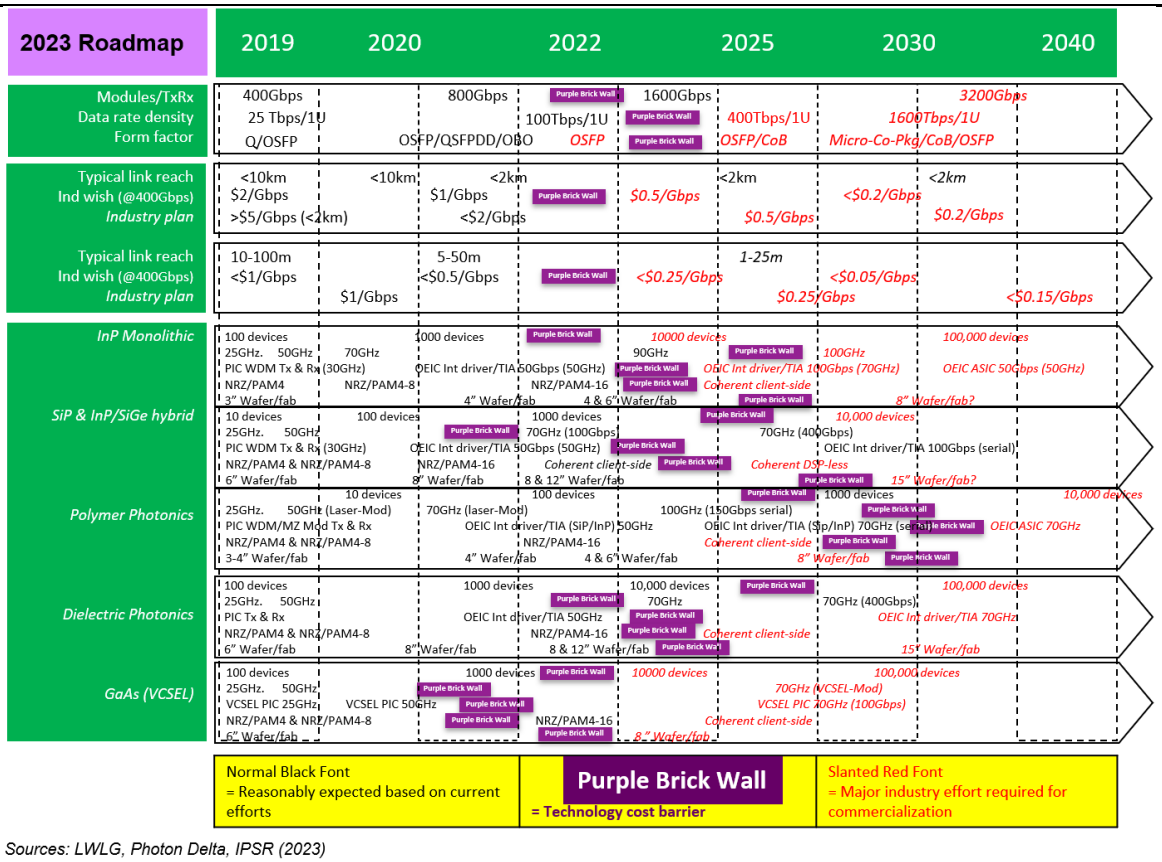

Sources: LWLG, Photon Delta, IPSR (2023)

Figure 1: Transceiver/module roadmap showing purple brick walls for technologies that will drive transceiver design (Sources: LWLG, IPSR, PhotonDelta)

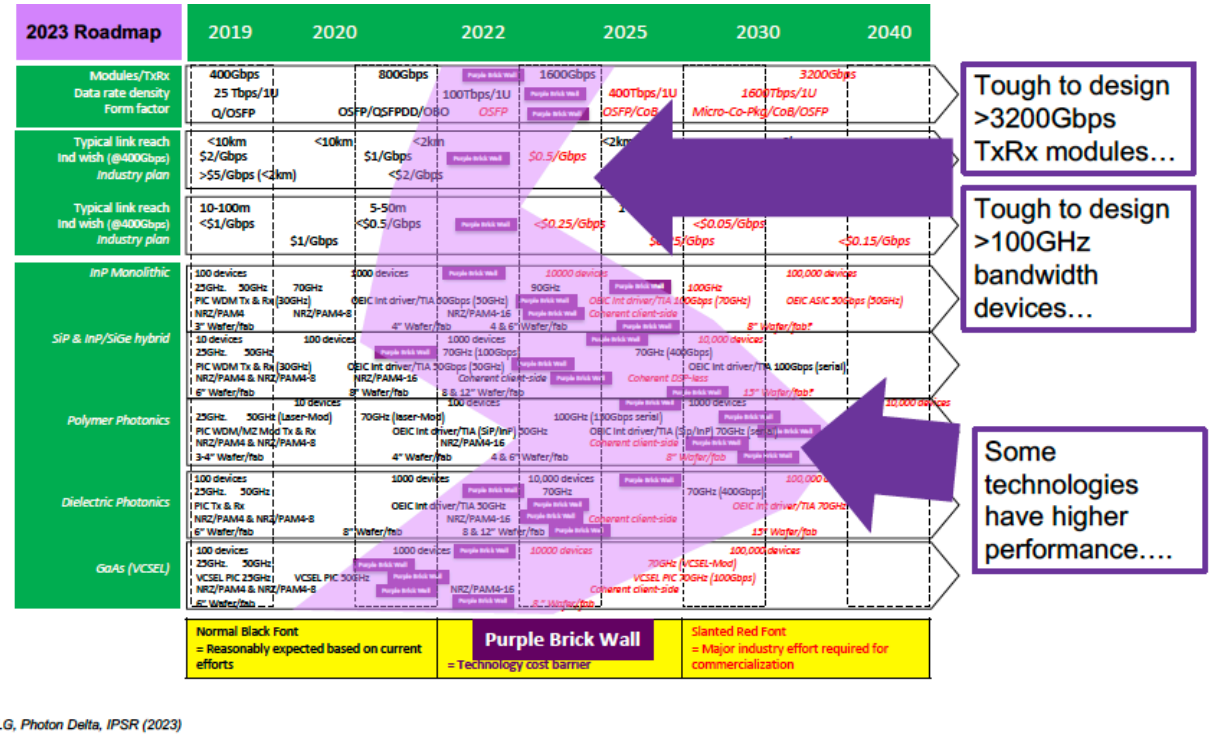

Sources: LWLG, Photon Delta, IPSR (2023)

Figure 2: Transceiver/module roadmap showing the key challenges identified by the purple brick walls for technologies that will drive transceiver design (Sources: LWLG, IPSR, PhotonDelta)

Table 1. Purple brick wall barriers – critical needs for transceivers/modules

Purple Brick Wall	Description
Product vehicle	>3.2Tbps and 6.4Tbps optical transceiver using multichannels
Form factor	Higher density and smaller form factor than OSFP today
Lasers	High operating temp (>150C); wavelength stable, uncooled 1310nm
Low noise integrated opt amplifiers	Noise factors for optical amplifiers <5dB
High speed, low power modulators	>100GHz 3dB bandwidth (EO); 1Volt or less for direct drive from CMOS
Waveguide loss	Low loss in visible wavelengths <0.01dB/cm
Temperature of operation	0-85C uncooled
Power consumption of TxRx	<1Tbps/10W
Full PIC implementation	Includes both Active and passive integrated components on PIC platform

PURPLE BRICK WALL FOR OPTICAL INTERCONNECTS

The optical interconnects roadmap, shown in Figure 3, is organized into Product Vehicles, Metrics, and Technologies sections which are described below after the transceiver section. Unless stated otherwise, the event dates mark the year of commercial adoption. Table 2 summarizes the interconnect critical technology needs also known as the purple brick wall barriers.

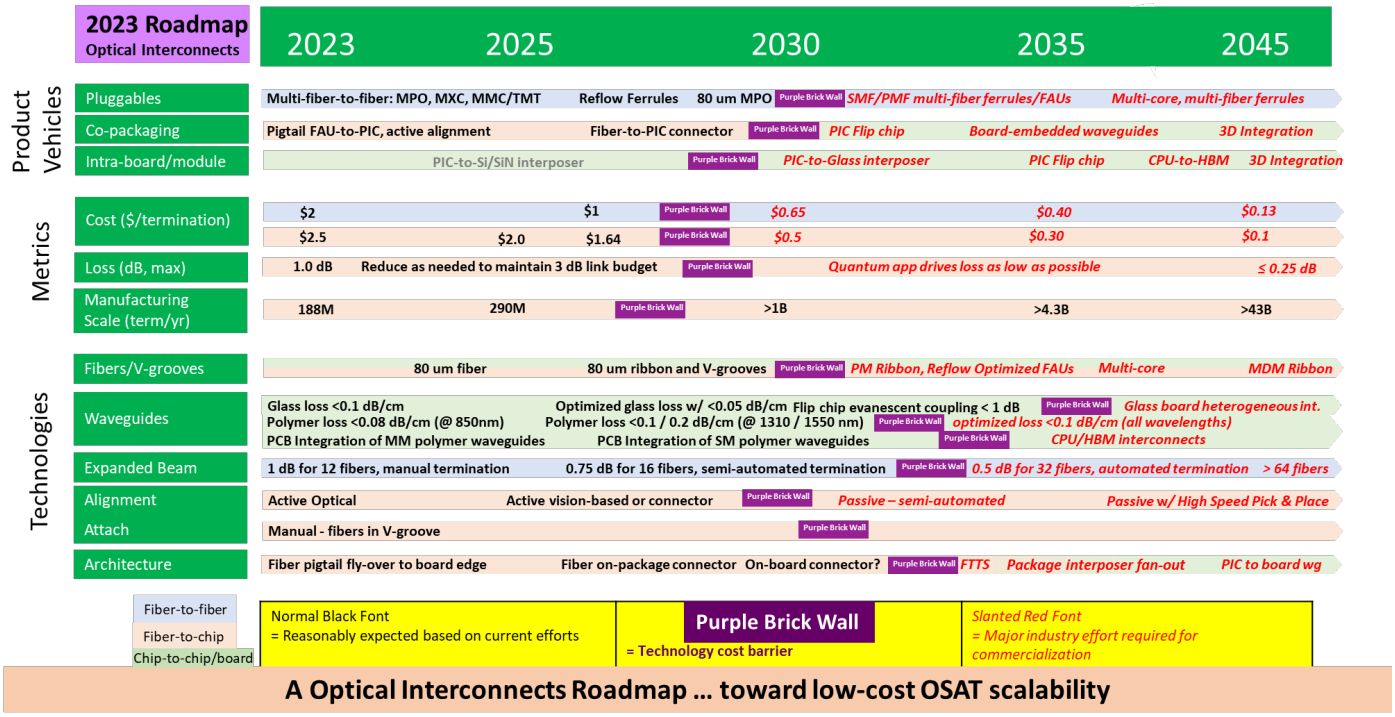

Figure 3: Optical interconnect roadmap showing the key challenges identified by the purple brick walls for technologies that will drive optical interconnect design (Sources: TWG Interconnect working group, IPSR, PhotonDelta)

Table 2. Purple brick wall barriers – critical needs for optical interconnects

Purple Brick Wall	Description
Pluggable	Next generation, small form factor connector
Co-packaging/intra-board	Flip-chip PICs w/ electrical and optical interconnects
Cost	< $1 per termination
Loss	< 1 dB max
Manufacturing Scale	Scaling to 1B terminations per year, intra-shelf fiber management
Fibers/V-grooves	PM ribbon, reflow-optimized FAUs
Glass Waveguides	Electrical/optical heterogeneous board integration
Polymer Waveguides	XSR/USR optical interconnects
Expanded Beam	0.5 dB loss for 32 fiber array, automated termination
Alignment	Passive – semi-automated
Attachment	Semi-automated – latching/mating
Architecture	Fiber-to-the-server costs

TRANSCEIVER ROADMAP CHALLENGES:

Several vendors are looking at how to address 200G line rates, and are planning spatial multiplexing (adding fibers), wavelength division multiplexing, (adding wavelengths in a single fiber), encoding with more complex symbols per bit (PAM, QAM, etc.), designing optical and electrical devices for higher bandwidth, lower voltage (power), and smaller footprint (size) as part of a PIC platform. These higher performance drivers are being accelerated by data hungry customers such as data centers, high performance computing, and shorter reach telecommunications.

Yet the main roadmap challenge for the fiber communications industry is to plan future data line rates that exceed 200G and extend towards 300G and 400G. These line rates are needed to achieve pluggable transceiver and/or co-packaged transceiver aggregate data rates of 400Gbps, 800Gbps, 1600Gbps (or 1.6Tbps), 3200Gbps, 6400Gbps and beyond. While companies are looking to complete the designs for 200G line rates in the 2023-2025 timeframe, they are also looking how to extend the data rates and line rates further[2]. Interest in line rates of 300G and 400G are still at an early stage, however, active electro-optic polymers offer modulator device demonstrations today that achieve these goals[3].

Processors, storage, networking and switching fabric performance, and their future trend toward increased integration, are all driving the packaging requirements for data centers. Whereas thermal, electrical, photonic and mechanical metrics constrain and direct development of that packaging. Of these, thermal and power management of the system continue to be major challenges for data centers. A key component of the total cost of ownership is the cost of the energy to power the high-end equipment. And, in some locations, the amount of energy that can be delivered to a data center is at the limits of what the electrical utility can deliver. To achieve reduced energy use, the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) is creating standards where the temperature and humidity can be higher than has been traditionally required for high-end equipment. ASHRAE recommends that server inlet temperatures be between 18 and 27 degrees Celsius (64.4 to 80.6 degrees Fahrenheit), with relative humidity anywhere between 20 and 80 percent. While less energy is required to cool the data center, there is a need for better thermal technologies within these

[2] Arista presentation from OFC 2022 for OSFP pluggable optical transceivers

[3] [3] https://www.lightwavelogic.com/presentation/polymer-modulators-with-50ghz-performance-for-power-consumption-reduction-at-400-800-and-1600-gbaud-aggregated-datarates

systems to maintain acceptable junction temperatures.
Likewise, active power management of individual components to maximize the power performance of the system is also a desirable capability. More efficient power conversion will also continue to be developed including voltage regulation close to the loads (such as in microprocessors) that significantly cut the power distribution losses within the system.

Packaging technology must also increase at a rapid rate to achieve the new performance requirements. As breakpoints are reached, lower loss laminates, smoother copper, and higher bandwidth connectors are required. Development of these components in high-volume manufacturing with cost-competitive materials and processes is ongoing. The reader is encouraged to read the Packaging and Assembly roadmap chapters for further details.

Optical interconnects are becoming more widely used with VCSEL data rates reaching 50 Gbps and silicon photonics becoming commercially viable. The need for increasing bandwidth to move the data as well as the growing size of data centers defining the distance the high bandwidth interconnect must travel is creating more need for optical communication. At this time 3 meters has the potential to be the breakpoint between copper cables and optical fibers, but this is expected to fall to 2 m and 1 m over the next decade. As aggregated data rates surpass 400 Gbps, and increase to 800 Gbps, and even 1600 Gbps, on-board photonics (i.e., printed circuit board) and more likely on-package photonics (i.e., co-packaged transceivers and electric chips on the package substrate) will be required.

However, the biggest challenge is transmitting ever increasing data streams with minimal latency to keep up with computational demands. With processor clock frequencies stagnated, data centers are designed as CPU/GPU clusters with all-to-all compute server access and ultimately disaggregation of compute, memory, and storage. Disaggregation requires ~10 ns CPU-to-CPU access and ~20 ns CPU-to-memory access which runs counter to the added latency (and increased power) of digital signal processing necessary to clean up electronic signals to maintain their integrity. Optical transmission which doesn't need this high level of signal conditioning encourages the move toward optical off-die data input/output. And while optical transmitter and receiver designs are driving power requirements down, adding more lanes/channels to increase bandwidth is limited by the expectation of decreasing price and overall power. This results in the current trend of quickly increasing data rates per physical channel.

As an aside, analog pre-emphasis as a low power equalization technique has been highly effective in boosting the bandwidth and data rates of VCSELs to 50 Gbps. While individual VCSEL transmission systems are composed of discrete components, arrays with combined functions are considered as integrated photonics particularly within the manufacturing scale interconnect terminations per year section of this report. Note that current VCSEL transmission systems are multimode and limited in distance by the chromatic dispersion of multimode fiber. For example, OM5 fiber has an effective modal bandwidth of 4700 MHz*km at 850 nm wavelength of operation or 100 meters of 100GBASE-SR4.

The above energy limiting factors in addition to efficiency, latency, and the growth of data transmission, place a challenge on the processing power needed to perform any data analytics because the total available power is split between the processing or compute power and the networking or transmission power between the processing units. Instead of adding general-purpose cores to a chip and enabling those cores to support additional threads, an emerging trend, perhaps more efficient, is to employ specialty cores to do specific tasks. The use of GPUs and FPGAs allow quick turn-around time to quickly adopt more compute and energy efficient algorithms while special ASIC cores are developed to address specific computational tasks. As the big-data era matures and 2 and 3 nm silicon nodes are developed, the packaging technologies to interconnect these components will also evolve. The challenge is to keep the momentum of scaling performance with cost, as the price of designing using 10 nm (or less) silicon nodes is becoming higher.

TRANSCEIVER TECHNOLOGY NEEDS

Data bandwidth demand is resulting in systems with ever faster interconnect speeds, even as processor speed is staying constant. The enormous size of the hyperscale data centers creates a huge demand for electrical power and an increasing focus on power efficiency. So, data center operators try to reduce power consumption and operation costs to manage the total cost of ownership. PIC-based transceivers as the low cost and low power solution have become the dominate solution at 500 meter intra-data center distances and are looking to expand to shorter, 100-meter high-performance computing, and longer, 2 km, campus-wide reaches.

Photonic packaging technology must create disruptive solutions to achieve the new performance requirements, in particular faster photonic devices such as 100 GHz modulators that can operate at 100 Gbps NRZ or 200 Gbaud PAM4. Today, one technology close to meeting these speed metrics with low power consumption are polymer modulators[4]. There are other modulators that are being developed with both the potential for high speed and low power that include: slot modulators, plasmonic slot modulators, silicon-based GeSi EAMs, micro-rings, Barium Titanate (BTO), and thin-film lithium niobate.

Data centers need faster optical devices (such as 100+ GHz components), higher integration levels, higher reliability, lower power consumption, a higher degree of scalable economics that address $/Gbps metrics, and smaller, more miniaturized foot-print platforms. Co-packaged solutions with hybrid integration between different technologies are also required. The integration of electronics, photonics, and packaging is an important topic that the industry will explore more intensity over the next decade.

Needs 2025

- Device speed increased (modulator bandwidths EO S21 of 100 GHz in PIC platform
- Drive voltage at 1V so that drivers can be eliminated to enable linear drive optical architectures.
- Co-packaging of electronics and photonics, OSFP and QSFP-DD lower power designs
- Hybrid integration of InP lasers with Si photonics

Needs 2030

- Device speed increased (modulator bandwidths EO S21 to 150 GHz+ in PIC platform
- Laser operation uncooled to >100 ºC
- Reduced power consumption (50% reduction), with architectures such as Linear Drive Optics
- Integration of photonics and electronics

Needs > 2040

- Device speed increased (modulator bandwidths EO S21 to 200 GHz+ in PIC platform
- Laser operation uncooled to >150 ºC
- Reduced power consumption (80% reduction), with architectures such as Linear Drive Optics
- Further integration of photonics and electronics
- Uncooled coherent laser with better than 2 GHz stability

INTERCONNECT PRODUCT VEHICLES

Product vehicles are organized by packaging architecture, not application reach. Application areas are described in the Appendix's Situational Analysis section.

[4] https://www.lightwavelogic.com/presentation/polymer-modulators-with-50ghz-performance-for-power-consumption-reduction-at-400-800-and-1600-gbaud-aggregated-datarates/

PLUGGABLES AND CO-PACKAGING

Optical interconnects in data centers, either between Ethernet switches or between switches and servers (i.e., high-powered computers), are created with transceivers at each end of an optical fiber link where the fiber between equipment racks is known as structured cable. Today, transceivers are available in numerous form factors (e.g., QSFP, OSFP), defined by multi-source agreements (MSAs) and IEEE standards, and plugged into the front panel of a server or switch shelf within a data communication rack. These "pluggable" transceivers have numerous benefits: ability to mix & match depending on customer need (e.g., choose right data rate for each individual modules); low cost due to MSA "standardization"; modularity for pay-as-you-grow economics; hot swappable for easy adds/changes and repair; and ever-increasing bandwidth within the same mechanical envelope. [3] And the industry has a path to 1.6 Tbps data rate with the recently announced OSFP-XD MSA staying with current and therefore less risky 100G optics and 16+16 I+O channels putting the speed increase burden on increasing the number of parallel channels. [4] [5] Yet signal integrity across the printed circuit board (PCB) electrical transmission lines from the server or switch to the front panel has become very challenging as the electrical interconnect data rate doubled from 28 Gbps, NRZ (non-return to zero modulation format) to 56 Gbps and now 112 Gbps, PAM4 (four-level pulse-amplitude modulation format). [6] Increased forward error correction (FEC), transmitter equalizer pre-emphasis, and receiver digital signal processing (decision feedback equalization, DFE, and/or feed-forward equalization, FFE, and continuous time linear equalizer, CTLE) overcome loss from electrical signal impairments like loss from the skin effect and radiation, multi-path interference (MPI) from impedance mismatch reflections, and crosstalk from neighboring channels. But this is at the expense of higher power consumption (more than 50% [6]) and resulting thermal dissipation. The benefits of co-packaging are so great that optical engines are even becoming available for previously released 25.6T Ethernet switches. [7]

Since the OIF's 112G VSR PAM4 CEI IA (Optical Interface Forum's Very Short Reach Common Electrical Interface Implementation Agreement) specifies a 20-cm reach across the PCB which is not quite sufficient for the full wide of the 19" front panel, the latest 51 Tbps switches use more power hunger LR SERDES (long reach, serializer/deserializer) with 1 meter reach. [8] The LR SERDES allow for a 100G PAM4 interface to direct attach copper (DAC), front panel pluggable optics, and/or co-packaged optics, but consumes far more power than switch ASICs likely to be introduced in the next generation that should be optimized for co-packaged optics with VSR SERDES. In other words, the power consumption required for electrical signal integrity is motivating the move to co-packaging the transceiver (stripped down as an optical engine without the mechanical envelope) on the same package substrate as the Ethernet switch ASIC and the use of VSR or even XSR SERDES, resulting in 50% lower power consumption down to below twenty or even ten pico-joules per bit. [6] [7] While this is a huge benefit, issues are currently being resolved with yield, rework, reliability, concentrated heat generation, and establishing a robust supply chain including the role of OSATs. While the roadmap refers to Ethernet switch ASICs within data centers since this is the initial off-chip bandwidth driving application, in the future many other high bandwidth ASICs such as FPGAs (e.g., for direct-RF sampling massive MIMO) [9] and CPUs/GPUs for artificial intelligence will rely on co-packaging. In fact, co-packaged optics (CPO) specifically designed for AI/ML has been announced with direct drive from the host ASIC SERDES to the CPO. [10]

Incumbent pluggables and emerging co-packaging optics (CPO) are and will continue to be similar in PIC footprint and channel count since transceiver companies will want to develop the same PIC for both applications. At present, the difference is in the length of fiber ribbon jumper between the FAU attached to the PIC and the outside world: ~1 cm from PIC-to-pluggable front plate; ~30 cm from PIC co-package optical engine to the shelf front panel. Besides jumper length, another difference for co-packaged PICs supplied with external lasers are fibers (typically polarization-maintaining) from external laser sources to the PICs. Various options exist for external laser source delivery to the CPO. [11] Pluggables have either internal lasers or micro-optic beam delivery. Co-packaged optical engines could move in this direction too if high-temperature laser operation can be achieved which is on the transceiver roadmap.

Since the pluggable face plate is somewhat more space constrained than the pass-through front panel connectors

of CPO, it is anticipated that pluggables will drive the connector cross-sectional space saving innovations that joints the pluggable fiber jumper to the outside-the-box fiber and that then it will be adopted to the CPO front panel application. Recently introduced MCM connectors with TMT ferrules offer the highest density, commercially available 1x16 or 2x12 fiber with 250 micron pitched MPO-type connectors. Since 80-micron diameter fibers are available, the next logical step in density could be developing fiber ribbon that when interleaved provides 125/127-micron pitch in a molded ferrule connector. Ultra-small connectors that require tools for installation rather than a human hand are another logical step but require industry consensus which has proven to be difficult. Beyond ~2035, increased space-division multiplexing I/O density within a standard molded ferrule connector can probably only be solved with multi-core fibers (MCF) or, less likely, mode-division multiplexed (SDM) few-mode fibers. Rotational orientation of a single multi-core fiber within a V-groove to align cores ~25 microns in pitch is commercially available and it has been demonstrated in fiber connectors. The purple brick wall challenge is the efficient orientation of multiple multi-core fibers. Fiber ribbon makes molded ferrule and V-groove fiber assembly efficient but individual fiber orientation within the ribbon is extremely difficult due to current processes. While MCF in short reach applications avoids the complication of amplification, depending on network architecture, it may need fan-outs from MCF to individual fibers. This adds cost and loss unless somehow integrated within the photonic circuit. Hence the best use of MCF if it can be ribbonized is in point-to-point links that reduce space particularly in structured cable overhead trays if the fanout can be integrated in the PIC. An implication is the need for a linear, not hex-packed, core geometry.

For co-packaging with typically 16 optical engines surrounding a central switch ASIC, [6] fiber management schemes and designs for physical separation (optically, electrically, or both) are critical. See the reference for a more complete description of options and tradeoffs. [11] Adhesively-bonded FAUs are manageable in the first generation but detachability along with latching is highly desirable for package-to-board assembly & rework, e.g., solder reflow and fiber routing. A recent announcement of a PIC connector in development, could mean ~2025 for commercial availability, assuming a 2- to 3-year development cycle. [12] Expanded beam optics (EBO) are also in the running and being actively developed. [13] [14] At some point, perhaps once more than one fiber ribbon connector per optical engine is required, more advanced fiber management schemes could include on-board connectors and ultimately a drive toward embedded waveguides and routing within, or on top of, the PCB.

The next likely development is flip-chip assembly of PICs for the benefit of efficient electrical interconnect packaging. Initially there may not be any change to optical interconnects but to take full advantage of the assembly efficiency, developing optical coupling either to, or on top of a package substrate interposer will be necessary and require significant technical development (hence a purple brick wall event). Interposers increase the available optical I/O shoreline by accessing the edge of the interposer rather than the PIC and, for the right dimensional-stable materials, can provide a dimensionally stable surface with alignment elements for passive fiber array alignment. In fact, the optical IO shoreline could be greater than that provided by a substrate with electrical vias to the BGAs attached to the PCB in a cantilevered design. Hence a glass package interposer substrate is on the roadmap prior to board-level waveguides. Eventually, as stated above, embedded glass waveguides within the PCB or specialized interposers are likely a solution to keep up with the ever-increasing number of PIC channels needed due to relentless switch ASIC bandwidth doubling every two years. Already, the number of fibers is 1024 for DR-based 51 Tbps switches. Interconnect routing will be simplified with low-loss glass waveguides once efficient, low loss coupling to those waveguides is developed. Polymer waveguides, on the other hand, which could have a play at shorter reach, probably are too lossy for routing to the front panel unless used at 850 nm wavelength for low-loss propagation and having e-o and o-e conversion at that wavelength within the board. And finally, in the roadmap in the distance future is included some type of 3D interconnects (to be determined) for stacked PICs and PIC to board-level waveguides. This prediction is consistent with existing 3D electrical IC stacking and corresponding through silicon via interconnects.

INTRA-MODULE AND INTRA-BOARD

The timing of optical intra-board (or intra-blade) and intra-module (or intra-package) products is expected to be delayed by incumbent electrical solutions. Electrical board-level VSR, XSR (extra-short reach), and USR (ultra-short reach) transmission line interconnects have standards and products at 100Gbps and working groups exploring 200Gbps channels. Electrical transmission across PCBs increasingly requires more power-hungry digital signal processing up to and including retimers but its cost is at least an order of magnitude lower than the optical alternative. Also, electrical fly-over cable within greater bandwidth capacity is a viable solution to further delay optical intra-board interconnects. The trigger for adopting optics within data center rack shelves will be driven by co-packaging fiber management problems. Once there are PCB-embedded optical waveguides and associated low-loss, low-cost, and scalable package-to-board optical coupling solutions available for switch-to-front-panel CPO interconnects, intra-board interconnects will begin to take-off. They will need to be cost competitive with electrical interconnects but not necessarily at cost parity. Full system-level cost-of-ownership costs that include electrical power consumption and heat management will likely hasten adoption once technical solutions are fully developed and are released into the market.

While the transition to optical interconnects from electrical interconnects has followed a continuous path of ever-shorter reach as the distance or reach times bandwidth product exceeds ~100Gb-m/s, that trend looks to be disturbed by the emergence of chiplet technology and heterogeneous integration. The demand of artificial intelligence for compute power necessitates closer/tighter packaging of compute chips like CPUs, GPUs, VPUs, Accelerators, FPGAs, ASICs
and memory (e.g., stacked DRAM). Multiple chips (e.g., up to 16) on a single multichip module (MCM) with high bandwidth, low latency interconnects addresses this need. Yet electrical interface standards such as High Bandwidth Memory (HBM), Advanced Interface Bus (AIB), and Universal Chiplet Interconnect Express (UCIe) struggle to provide the needed bandwidth (>2 Tbps), over the MCM substrate reach (>50 mm), at low latency (<20 ns) and power (<5 pJ/b). An approach that disrupts the gradual transition to optical interconnects could use of one of these standards from the chips to an optical transceiver engine followed by optical transmission across the MCM substrate and subsequent optical-electrical conversion and electrical transmission to the receive chip. The remaining challenge is low-loss, low-cost, scalable PIC transceiver to substrate coupling. This is an active R&D area and even the availability of an early commercial solution. [Lightmatter Passage brings Co-Packaged Optics and Silicon Photonics to the Chiplet Era (servethehome.com)]

The purple brick wall graphic lists events in the year of general commercial adoption not in technology development, a press release, conference demonstrator, or limited/general commercial availability. The exception is intra-module/board where the next several years is not expected to have commercial traction (hence the text prior to 2028 is grayed out). Events prior to that such as PIC-to-SiN interposers may be adopted for co-packaging but will be technology building blocks for the eventuality of optics displacing copper at these ever-shorter reaches. Using history as a guide, that should occur for intra-board, ~50 cm reach, once data rates exceed 200 Gbps, for a 100Gbps-m, distance times bandwidth product (aside from the MCM interconnects described previously). PIC flip chip will improve assembly manufacturing scale but needs significantly more process development to work out the optical I/O and come to fruition (example shown in Figure 4). While board-embedded waveguides could be beneficial with on-board connectors combined with short PIC-to-board jumpers, the combination of embedded waveguides and PIC flip-chip will likely be needed for intra-board optical interconnects to be viable. For optical interconnects to displace electrical transmission lines within the package, e.g., for CPU-to-stacked high bandwidth memory, optical engines will probably need another order of magnitude drop in cost. An intriguing alternative is micro-LEDs that today are widely used in display applications. [15]

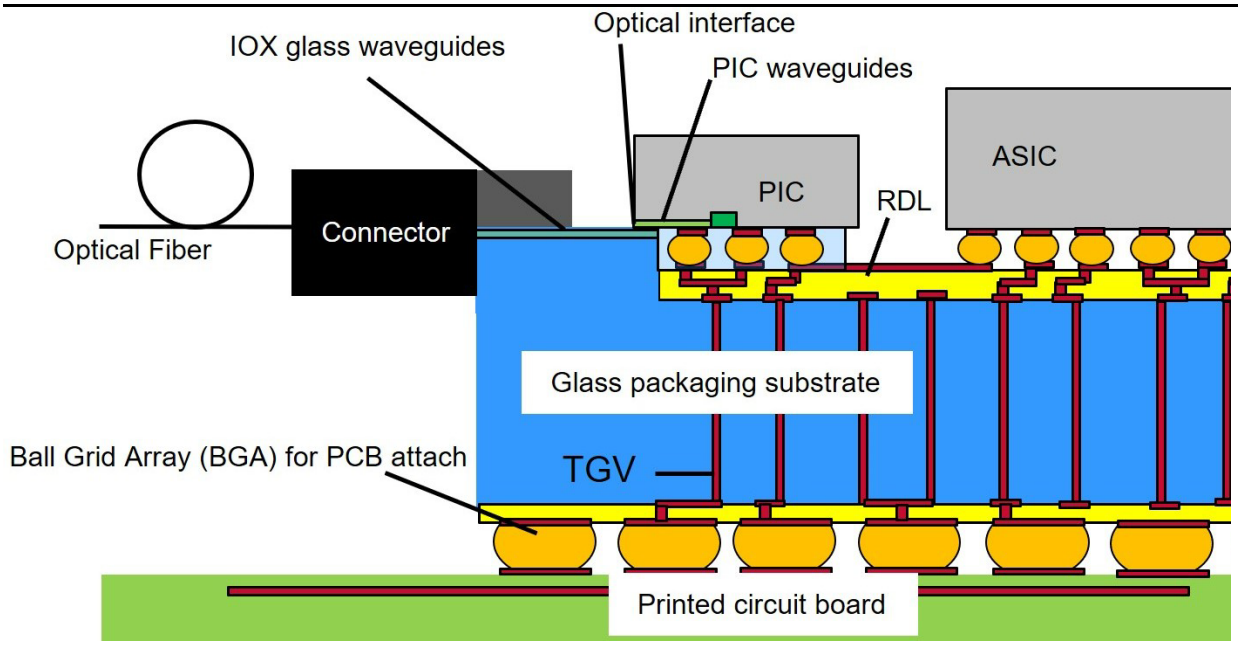

Figure 4. Example of a PIC flip chip onto glass package substrate with evanescent optical coupling interconnects to ion-exchange glass waveguides on the package surface and through glass via (TGV) electrical interconnects through the package to the PCB. [16]

KEY METRICS

COST ($/TERMINATION)

The most important metric is cost (meaning end customer price) since it heavily influences the buyer's decision. More specifically, what drives technology adoption is the total cost of ownership for the dominant player within the value chain. In the case of PIC interconnects, that is the hyperscale data center operator. Total cost of ownership is capital costs plus operating costs and the conversion costs for adoption if it is a new technology. New technology comes with unforeseen risks, so the adoption barrier, real and perceived, needs to be overcome by a strong value proposition.

As the PIC chip costs are driven down by foundry volume scaling, and testing costs are driven down by automation and sampling, the optical interconnect costs will come under increasing pressure. And for new PIC sensing and imaging applications, which tend to be more price sensitive, that cost pressure is even stronger. But only so much cost reduction comes from volume scaling of data center interconnects and adopting them to more price-sensitive PIC applications.

Cost is a notoriously hard metric to quantify given companies' reluctance to share this sensitive information. Additionally, any outside cost model requires a multitude of estimates and assumptions within the categories of bill of materials, labor costs, equipment costs, overhead costs, and profit. An attempt at quantifying the price of MPO and FAUs has been made for this roadmap; MPO and fiber array unit costs are obtained through quotes and the internet. However, the numbers are intended to be for guidance and directionally correct and not based on any specific company.

To begin, a current internet search finds that, in single quantities, 12 fiber MPO jumper prices range from $25 to $50 and 24 fiber MPO jumper prices range from $50 to $100 with both being $2 to $4/termination where termination is defined as the PIC-to external waveguide joint, in this case an optical fiber. Volume ordering should decrease the

unit price counter-balanced by the desire to buy at the upper end of the price range due to quality concerns from the high labor costs and risk of the jumper not working in a large installation. Hence the fiber-to-fiber termination price on the purple brick wall is assumed to start at $2/termination. This also assumes that installation costs are negligible with just a few seconds to clean and make the connection. Of course, customers will demand cost reduction over time, perhaps 5-10%/year. Additionally, higher fiber count interconnects, 2x12 and 2x16, with potentially similar assembly costs should drive a 2X reduction over the short-term time frame. Beyond the next several years, lower costs from such a well-established, incumbent technology will need to come from automation. Volumes will need to justify any capital expenditures for automation equipment and are expected to do so with the growth of artificial intelligence and the need for computationally intensive training algorithms. A purple brick wall event demarks manual-to-automated conversion due to the complication of converting intricate manual assembly to semi- and then fully automated assembly, particularly fiber loading into the MT ferrules, along with the unknown market timing of when volumes justify that conversion. Pricing on the purple brick wall assumes a 2X price reduction in 2026 due to 2x greater number of channels within the MPO interconnect at the same price and 10% per year reduction after that.

By contrast, fiber-to-chip termination is dominated by fibers loaded in glass V-grooves and actively aligned to the PIC for edge coupling. The bill of materials is expected to be greatly reduced over a full MPO connector, but the active alignment and particularly bonding/curing time adds significant time on an expensive, 6-axis alignment machine. A recent price quote for 1-9 quantities of glass V-groove fiber array units is ~$16/termination. This is much higher than a volume (~1M unit) price target estimated at ~$1/termination. On top of this, the attachment and bonding are estimated to be $1.74 if manufactured in the US and $1.34 if manufacturing in a low-labor country. Assumptions include 5-year depreciation on a 6-axis active alignment tool, 5-minute tact time, 2 working shifts per day, 76% machine up-time (accounting for maintenance, repair, plant shutdown, and daily overhead), 12 terminations per unit, and 100% cost-to-price markup (for profit, SG&A, R&D, taxes, utilities & space, and repair & warehousing). Under the assumption that production starts in the US and moves to a low-labor country halfway through the depreciation cycle, the current estimated attachment cost on the purple brick wall is $1.50. Hence the total unit plus attachment price starts at $2.5/termination.

Given that the incumbent V-groove technology is more than 30 years old, opportunity for cost reduction is limited. Production efficiency, minimizing CAPEX, and low-labor market production should keep a 5-10% price reduction going in the short term. While many new approaches are being explored longer term, none have emerged as the clear winner. Hence a purple brick wall event is anticipated which will get resolved through the hard work of yield improvements upon manufacturing scale-up and rigorous reliability testing of a new, winning technology. A large drop in price should occur, placed in the late 2020s, and after that subsequent 5-10% price reduction.

Over time, glass V-grooves may migrate to silicon V-grooves on the PIC or interposer (particularly for quantum, cryogenic applications), before moving to either evanescent coupling or expanded beam coupling (from the edge or surface). Alignment may evolve into vision-based and then passive alignment. Assembly will likely migrate to semi-automated, then automated, and ultimately full compatibility with the silicon electronic industry's OSATs (outsourced semiconductor assembly and testing) and traditional high-speed pick & place tools. This should keep the $/termination at an ever-decreasing level such that termination does not begin to dominate the packaged PIC cost and maintains its relative percentage of the total package cost.

COUPLING LOSS

Insertion loss (in dB) is the attenuation between two fixed points in an optical system and for this roadmap is the worst-case manufacturing distribution over the transceiver operating temperature and lifetime. For pluggables, the two points are between the PIC waveguide just before the input/output terminus and the pluggable-side connector interface (PIC-to-fiber coupling loss plus negligible fiber propagation loss from a short (~1 cm) fiber jumper); for co-packaging, the two points are between the PIC waveguide just before the input/output terminus and the front-panel side connector interface (currently PIC-to-fiber coupling loss plus a bit more fiber (<0.5 meter) but still negligible

fiber propagation loss). This definition ensures compliance with the IEEE 802.3 3 dB link budget for 500-meter reach, low-cost transceivers. SiP has been able to meet this standard and due to integration cost benefits has become the preferred transceiver.

Coupling loss (pluggable or co-packaging) has contributions from mode converters (if used), mode-field diameter mismatch, and alignment error between the PIC waveguide and the fiber. For edge coupling, inverse tapers of the silicon waveguide expand the mode field to nearly that of single mode fiber. It especially works well if a top layer of SiN is used as an intermediate coupling medium and/or a thicker 3 micron buried oxide layer is used. Active alignment with precision 6 axis tools minimized any error resulting in coupling loss about as low as possible for any existing technology. Any incremental improvement for actively aligned edge coupling, while not expected, could help the yield hit for transceiver output power and result in lower transceiver cost. Passive alignment will provide scalability but the gap in coupling loss performance between active and passive needs to be closed (e.g., ~0.7 dB [6]).

In the future co-packaging insertion loss will include any intra-rack connectors and any PIC-to-interposer coupling and propagation loss. This is not a problem for first-generation coupling technology being commercially introduced which uses actively aligned, fiber array flyovers to front panel connectors. Yet any new technology to address cost, manufacturing scale and deployment speed must maintain this low insertion loss or push for a new standard that adds link budget, e.g., new technologies like connectors or other passive alignment in the intermediate term and fiber- or interposer-to-board-level waveguides to address fiber management in the longer term.

But by 2045 if not before, the accessibility of very low-cost SiP PIC transceivers there will be a desire for inter- and intra-rack link budgets to become the same. Additional waveguide propagation loss will make this a significant challenge and more likely will be met with shorter intra-board or inter-module reach than longer inter-rack optical backplane reach. The end result is that fiber-based optical backplanes are much more likely than waveguide-based ones.

For quantum applications, coupling loss needs to be as low as possible (e.g., ≤ 0.25 dB). Every entangled photon quantum bit is valuable; loss of it requires retransmission, greatly impacting the data rate. An added requirement is operation near 4K so alignment at room temperature with a different fiber pitch-defining host material will have significant movement at the operating temperature due to coefficient of thermal expansion (CTE) mismatch and thermal walk-off. Hence silicon V-grooves with no thermal walk-off alignment errors are appropriate but a highly reliable solution to strain-relief of the fiber ribbon to the silicon V-groove is needed [17] and post-bond positional shift needs to be minimized. Another need is fiber with much better geometry specifications (cladding diameter and ovality, and core-clad concentricity). Fibers exist today with much better dimensional control than the standard single-mode fiber SMF-28 specification. Innovations in quantum PIC coupling, particularly tighter fiber geometry could be adopted for datacom as volume increases and price drops.

For some sensor applications, the sensing transducer is a disposable component separate from the source/detector PIC. Extremely low coupling cost and the need for rapid install/deinstall may afford a higher coupling loss specification.

DENSITY – LINEAR AND AREAL

Switch ASIC bandwidth has been consistently doubling every two years putting a real strain on optical transceivers to keep up with a combination of increased channel count, more wavelengths, and/or more modulation levels. Historically, increased channel count is initially the low-cost solution and then newer transceiver generations at the same data rate increase wavelengths and/or modulation levels. The new OSFP-XD MSA for 1.6 Tbps pluggables follows that tradition but does not yet constrain linear, shoreline (PIC edge) density. It will not increase until the PIC waveguide layout consumes the entire PIC edge. With an assumed 10 cm edge and 127-micron pitch to match that of the fiber allowing for more than 64 parallel channels, it is not expected for the foreseeable future. If needed, 80-micron fibers interleaved in V-grooves provide a 50% density improvement. Other solutions are interposer fan-outs

and 2D surface coupling. Finally, in the long term, coherent transceivers could be the disruptor for 500-meter DC interconnects that further mitigates increased density requirements due to much better spectral efficiency with the QAM modulation format. The recent commercial release of 100ZR QSFP28-DCO (digital coherent optic) SiP transceivers for up to 100 km is significant in getting to a compatible pluggable form factor with less than 5 W power dissipation including DSP but costs will also need to drop significantly to be competitive to direct detection SiP transceivers. [18] The summary is that density, while an important design consideration, is not likely to be a critical constraint so is not represented on the purple brick wall.

MANUFACTURING SCALE (TERMINATIONS/YEAR)

The relentless growth of the internet and the strong emergence of artificial intelligence and machine learning will be mirrored in the number of PIC-to-outside world input/output (I/O) terminations. As of 2023, there are 5.3 B internet users, three times the number of connected devices than the global population, and 14.7 B machine-to-machine connections. [19] Additionally, ~85% of data remains local East-West traffic within data centers and new AI-specific interconnect architectures will power PIC transceivers and resulting lane/channel terminations well into the future.

Terminations per year are based on market report volume and forecasting from 2018 to 2028 for the various GaAs, InP, and SiP integrated transceiver types (defined broadly as transceivers with at least two integrated functions). [20] Transceiver volume is converted to the number of terminations, (e.g., 8 terminations, Tx + Rx, for PSM4 and DR4). Then data was extrapolated to 2035 based on the individual material system's growth rate from the last year of the forecast, 2027 to 2028 of 42%, 8%, and 19% for SiP, InP, and GaAs material systems, respectively. Transceiver markets include wavelength division multiplexing (WDM), Ethernet, Fiber Channel, front haul, back haul, fiber-to-the-x (FTTx), and active optical cable/electro-optic modules-co-packaged optics (AOC-EOM-CPO) with Ethernet being by far the largest at 80% of terminations in 2023. The breakdown by market for SiP is shown in Figure 5. The terminations/year split by material systems in 2023 is 17%, 27%, and 56% out of a total 188M terminations for SiP, InP, and GaAs, respectively. However, SiP, InP, and GaAs growth rates over the next 5 years of 55%, 12%, and 16%, respectively, changes the mix to 48%, 15%, and 37% in 2028 (601M total terminations) and 79%, 4%, and 17% in 2035 (4.3B total terminations). This is due to the growth of SiP 500-meter 400G DR4 transceivers, followed by SiP 500-meter 800G DR8 and SiP 500-meter 1.6T DR8 transceivers. For example, 800G SiP terminations outpace 400G for all SiP transceiver types in 2027. SiP terminations dominate because they are the preferred transceiver type for 500-meter reach and 500-meter reach is the sweet spot for rack-to-rack DC interconnects.

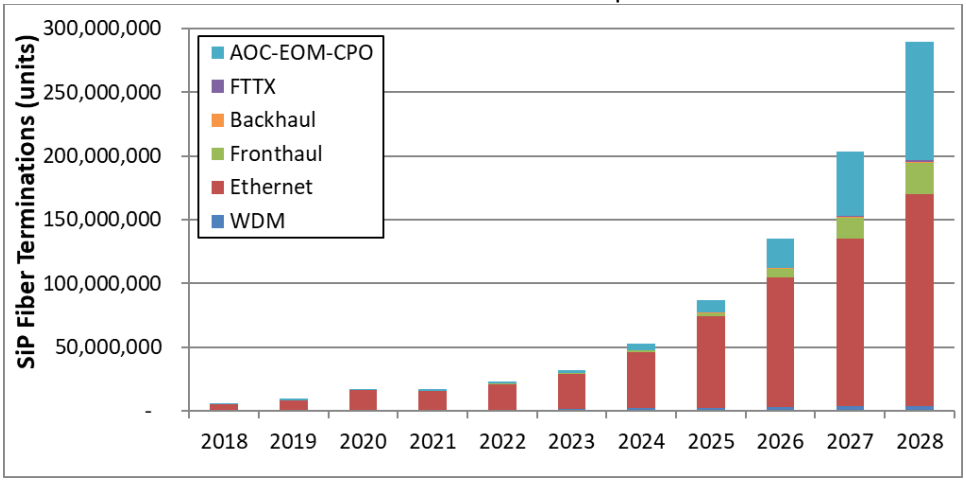

Figure 5. Fiber terminations by year for SiP transceivers by market segment

An extrapolation of the year 4 to year 5 material type growth rates to the 5-to-10-year range is uncertain but should be directionally correct. That uncertainty comes from the transition from pluggables to co-packaging, timing and product mix of 3.2T transceivers, market adoption of fiber-to-the-server, and an unclear path beyond 200G PAM4 optics. Historically, the first generation of a higher speed transceiver occurs with more parallel

lanes/channels followed by increases in wavelength, modulation speed, and/or modulation level. Increased parallelism is expected for 3.2T since both FR8 and PAM4 are built into the 1.6T data rate and more than 8 wavelength and 4 levels has significant technical challenges. Nevertheless, 3.2T transceivers with an expected 2X more terminations in their first generation are not reflected in the data even in the out years because there is no significant volume expected in the 5-year time horizon used for extrapolation. This reflects a conservative approach to the projections. Beyond 10 years, the number of terminations may dramatically increase as board- and module- level interconnects come into play. Perhaps a 10X bump at this lower level of interconnects can be expected but with the termination characteristics radically changing with flip-chip surface mount and high-speed pick & place technology. This 10x increase is reflected in the 2045-year terminations projection of the interconnect's purple brick wall.

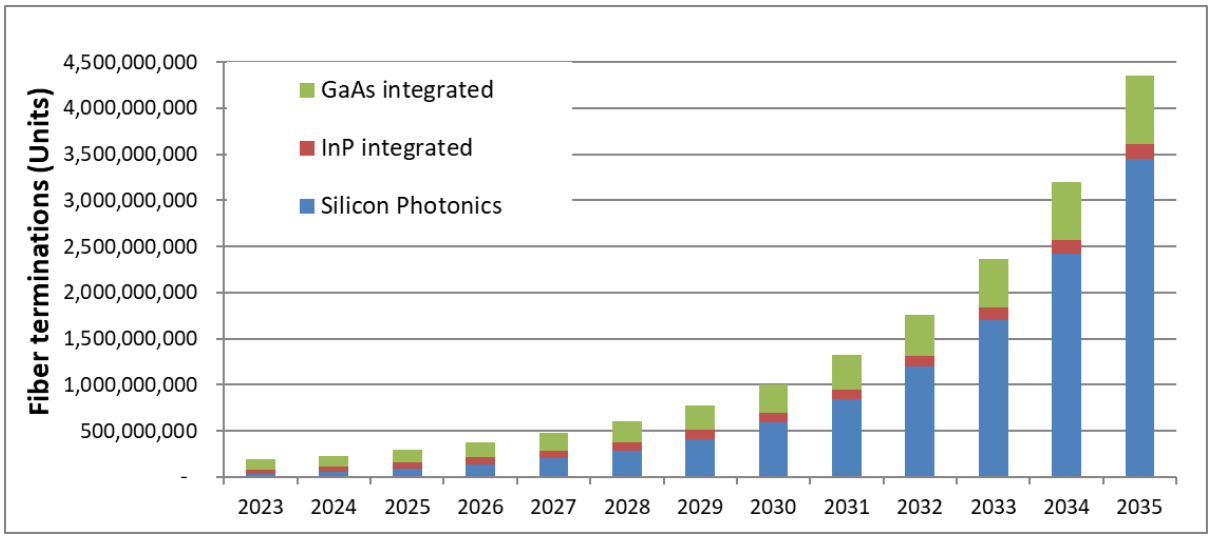

Figure 6. Fiber termination by year for the three material system transceivers

INSTALLATION SPEED

Installation speed becomes increasingly important with terminations per year growing from 188M today to 4.3B in 2035. Even for the unrealistic scenario of 24x7 installation, that translates to 6 termination/second across all manufacturing lines increasing to 138 termination/second. Clearly the assumed 5-minute current tact time for fiber array placement and bonding will struggle to scale - more tools in parallel is insufficient. Yet the recent announcements of co-packaged PIC-to-fiber connectors in development and accounting for a 2-year development cycle prior to commercialization, installation speed should have a dramatic drop in about 2025 to 10s of seconds, accounting for a connector cleaning step. The next jumps in installation will come first from semi-automated connector assembly and placement on the PIC with manual loading and unloading, then to full automation, and finally high-speed pick & place, flip-chip connections. While the interconnect's purple brick wall lists installation speed as the metric, a key installation challenge for fiber jumpers is fiber handling. A second further out is attachment consistent with pick & place tooling. Additionally, pick & place placement speed is highly dependent on required placement accuracy. Micron-level accuracy is needed for single mode insertion loss requirements. Technologies will need to be developed to provide alignment features or other mechanisms to translate the 5+ micron alignment accuracy of standard high-speed pick & place tools to micron level. Alternatively expanded

beam connectors relax the translational alignment tolerance to levels already available in standard pick & place tools.

OPTICAL INTERCONNECT TECHNOLOGY NEEDS

FIBERS/V-GROOVES

Fiber arrays in glass V-grooves (fiber array units, FAUs) have been commercialized for over 25 years and are the prevailing approach to multi-lane PIC I/O (either edge or grating coupling). They are highly reliable and provide low loss for matched mode field diameters on both sides of the interface (≤ 0.15 dB). Fibers, often in 12-fiber ribbons are stripped, cleaned and placed within V-grooves machined into a small glass block by a computer numerical control (CNC) machine. The assembly is completed by epoxying the glass lid/cover to the fiber array and the fiber ribbon to the glass block for strain relief, followed by dicing, and grinding/polishing the front surface, either flat or angled. The other end of the fibers can be bare pigtails or connectorized. Various fiber types including the addition of individual polarization maintaining fibers are available. The glass material choice allows for a coefficient of thermal expansion (CTE) closely matched to silicon to minimize thermal walk-off of the fiber pitch over the device operating temperature range. The standard fiber pitch is 127 or 250 microns with 127-micron pitch achieved in standard 250-micron ribbon by interleaving two fiber ribbons. While arrays up to 96 are available, accuracy degrades: ±0.5 microns for ≤ 16; ±1.0 microns for ≤ 48; ±1.5 microns for ≤ 96. Fiber cladding diameter and ovality, and core-clad concentricity contribute to positional uncertainty independent of array size whereas translational error of the CNC tool accumulates with distance. It is conceivable that in the future displacement measuring interferometry could be applied to improve tooling positional accuracy. Such measurement systems can detect movement down to 1 pm but other sources of error besides those of the fiber may still make it challenging to have the same positional accuracy independent of array size. Besides, the warp of the PIC chip from its thermal history and material layer CTE differences serves as another limit to accurate positioning of the fibers to PIC waveguides for large channel count.

Grinding of the grooves by CNC machine offers both high accuracy and flexibly in fiber placement in addition to the ability to create precision depth grooves as alignment features such as for dummy fibers or alignment pins or other elements. Another currently available innovation is lidless FAUs. When a lidless FAU is flipped upside down, the highly accurate fiber core center to cladding distance can be used for alignment to PIC waveguides. Another possibility, not yet commercially available but disclosed, is stacked fiber FAUs. [21] While it benefits from avoiding the cost of the V-groove (~$1), automation of assembly and the associated costs and ability to achieve high yield is to be determined. More likely in the near-term is the adoption of 80 mm diameter fiber and fiber ribbon for an improvement in linear I/O density. This becomes increasingly valuable for the shoreline limitations of co-packaging due to package substrate size constraints (from reliability issues of chip-to-package and package-to-board electrical interconnects). Longer term and significantly more challenging is PM ribbon (or SMF/PM composite ribbon), needed if externally laser sources continue to be an option for co-packaging – the challenge being alignment of the PM fiber axes of all fibers in the ribbon. It is on the other side of the purple brick wall. (See the reference for a detailed analysis of internal vs. external lasers. [18]) Less challenging is fiber ribbon able to withstand solder reflow. Materials exist (like PMMA) but they are not easily applied. New ink formulations would also be required. Beyond this, multi-core fiber for even greater linear density (~30 mm pitch) could be needed for co-packaging with ever increasing channel count. Even further out could be the space-division technology of mode-division multiplexing shown in the 2045-time horizon. Of course, a more straight-forward but more energy intensity approach is coherent transceivers and QAM modulation format to greatly reduce I/O channel count.

With all the advantages as the incumbent technology, the disadvantages of FAUs are further cost reduction, increasing the scale of manufacturing and decreasing installation speed. Fiber loading, epoxying, dicing, polishing, and testing are currently labor-intensive, manual processes so the move to manufacturing in low-labor cost countries has already occurred. Equipment is already fully depreciated, and automation requires CAPEX that raises the equipment cost as it increases production volume. And installation speed is limited by FAU-to-PIC epoxy cure time and has the burden of its individual depreciation time on an expensive 6-axis tool.

An alternative to separate, actively aligned glass V-grooves are silicon V-grooves integrated in the PIC. They eliminate the need for expensive alignment equipment by enabling passive alignment with a pick & place tool (with additional care for the pigtails or fiber spools.) Also, silicon V-grooves within a silicon PIC have no thermal walk-off which is critical to quantum applications. Finally, they are easily scalable to high fiber counts and is a standard offering on a 300 mm wafer process platform but take up valuable silicon real estate, are hard to rework, still require adhesive curing, and currently have loss just under 2 dB, higher than the purple brick wall target. [22]

FAUs will continue to be a dependable PIC interconnect technology for a long time to come but newer approaches will take over as cost reduction and scale-up of 400G and faster transceivers drive the market. One such approach is laser bonding rather than epoxy to avoid lengthy cure times. [23] Another approach is novel use of a silicon interposer with dry-etched U-grooves for dummy guide fibers to passively align the PIC and FAU which are flip-chip assembled. [24] It is a passive alignment concept for potentially high-speed assembly that has not yet been demonstrated.

GLASS WAVEGUIDES

For integrated optic waveguide fabrication, several material platforms can be used including silicon, indium phosphide, silicon nitride, lithium niobate, polymer, silica, and glass. [25] These differ greatly in their advantages and disadvantages and therefore are selected depending on the application. If low cost, high dimensional stability and both passive electrical and optical interconnects are required, glass substrates and waveguides made by femtosecond-laser writing (fs) and ion-exchange are very well suited [26] for both wafer and panel formats. For fs-laser writing, the energy of a single photon is lower than the bandgap and the refractive index modification is achieved by multiphoton processing. [27] This method can write full 3D waveguides. It can also create alignment features self-aligned with the waveguides with a laser damage and etch process. However, it is a sequential process with throughput limitations. Ion-exchange in molten salts is a well-known process in the glass industry (e.g., smartphone cover glass) for batch chemical strengthening of glass by exchanging potassium for sodium ions. In 1972, Izawa and Nakagome were the first to use the increase in refractive index to fabricate waveguides using ion-exchange. [28] Since then, the process has been further developed resulting in many publications showing the high potential of this technique and demonstrating a wide range of components. [29] The silver ion-exchange waveguide fabrication on large panel sizes (457 mm x 303 mm) was demonstrated on display glass with low propagation loss of 0.06 dB/cm. [30] The process has the potential for up-scaling to larger formats. The ion-exchange process parameters are dependent on glass diffusivity characteristics and need to be adapted to new glasses. The overall process flow for making glass waveguides by thermal ion-exchange are shown Figure 7: (a) glass panel, (b) thin film deposition, (c) resist dip-coating, (d) laser direct imaging and etching, (e) silver ion-exchange, (f) mask removal, (g) reverse ion-exchange with sodium ions and (h) laser singulation.

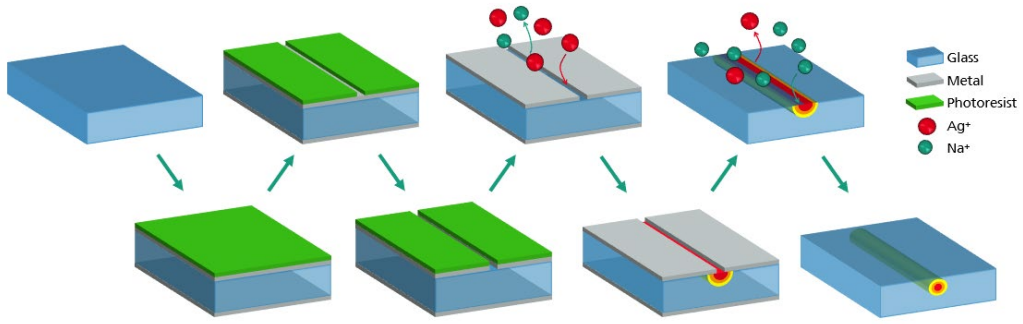

Figure 7: Process steps for ion-exchange on glass starting with glass panels (step a) which have been metallized (step b) to provide a diffusion mask and lithographic patterned in step d. The ion-exchange is done in steps e and g [26].

Because of the performance and reliability requirements, a glass with very low bulk absorption and waveguide longevity over the operating temperature is required for co-packaged optics. The waveguides are exposed to elevated temperatures over many years inside a package of a high-performance computing chip with significant heat dissipation of hundreds of watts. Of course, alkali-glass for ion-exchange from display applications could be considered but with trade-offs in waveguide longevity or optical loss performance. An optimized glass will be required with limited ion mobility at operation temperature range to avoid additional coupling and bend propagation losses due to change in waveguide characteristics by mode expansion over time. Improved glass compositions show promising results meeting the longevity requirements of 5 years at 110°C. [31]

For cutting glass waveguide circuits, high-speed laser singulation with optical-quality end-facets are required to achieve low-loss optical edge coupling to optical fibers. Laser cutting has significant advantages compared to other processes because they provide non-contact, near-zero kerf, no tool wear or the need for lubricants. [32] Laser singulation with controlled crack propagation by CO_2 laser creates optical end-facets along straight lines. Alternative, laser-based methods utilizing quasi non-diffracting beams (e.g., Bessel beams) can be combined with ultrafast lasers [33] for partial end-face perforation in combination with application of mechanical force or thermal expansion (e.g., with a CO_2 laser beam). The technique enables high accuracy, free-form cuts with reduced dependency on glass composition for cutting glass of all sizes. [34]

Two established methods of interfacing fiber and photonic integrated circuits are grating and edge couplers. [35] [36] Additionally, evanescent coupling in combination with intermediary waveguides and optionally fan-out to FAUs is a promising approach which is scalable to high optical port counts with pitches of 50 microns or less and chip assembly can be automated by pick and place machines. [37] Minimum coupling loss of 0.5 dB was demonstrated between ion-exchanged glass waveguides and flip-chip assembled SiN PICs. [31]

Besides optical integration, through glass vias (TGV) in combination with multi-layer electrical circuits on both sides of the glass core need to be manufactured with advanced panel level processing for high-density, high-bandwidth, parallel electrical off-package interconnects. High resolution laser lithography and deposition technologies (e.g., sputtering, spray-coating, plating) are required to achieve line/spacing of down to two microns as demonstrated for embedded multi-die interconnect bridge (EMIB). [38]

For Level 2 electronic packaging, glass boards can host optical and electrical signal distribution, optical and electrical interfaces to photonic integrated circuits, and connectors at the edge as interconnect to fiber cables (Figure 8. With heterogenous integration using glass substrates as printed circuit board substrates, either individually or embedded within in laminated layers of copper and an organic (e.g., fiberglass and thermoset resin), waveguide integration and advanced glass processing must be merged with electronic packaging and high-precision automated assembly. Given the ecosystem changes needed, this is considered a purple brick wall event.

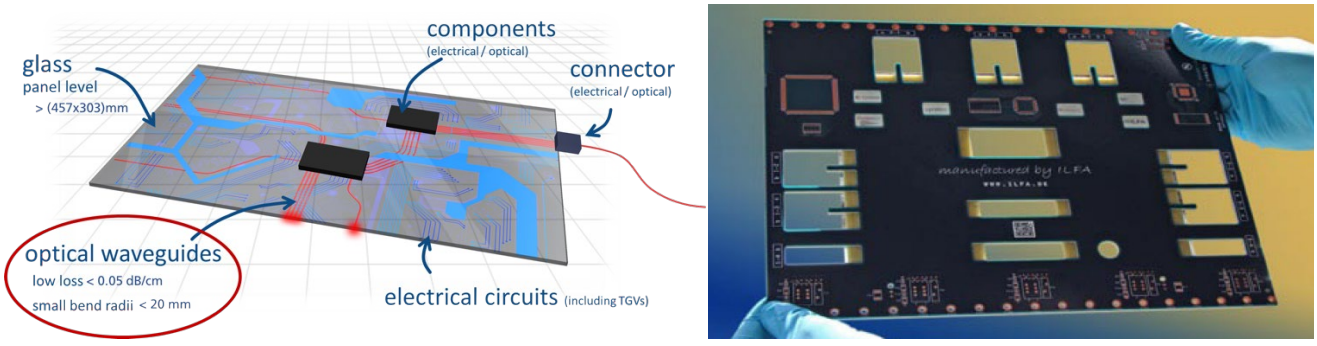

Figure 8: left: Concept drawing of the electro-optical glass layer of electro-optic circuit board (EOCB) containing optical waveguides (red), electrical wiring (blue), edge connector indicating some values for performance and market acceptance and (right) EOCB backplane with single glass waveguide panel and connector layout [39]

POLYMER WAVEGUIDES AND COUPLERS

In recent years, polymer photonics gained tremendous attention due to their manufacturing simplicity. Polymer optical lenses, for example, are employed in modern smartphone cameras due to the easier fabrication, robustness, and low costs. Other examples are total internal reflection [14] and spherical mirrors for expanded beam optics, evanescent couplers, [40] [41] alignment grooves or other features for fibers in FAUs, and FAU-to-PIC passive alignment. Moreover, the use of planar polymer waveguides offers a promising route for on-board optics and intra-module optical interconnects. Fabrication techniques for polymer waveguides include open-air polymer solidification from a glass micropipette, [42] laser-scanned two-photon polymerization having 3D capability, [43] and micro dispensing a core monomer within a not yet cured cladding monomer resin. [44] However, like laser-written glass waveguides, these are a sequential process with throughput limitations. The 3D capability of two-photon photopolymer lithography, also known as photonic wirebonding, makes it a useful prototyping and small volume manufacturing tool. While the writing speed had not yet been optimized, it was recently reported as 3 minutes for a two-photon polymerized photonic wire bond. [45] Additional, the insertion loss from heterogeneously packaged laser chip-to-PIC was shown to be as low as 0.4 dB, [45] the PIC-to-fiber insertion loss of 1.7 dB is higher than the 1 dB target. [46] In contrast, polymer waveguides can also be manufactured using panel-size photolithography. Making use of the tooling and standards from the printed circuit board industry provides an easier, more common process flow, and cuts down the production costs. [47]

A benefit of using polymer waveguides is the fact that they exhibit roughly the same dimensions and optical properties (e.g., refractive index contrast and mode-field diameter) as standard optical glass fibers, which makes combining with commonly used pigtailed laser sources or detector modules straightforward. Similar to optical fibers, polymer waveguides can either be used in multimode (larger core dimension, typically 30 – 500 µm) or single mode (2-5 µm core size) operation. In both cases, they are compatible with the according fiber type. Recent advances in dry-film photopolymers have enabled low-index contrast, single-mode polymer waveguides that can be directly written using maskless lithography to connect PICs to other PICs, optical printed circuit boards [40] or fiber arrays. Further work is required to improve the thermal stability of this waveguide material system so they can withstand solder reflow processes.

Due to the compatibility of polymer waveguides with electrical PCB technology, they offer a promising route towards on-board optics systems, within which light is routed intra-rack at least at the module/package and daughter card level.

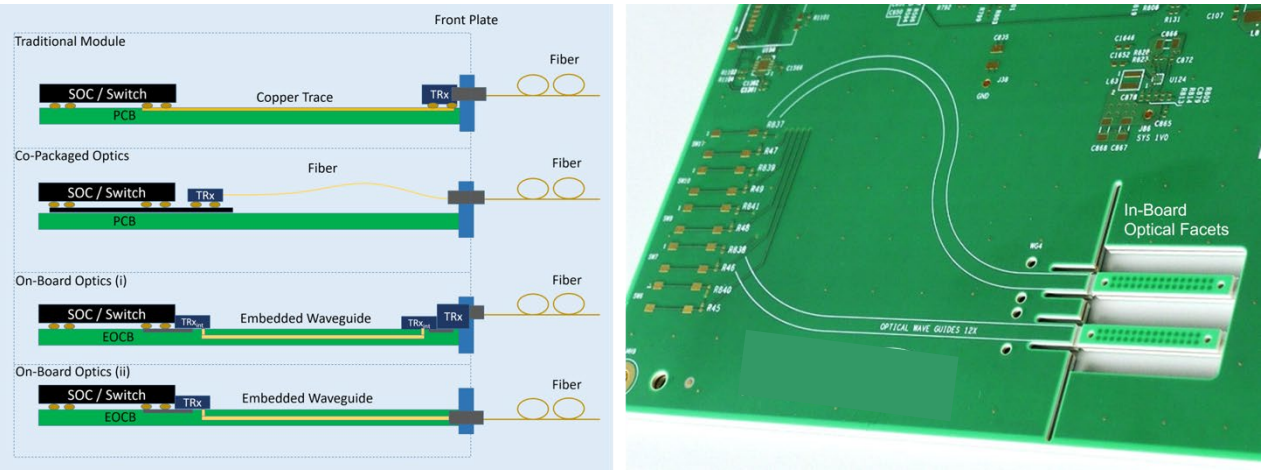

Figure 9: Comparison of various CPO and on-board optics concepts (left). Image adapted from [48] Example of a PCB-integrated polymer waveguide board (right).

A lot of effort in recent years has gone into maturing polymer waveguide technology with respect to their environmental stability. These days, polymer-based electro-optical circuit boards (EOCBs) can undergo reflow soldering (260°C) and have successfully passed typical Telcordia tests (85°C/85% relative humidity). While the

propagation loss at 850 nm is already very good (≤ 0.05 dB/cm), further developments in reducing loss at 1310/1550 nm to below 0.1 dB/cm is on the roadmap. Achieving these values will require the use of polymers with greatly reduced hydrogen content, as it is the vibrational overtones of X-H (where X is C, O, N, etc.) that cause the losses in the 1310/1550nm region to be much higher than those in the 850nm region. Recently reported attenuation for multimode waveguides is 0.13 dB/cm [49] to 0.42 dB/cm [50]) at 1310 nm and 0.22 dB/cm [49] to 0.8 dB/cm [51]) at 1550 nm wavelength of operation. As a demonstration of feasibility, completely fluorinated, buried-channel, polymer waveguides using conventional lithography showed 0.05 dB/cm loss at 1310 nm and 0.07 dB/cm loss at 1550 nm [59]. While it replaces the highly absorbing C-H bond with very low absorption loss C-F bonds, it is expensive material technology yet to be commercialized.

Multimode coupling to PICs and other transceiver components has been demonstrated with further standardizations expected through large industry consortia, which are being formed. One concept considers very-short reach, client optics across the board at 850 nm wavelength and then o-e-o conversion (optical-electrical-optical) to short reach client optics at 1310 nm. [52] Cost of the very-short reach, client optics will have to be very low for this to be an attractive alternative. Aside from that, perhaps extra-short reach (XSR, up to 50 mm) across the package from embedded optics on the switch ASIC or ultra-short reach (USR, up to 10 mm) across the package from compute-to-memory (CPU/GPU to high-bandwidth memory (HBM) is the best application reach that will embrace the existing infrastructure planar processing compatibility of polymer waveguides.

EXPANDED BEAM

Recently, an expanded beam connector (EBC) for optical fiber-to-fiber arrays became commercially available to address two key issues in structured cable, data center interconnects: dust and higher channel count. [13] A lens array at the front surface of the connector is used to expand the beam from the fiber array. Any dust on that surface has less optical loss because there is less overlap with the optical beam. Dust/debris can also damage fiber end faces when mated so EBCs reduce or prevent costly cable cleaning and inspection procedures in the field. EBCs enable higher channel count than MPO connectors which rely on physical and optical contact between fibers slightly protruding from adjacent molded MT ferrules. The Hertzian force for contact scales with the number of fibers limiting what is commercially available today to 2x16 for single mode fiber ferrules whereas EBCs with their low mating force are available up to 144 fibers. Another advantage is relaxed lateral alignment tolerance in the collimated, expanded beam region between the lens arrays of each connector. Yet this has not translated to lower loss in commercially available products due to additional optical surfaces and alignment degrees of freedom (EBC of 0.7 dB vs. MPO of 0.25 dB). In addition, there is more parts and alignment assembly requirements for the lens arrays, need for AR coatings, and difficulty reducing the pitch below that of the fiber at 250 microns due to the diameter of the lens array and desire for large beams for more relaxed lateral alignment tolerance.

However, the benefit of reliable assembly in non-clean environments typical of OSATs and detachability or mate/demate functionality are attractive for co-packaging. Far better alignment tolerance (greater than ±20µm for 1dB additional loss) is the key enabler for a detachable solution for optical chip connectivity. The connection can be tested then detached before package-to-board attach, typically through a solder reflow process, to avoid exposing the fibers and other components to high temperature (~250°C). EBCs also help with fiber deployment within the shelf compared to permanently attached pigtails. Further, any thermal alignment walk-off between the PIC and fiber array due coefficient of thermal expansion mismatch over the wide operating temperature range is mitigated by the wide, collimated beam waist between the PIC and fibers.

Molded multi-mode fiber (MMF) array lenses, ferrules, and connectors exist for board-mounted parallel optics and an SMF version has been investigated but the tolerance and CTE are extremely challenging. So, the two main approaches for beam expansion for optical connectivity are a micro-assembly of discrete micro-lenses or lens arrays such as in a glass molding process, [53] and a new approach using photolithographic alignment of beam-expanding lenses. An example of the former is a two-component, separable, expanded-beam connector that is based on a novel ceramic manufacturing process. A fiber array is permanently bonded to a ceramic ferrule that mechanically self-aligns into a ceramic cradle that is permanently bonded to the PIC. The ferrule and cradle are separable, capable of undergoing several mating cycles, and can either be held together with a latch or

permanently bonded together at a final assembly stage. [13] For an off-the-shelf lens array glued to MT ferrule, a ~3 dB coupling loss is reported so work to reduce loss is needed. The advantage of the latter approach, photolithographic alignment of beam-expanding lenses, is that it facilitates coupling to polymer waveguides on the surface of the PIC paving the way for future chip-to-chip, on-board optical connectivity. Both approaches proport to enabling wafer-scale assembly and testing and having paths to high volume production. With multiple companies developing EBO solutions, expect a 1 to 3-year development cycle for first products. For adoption, challenges to overcome besides those described for fiber-to-fiber connectors are size and attachment processes for a socket or connector housing surrounding, or on top of, the PIC. Also lens array pitch match that of the fiber unless the mirrors or lenses have magnification and a waveguide pitch of 250 microns consumes valuable PIC real estate.

Further generations will likely improve loss, cost, and channel count, e.g., automated termination of 32 fiber arrays with maximum 0.5 dB loss. Further out could be industry adoption of evanescent polymer couplers. In that case, a turning-mirror-to-expanded-beam configuration on the PIC surface would provide wide bandwidth and low cost if the mirrors can seamlessly be integrated into silicon photonic foundry process flow and socket or connector housing attachment can be automated.

ALIGNMENT & ATTACH

Surface coupling can be broken into grating, evanescent, and turning mirror-based; edge coupling can be broken into proximity and lens-based. Much has been written about the advantages and disadvantages of each. [35] [37] A detailed review is beyond the scope of this roadmap where the focus will be on alignment tolerances. Proximity (fiber array-to-PIC) is currently widely used for SiP with inverse tapered silicon waveguides in the PIC to match its mode field diameter to single mode fiber. The 1 dB loss lateral misalignment tolerance is < ±2 microns. Expanded beam approaches have recently become commercially available that use a lens to expand the beam and a lens to refocus the beam greatly improves lateral assembly tolerances to better than ±20 µm for 1 dB lateral misalignment with a controllable decrease in angular misalignment tolerance. Future evanescent coupling approaches can have lateral misalignment tolerance ~ 2 microns, vertical misalignment tolerance between 0.5 and 3 microns depending on the design, and longitudinal misalignment of 100s of microns.

The future of alignment can be divided into two possible paths. The first are coupling approaches having misalignment tolerance of a few microns with the following time sequence: 1) currently active alignment, i.e., optical power peaking and a six-axis alignment tool, moving to alignment elements or features on the PIC, interposer, and/or fiber array (in one form a complete optical connector directly attached to the PIC); 2) then improvements in speed though partial or full automation; 3) finally leading to pick & place compatible high-speed placement. The second are coupling approaches having misalignment tolerance of a 10s of microns (i.e., expanded beam) with the following time sequence: 1) work out details of beam expansion on the PIC side that are transparent to semiconductor processing and are low cost; 2) also move to improvements in speed through partial or full automation; 3) finally lead to pick & place compatible high-speed placement. Semi- or partial- automation could be used for manual material loading/unloading followed by automation of the alignment process. Cost, manufacturing scale, and deployment speed requirements will define the timing of these technology transitions with estimates provided on the purple brick wall roadmap.

Closely associated with alignment is attachment of the optical interconnect assembly. The analysis in the cost section highlights the high cost of the existing FAU approach even though it is highly reliable and well-established. There is little that can be done to drive the cost down - industry needs to move away from adhesive attachment for increased volume scale up and ultimately OSAT transparent production. Semi-automation is the first step where alignment processes are likely to be worked out first before any automated interconnect latching/mating. A more significant challenge is developing processes to test a known good-attach prior to committing the optical engine to the co-package. [9]

Since co-packaging uses 16 optical engines in a typical current configuration, OSAT processes such as solder reflow and non-clean-room assembly are a critical long-term goal. Only then will integrated photonic assembly fulfill the

promise of leveraging the silicon industry across all manufacturing processes (wafer, packaging, assembly, and test). The path forward is innovative latching mechanisms analogous to fiber-to-fiber connectors that meet loss and cost metrics and eventually pick & place machine compatible to meet manufacturing scale and deployment speed metrics.

ARCHITECTURE

The single biggest impact on data center optical interconnects over the last 10 years has been the transition from MMF to SMF. Hyperscale operators have seen the cost benefits of 500-meter, single-mode transmission using SiP transceivers that have dropped below the $1/Gb cost target combined with the future-proofing benefits of single mode structured cable. While enterprise data centers and the short reach (≤ 100 meter) of high-performance computer interconnects have preferred multi-mode fiber, cloud services and new AI architectures along with hyperscale DCs will insure a higher CAGR for single mode fiber transmission systems

Over the next 5 to 10 years, the biggest impact will be the transition from top-of-rack Ethernet switch ASICs and direct attach copper to servers in the rack to middle-of-the-row (MOR) or end-of-row (EOR) switches and fiber-to-the-server (FTTS). Fiber-to-the server deployment will be delayed until 800 Gbps server links by the emergence of active electrical cable which has a 7 meter reach for 8x50G PAM4 and 30-50% lower costs than active optical cable. [3] Cost effectiveness had previously been achieved with SiP transceivers in active optical cables (AOCs). AOCs save termination costs and have higher yield as a dedicated link but suffer from installation speed which has hindered wider deployment. The dropping price of SiP transceivers and relentless doubling of switch bandwidth driving ever-increasing server-to-switch bandwidth demand will contribute to FTTS adoption. Specifically, full utilization of the highest available bandwidth switches provides the lowest cost architecture. The latest switch generation fabricated on the smallest available semiconductor technology node always has lower power consumption and lower network diameter* (and hence lower latency) avoiding a complex and costly CLOS switching architecture based on lower bandwidth switches. [3] (* The diameter of a network is the maximum length of any shortest path between an input and an output.) Added benefits of fewer switches are less network management and less rack space for the switches. A generalized cost model shows that “as soon as optical links are twice to thrice as expensive as electrical links, FTTS merits serious consideration.” [54] FTTS is a big deal - it could open a 10x increase in optical interconnect demand not shown in the roadmap projections.

Another important architectural driver for demand of optical interconnects is in machine learning and artificial intelligence networks. They require the full parallelism of an all-to-all mesh topology between servers for optimum performance. [10] This forces a lower bandwidth granularity to match that of server’s I/O and hence increased number of overall interconnects rather than bundling of higher bandwidth optical streams on a single fiber.

Purple brick wall challenges start with fiber-to-the server which will need a significant drop in price to compete with direct-attach copper. They continue with co-package fiber management issues [11] that could drive glass package substrates with waveguide fan-outs and finally embedded PCB waveguides for front panel routing.

CONTRIBUTORS FOR TRANSCEIVERS

2023 Report

Michael Lebby, Lightwave Logic – Co-chair
Atikem Haile-Mariam, Lightwave Logic – Co-chair
Alan Evans, Evans Consulting and Corning (retired), USA – Co-chair
Graham Berry, Huawei Technologies R&D UK Ltd.
Richard Pitwon, Resolute Photonics
Michael Robertson, Huawei Technologies R&D UK Ltd.
Amit Agrawal, Cisco
Keith Newman, Hewlett-Packard
Russell H. Lewis, HP
Sherwin Kahn, Alcatel-Lucent
Debendra Malik, Intel
Paul van Dijk, LioniX International
Alberto Bianchi, Ericsson
Fabio Cavaliere, Ericsson
Peter Maat, AstronBoudewijn Docter
Henk Bulthuis, Broadex Technologies UK Ltd.
Mike Wale, UCL & Technical University of Eindhoven
Nestor Coronado, Value loops
Maarten van Egmond, Genexis
Timo Aalto, VTT
Shangjian Zhang, University of electronic science and technology of China
Erwin Vergeest
Stephane Barnebe, CEA LETI
Zhaohan Liu, VTEC
Hans Hilgenkamp, University of Twente
Edwin de Jong, NIBC
Eric Mounier, Yolé developpement
Hann-Huei Tsai, NARLabs
Jeroen Duis, PHIX
Sylwester Latkowski, Technical University of Eindhoven & PITC
Tobias Gnausch, Jenoptic
Martin Korst, Fraunhofer
Jinzhong Yu, Chinese academy of Science
Rogier Verberk, TNO
Andre Richter, VPI Photonics
Bert-Jan Offrein, IBM
Ruth Houbertz
Martin Moehrle, Fraunhofer Heinrich-Hertz Institute
Zhiyong Li, Chinese academy of sciences
Huub Ambrosius, Technical University of Eindhoven
Makoto Okano, AIS
Arnan Mitchell, RMIT University
Koji Yamada, AIST
David Marpaung, University of Twente
Jan Mink, VTEC
Willem Vos, University of Twente
Harald Haas, The University of Edinburgh

CONTRIBUTORS FOR OPTICAL INTERCONNECTS

2023 Report

Alan Evans, Evans Consulting and Corning (retired), USA (TWG Lead), evans5836@gmail.com – Chair
Michael Lebby, Lightwave Logic
Atikem Haile-Mariam, Lightwave Logic
Lars Brusberg, Corning, USA
Jaime Cardenas, University of Rochester, USA
Nikaus Flöry, Vario-Optics, Switzerland
Atikem Haile, Lightwave Logic, USA
Vladimir Kozlov, LightCounting, USA
John Lively, LightCounting, USA
Jurgen Michel, MIT, USA
Amit Miron, Teramount, Israel
Dick Otte, Promex, USA
Henning Schröder, Fraunhofer IZM, Germany
Julian Schwietering, Fraunhofer IZM, Germany
David Stegall, 3M, USA
IPSR-I 2020 Interconnects Roadmap

APPENDIX

MARKET POTENTIAL

Growth in datacenters and telecom will support the ever-increasing internet traffic: historically 50% CAGR from cloud computing, video streaming, mobile data traffic and the general expectation by consumers to have higher connectivity options and now additional growth from Generative-AI, "G-AI", applications. This growth will also require limited increases in fronthaul and backhaul network capacity. As datacenters evolve from concentrated core megacenter locations to clusters of small, distributed datacenters, there will be the need for high-capacity links between these edge computing data centers.

Integrated photonics is of great interest as an enabling technology for data center resource disaggregation. Disaggregation promises improved efficiency through resource utilization and modular upgradability for the data center but is limited by the cost and performance of the links. Additionally, latency requirements impose hard limits on distances over which certain resources can be disaggregated.

The trends identified by Bloomberg and other analysts noting the explosive growth of G-AI, as well as trends noted in the IPSR roadmap have accelerated over the past decade. While the IPSR roadmap primarily focuses on the number of devices connected to the data center that could lead to the adoption of billions of IoT devices for appliances as diverse as HVAC, medical instruments and factory automation, Bloomberg and others focus on the need to accommodate increasingly complex photonic devices that will be widely deployed in either edge or core datacenters.

The massive amount of data from these fast-growing elements is unstructured and the demand for analytics and fast movement of data is also accelerating. In addition, the importance of security has been repeatedly highlighted with multiple high-profile situations gathering world-wide attention.

Data centers of today are seeing a transformation from independent computing, storage and networking systems to integrated systems. As noted above, this change is driven by the twin demands of G-AI focused devices and mass quantity IoT devices. This more integrated system will increase the number of devices that are interconnected thereby increasing the amount of data that is transmitted, sorted, analyzed and distributed.

Internet Protocol (IP) traffic is the key metric to show activity in telecommunications, both trends in end-user connectivity as well as data communications within datacenters and high-performance computers (see Figure 10 (source: Cisco VNI). In Figure 11, global IP traffic was forecasted to approach 400 exabytes (10^{18} bytes) per month by 2022. The graph shows a strong growth of 26% CAGR (2017-2022). Most of the traffic is being driven by internet video and is very fast approaching the metric of Zettabytes per month, which is 10^{21} bytes of data per month with no signs of slowing down. Some estimates are discussing the further metric of Yotta which is 10^{24} bytes of data over the next decade, which is also expected to be driven by internet video but also by emerging web/data, gaming, and Internet of Things (IoT) devices.

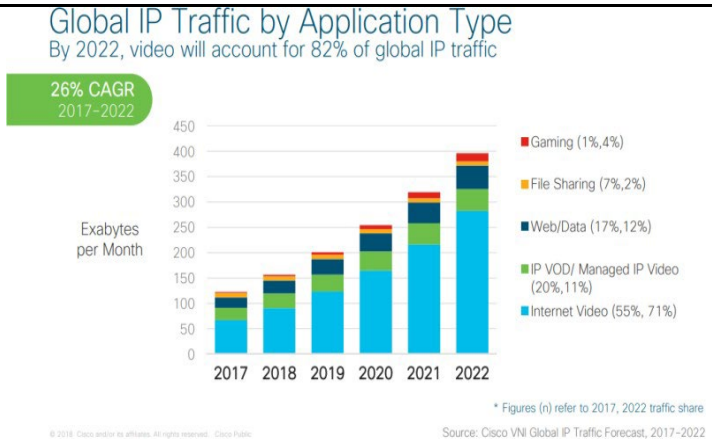

Figure 10. Rapid increase in data traffic measured in Exabytes (Source: Cisco VNI 2020).

As noted above, it is also useful to track end user connectivity trends as user expectations will impact IP traffic growth patterns.

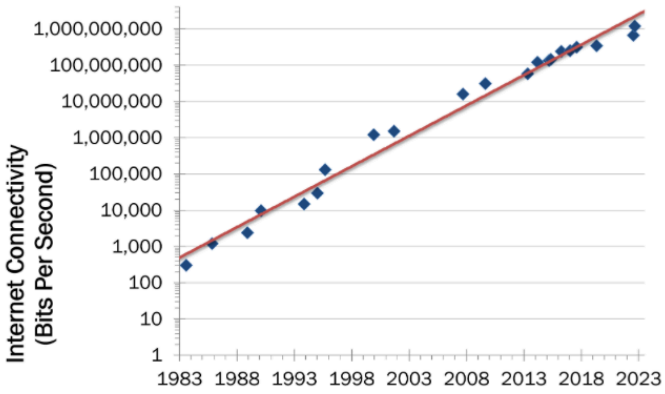

Figure 11: "Nielsen's Law" Internet Connectivity Grows at 50%/yr (Source: Nielsen Norman Group 1/23/23)

However, "traditional" Internet and user traffic data and forecasts are quickly becoming overshadowed by intra and inter data center traffic patterns which are driven by the rapid adoption of G-AI. While publicly available data is limited, secondary evidence of this growth can also be seen by the growth in computing power and high speed = largely photonic enabled component - link sales

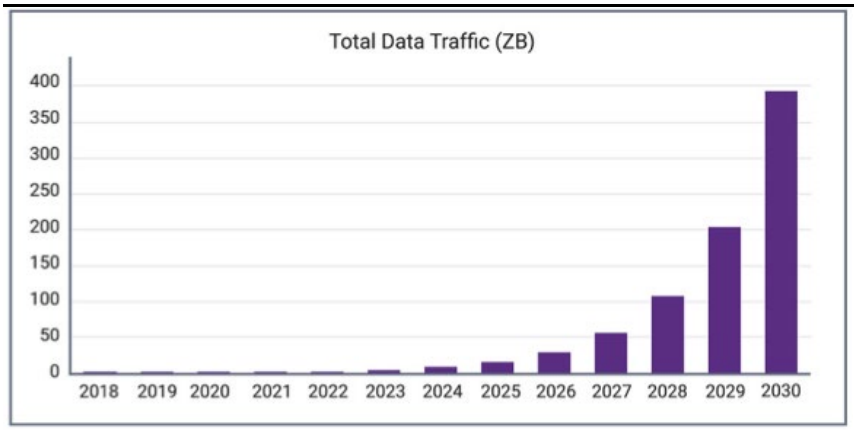

Figure 12: Total Data (center) traffic. Impact of AI on Electronics and Semiconductor Industries", IBS, April 2020.

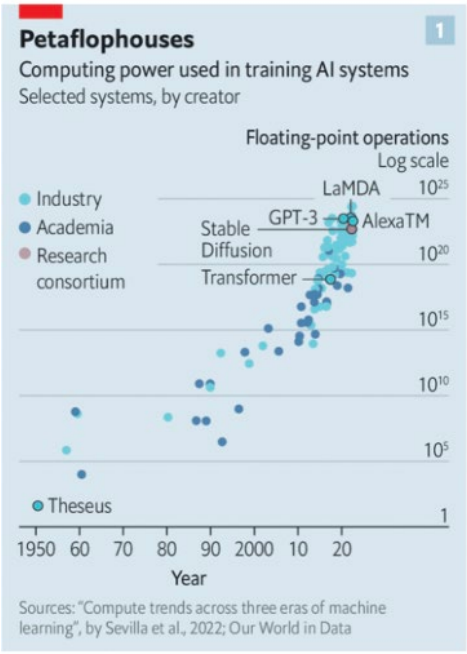

Figure 13: "Who Will Dominate the Generative AI Search Market?", Wally Boston, February 21, 2023

To accommodate the strong growths in data traffic, the fiber-optic infrastructure that allows data to be communicated between network nodes and within datacenters, has to be upgraded. Today, fiber-optic networks are a combination of long, medium and short optical interconnects that range from approximately 3 m to over 1000 km depending on application in the optical network. Photonic components are used to build the fiber-optic infrastructure; they comprise devices like laser diodes, photodetectors, multipliers, modulators and transceivers.

These are known as discrete components, while a mix of these components that are integrated or connected on a single substrate (such as silicon, InP or GaAs) are known as PICs. The summary assessment of the entire photonics-enabled marketplace in 2022 is estimated to be over $2T[5].

5 SPIE Optics and Photonics Industry Report, 2022

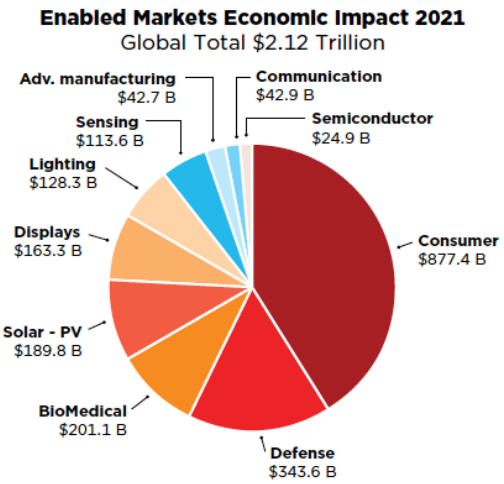

Figure 14. The summary photonics components market forecast for the next decade (Source: SPIE 2022)

In Figure 15 the forecast of Ethernet-datacom transceivers is shown by data rate for the next decade. As data rates keep increasing on a yearly basis, it can be seen that the main drivers will be 800G applications. By the end of the decade, revenues from 1600G are expected to grow quickly.

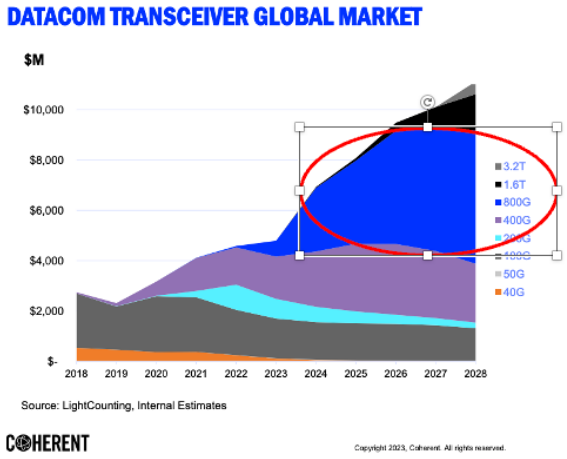

Figure 15. Forecast of -datacom transceivers by data rate to 2028 (Source: LightCounting/Coherent 2023).

It is expected over the next decade that transceivers will be an excellent platform for the accelerating trends of PICs in both telecom and datacom applications; it is well known that transceiver trends over the past decade have been towards smaller devices i.e. smaller transceiver formats and footprints, with higher densities of photonics components.

APPLICATIONS

As compared to the overview above, a more specific technical sub-division of the market for photonic components can be given. In this section a more detailed overview is given of the following applications:

- Optical Interconnects
- CATV and Radio
- RF Analog applications
- Active optical cable (AOC)
- Fiber to the X (FTTX) (X=curb, building, home, cabinet etc.)
- 5G front and back-haul
- Optical wireless (Li-Fi)
- Undersea and long-haul systems
- Metro and optical transport
- Datacenters and High-performance computing
- Optical transceivers
- Situational analysis for optical interconnects

At present, the optical interconnects are implemented by a number of low channel count optical pluggable modules or board mounted optical engines. Next generation Digital ASIC for 5G and data center systems will have a much-increased processing capacity reaching tens of terabit/sec in a single unit. To achieve that the signal data rate increases from the currently used 25 Gbps to 400 Gbps and beyond, the bandwidth density and the energy efficiency must also increase.

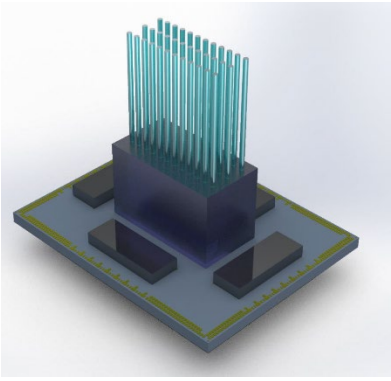

Figure 16. 16 x 112 Gbps full-duplex 4xDR4 HDI/O chiplet with in-house Silicon Photonics, Equalizing Driver, TIA (Nubis Communications 2023)

In such conditions, a new ASIC optical interconnect technology is needed based on co-packaging digital ASICs with multi-channel optical transceivers in the same multi-chip substrate to form an optical multi-chip module (OMCM). OMCMs may well be classed as co-packaged units that are part of a photonics integrated circuit (PIC) platform over the next decade.

These OMCM modules will include a high processing capacity electronic ASIC and a high scale integration silicon photonics interposer as part of a photonics integrated circuit (PIC) platform and will comprise of a bank of high-

speed optical transceivers. The analog electronic integrated chip (EIC) to drive the optical transceivers is 3D integrated on top of the photonic interposer (see Figure 4).

This can be particularly beneficial in the situation where one encounters many parallel connections, like in the OBO-switch in datacenter where a single Tomahawk BCM56980 is supported to deliver 12.8TBs switching capacity. See https://www.onboardoptics.org/cobo-switch . Or for DWDM-coherent systems when large number of optical lanes are available to be multiplexed on a single fiber. For these and other applications on board optics may emerge as a way of reducing footprint and reducing optical power because you locate the high-speed electronics as close to the actives (lasers and detectors) as possible, avoiding all kinds of signal degradation that require further electronics for recovering the integrity of the signal.

CATV AND RADIO

Market Potential

The revenues development for CATV and Radio TxRx devices is expected to be in the magnitude of $2B by the end of the decade (2030). The CATV market segment is composed of several fiber optic components utilized for CATV networks. These can be categorized as CATV amplifiers, CATV laser diodes and receiver photodetectors, and CATV passive photonic components. For Radio, there are 4 main data rates speeds for the photonic components and those are 3 Gbps, 6 Gbps, 10 Gbps, and 25 Gbps. These photonic parts are segmented in a number of distance categories that include <40 km, <10/15 km, <2 km, and <300 m. It is expected that PIC technologies are slower to penetrate this market as it is price sensitive compared to traditional data-communications and telecommunication market segments.

RF ANALOG APPLICATIONS

MicroWave Photonics (MWP) is an inter-disciplinary field that bridges photonics and microwave electronics. It aims to apply photonic solutions to microwave applications to achieve superior performance in terms of frequency agility, bandwidth, insertion loss, dynamic range, efficiency, size, weight, power and EMI robustness. MWP serves as an enabling technology in a wide variety of applications such as GHz and THz signal generation and distribution, high-speed wireless communication networks and radar systems.

ACTIVE OPTICAL CABLE (AOC)

An active optical cable is a short distance interconnect that connects racks and switches typically inside a datacenter or in a high-performance computing environment. The fiber-optic cable is multimode and is limited to a distance of a few 100s of meters depending on the specifications and optical link budget designed. The key for AOC cables is that the transceiver is designed to be part of the connector so that the optics are completely hidden from the user. In a typical AOC cable, each end of the interconnect cable has embedded transceivers where the output of each end is not a fiber optic connector but electrical connectors. The design of AOC is meant to bring down the $/Gbps metric so that markets such as datacenters, high-performance computing, PC and consumer markets can be reached at competitive pricing. AOCs are designed to fit into existing network infrastructures by interfacing to systems via a wide range of standard MSA connectors including CXP and QSFP+. The electrically connectorized cable ends are electrically compliant with InfiniBand, Fibre Channel, SAS 3.0 and 2.1, and other consumer protocol applications.

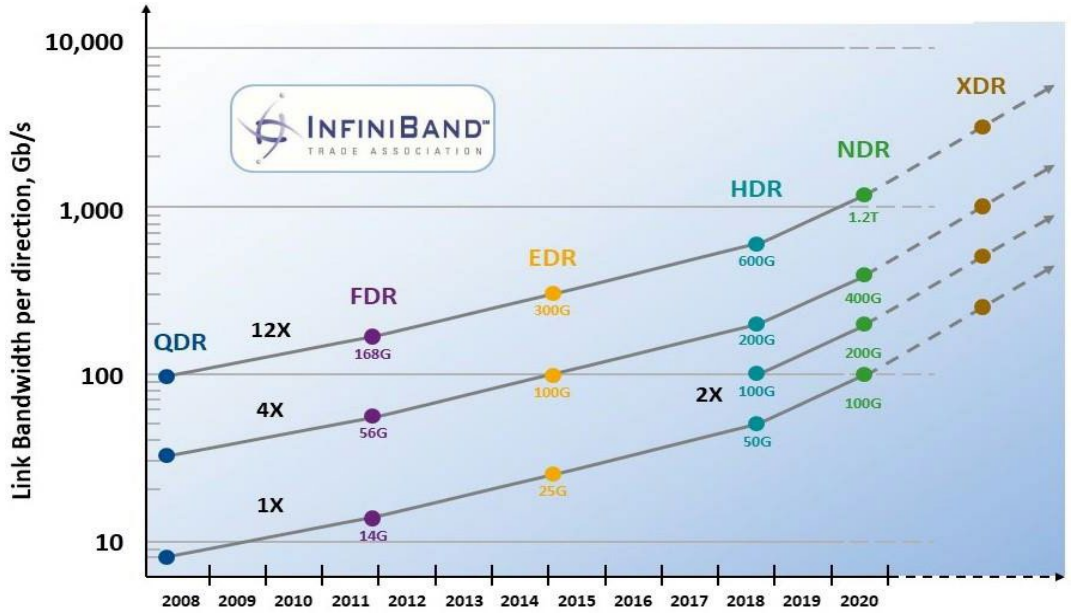

Figure 17. Market demands wrt data rates for Active Optical Cables (Source: Infiniband Trade Association)

Infiniband has been one of the most successful protocols for AOC cables and a graph of the increasing data rates from the Infiniband trade association is shown in Figure 17.

The typical cost for CXP and Infiniband products has been steadily decreasing each year since their introduction. Average price erosion is around 10% per annum. That being said the average selling price (ASP) of AOC CXP engines are in the range from €1-2/Gbps for slow speed (1 to 5 Gbps) CXP to €2-5/Gbps for high speed (10 to 12.5 Gbps) CXP engines depending on volumes.

The key challenges for AOC technology over the next decade will be to continue to increase the data rates towards 400 Gbps while keeping the size or footprint, power consumption and reliability performance similar to that of today's product set.

The goal of the industry is to design higher performance AOC solutions that bring the €/Gbps to under €1/Gbps at 400 Gbps, and closer to €0.50/Gbps at 400 Gbps over the next decade.

FIBER TO THE X (FTTX)

Market development

The FTTX market segment is composed of several fiber optic components utilized for Fiber to the home (FTTH), fiber to the curb (FTTC), fiber to the local box in the street (FTTX) based networks. The photonic components in these segments can be categorized as Passive Optical Network (PON), ONT (Optical Network Transceiver), PON OLT (Optical Line Transceiver) and PON components such as optical splitters. The data rates for the PON transceivers typically are 2.5 GPON, XG-PON1, XGS-PON, NG-PON2, 1G-EPON, 10/1 EPON, 10/10 EPON transceiver modules.

Table 2. Functional requirements FTTX

FTTX	[unit]	2025	2030	2040
Annual revenue	[€/year]	7B	10B	20B
Cost price	[€/unit]	80%	65%	40%
Energy consumption	[W]	1	0.7	0.3
Wavelength range	[nm]	1200-1600		
Footprint	[mm²]	SFP-like and chip scale packaging modules		
Output power	[dBm]	3-10		
Life cycle	[years]	3-7		
Bandwidth	[bps]	10G	50-100G	200G-1T
Swap time	[seconds]	n.a		

The trend is visible towards higher bandwidth devices; soon the 1 Gbps date rate will not be adequate anymore, especially to 'power users'. It is expected that within a period of about 5 years, a data rate of 5G is required, with a continuously upgoing trend. It must be noted that in standardization, in general 'jumps' in the maximum data rates are defined, e.g. from 100 Mbps through 1 Gbps to 10 Gbps.

Despite the larger bandwidth which will be attained, the power consumption of FTTH-solutions is subject to a downward pressure, not only required by the consumers but particularly due to regulatory forces. This is true for the entire FTTH-solution as a whole and particularly for the transceiver devices. When the output power is increased, energy consumption and non-linearities in the devices will increase as well. For this reason, the output power budget for the application is of greater importance; e.g. when a more sensitive receiver is used, the output power can be reduced thereby reducing power consumption and heat dissipation.

5G FRONT-HAUL AND BACK-HAUL

A 5G network will look as in Figure 18 below. The main difference between existing mobile networks and 5G networks is the low latency and high bandwidth requirement for the end user in a 5G network. This means that the passive power splitter that is currently used in the field, to distribute the signal to the end user, or remote antenna unit, can no longer be used. Either a lot of cabling, using parallel fiber is deployed or a wavelength splitter can be used as the ODN (Optical Distribution Unit) in the passive optical network on the front haul. This multi wavelength or WDM-PON scenario for 5G has been supported by various standardization bodies, like IEEE, and ITU-T. Relevant standards are SuperPON (IEEE P802.3cs) NGPON2 ITU 989.2 standard and Metro WDM-PON in ITU
689.3 and 689.4. The WDM-PON works with colored transceivers, which makes the logistics complicated. Ideally every end user has the same transceiver unit which can be achieved by employing a tunable laser that can select any wavelength of the typical 20 wavelength channels for upstream traffic.

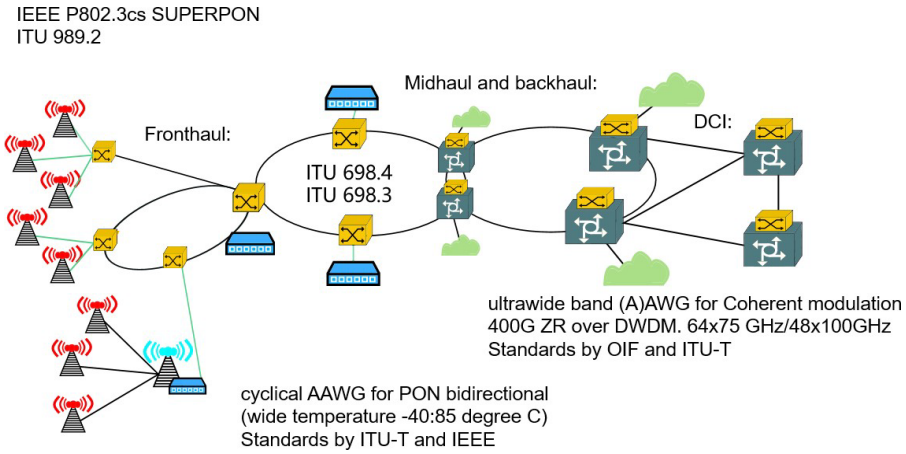

Figure 18. 5G architectural network layout

At present optical switches are based on Wavelength Selective switches (WSS) which are based on 3D MEMS in free space optics that are costly and bulky devices. In the future, new types of photonic devices of a much lower cost, a higher level of device miniaturization as compared to the current optical metro /long-haul modules will be designed.

In fronthaul, centralization of baseband processing is emerging to secure performance, flexibility, and scalability of RAN. A centralized pool of baseband processing devices, the radio equipment controllers (RECs), may serve several distributed radio equipment (RE) aggregated in clusters optimizing the use of computational resources and enabling a significant energy saving.

Starting in 2023, Google has been deploying "OCS" or "Optical Circuit Systems" in datacenters. OCS could find applicability in other parts of the network including RAN networks.

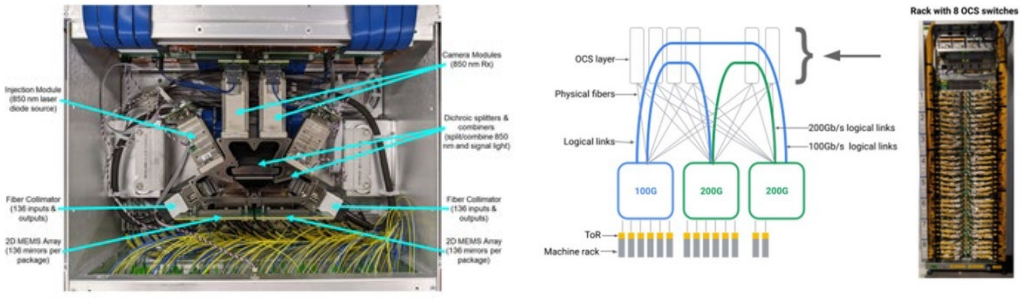

Figure 19. Google OCS Solution.

UNDERSEA AND LONG-HAUL PHOTONIC COMPONENTS

The market potential for long haul components that are utilized in undersea networks is shown below between the use of discrete photonic components and PIC based technologies. The undersea and long-haul market segment is composed of several fiber optic components utilized for these applications. The photonic component in these segments can be categorized as RGM passives which include isolators, gain flatting filters, power monitors, C and L band amplifiers where the types include EDFA and Raman, pumps (980nm, 1420nm, and 1480nm), dispersion compensators, dynamic gain equalizing filters, passive amplifiers that include filters, WEDs, power monitors etc.).

The expectation for amplifiers and pumps is to achieve a market magnitude of around $1B by the end of the decade (2030). There is an expectation that there will be a decline in the reduction of cost and ASP of photonic component parts over the next decade, in part due to an advancement in photonics technology, especially the penetration of PIC based technologies in this type of applications. PIC technologies are expected to penetrate this market very quickly from about the year 2024 and grow quickly so that a large part of the photonics-based solutions will contain integrated photonics technology in longhaul based networks. PICs for this market will tend to become hybrid solutions, and are expected to be not only a InP PIC, or a SiP PIC but a combination of the two. The trend over the next decade in this segment of the market will be high performance, as the optical components will be in situations where replacement is difficult. While InP and silicon components today are being utilized in a discrete format, hybrid PIC solutions are anticipated with InP sources, coupled to passive integrated

components in materials such as silicon, dielectric and even polymers as they mature to become additive to InP and SiPh technology platforms.

METRO AND OPTICAL TRANSPORT

As network traffic continues to grow (doubling approximately every 2 years), efforts are focused on increasing the bandwidth of the existing fiber infrastructure through more efficient use of the spectrum available. Generally, higher modulation speeds and new modulation formats are favored over WDM, as in WDM systems extra components such as multiplexers and wavelength lockers are required introducing extra cost. Increasing bitrate also carries a cost as higher speed electronic driver technology is required as is co-integration of electronics and photonics to reduce RF loss and crosstalk at bitrates in excess of 100 Gbps. Innovation using new materials and physical effects are also required for optical modulation >100 GHz as conventional technologies struggle to achieve these speeds. Photonic integration, particularly hybrid integration will be required to deliver these devices in compact form at low cost with low power consumption.

Photonic integration has had a chequered history. For simple FTTH transceivers it was found to be more expensive than using individual lasers and detectors in simple packaging. For coherent receivers, however, the requirement for phase stability means that it is well-suited. Because of the high cost of InP material, small devices interconnected with lower cost passive waveguide material are more economical if the interconnection can be simple and efficient.

Another key area to be considered is the thermal compatibility of photonics and electronics. Semiconductor lasers generally need to be cooled to near room temperature to work efficiently, whereas electronics can operate successfully up to 200 ^{o}C and can generate considerable heat. Also, dielectric photonics are less temperature sensitive than III-V material. This has led to one approach to photonic integration of having an "optical power supply (eg a comb laser)" separate from the PIC.

Coherent communications have developed quickly over the past decade and penetrated many telecommunications fiber optic interconnects and connections. These links tend to be many kilometers in length, especially over 100 km where performance tends to drive acceptance over that of low cost. While cost is becoming more of an issue for 100 km, 400 km, 600 km, >1000 km, and even 10,000 km fiber optic links, electronic DSP chips have alleviated a number of optical design issues and allowed costs to be kept reasonably under control. It is interesting to note that DSP chips are not low-cost and tend to be one of the most expensive parts of a coherent interconnect link, no matter what the reach is (over 100 km). While this trend has been successful over the past decade, the next decade will have a heavy focus on low-cost coherent for the >10 km fiber optic interconnect reach links. In the range of 10 km to 100 km, which includes classic reach distances of 25 km, 40 km, and 80 km, a number of datacenter operators as well as the telecommunications industry are seeing an interest in 'coherent-lite' product solutions in the marketplace.

To satisfy future telecom requirements, PICs will need to focus on addressing the following issues:

- Lower electrical power consumption
- Higher data rates
- Lower cost
- Smaller size
- Increased reliability
- Packaging (simplicity/cost)

To meet future optical systems requirements by 2030, transceivers are predicted to require the following:

- Modulator bandwidth of >150 GHz, low Vpi < 1 V

- New reliable material systems and integration platforms
- Driver-modulator co-design and integration
- Matching detectors with high responsivity (>80% efficient), low dark currents (<10nA) across L and C band
- High optical power (>10mW) narrow linewidth lasers (<100kHz) for multiple bands
- High operating temperature (150°C) lasers
- Temperature stable lasers with respect to temperature
- Integrated polarization rotators, splitters, and isolators > 25-30 dB ER (from all passive components)
- Low noise optical amplifiers with NF< 5 dB (integrated components)
- Small footprint electro-optical packaging (DCO-like)

It is expected that this can be achieved using these platforms:

Table 3. Transceiver requirements expected for 2025 and 2030

2025	InP monolithic integration
	Hybrid combination of silicon and InP
	Wafer bonding for electronics co-integration with photonics
	Non-hermetic sealing of chip scale packages
	Alternative modulator platforms: E/O polymers, thin film $LiNbO_3$ SiGe, etc.
	Flip-chip of electronics and photonics
2030	New material platform for lasers (better temperature performance)
	New TEC materials (integrated microTEC)
	High E/O effect materials eg polymers with temperature stability >200C

Table 4. Functional requirements for Metro and Optical transport

Metro and optical transport	[unit]	2025	2030	2040
Annual revenue	[M€/year]	28000	46000	62000
Cost price (40 channels)	[€/unit]	7000	5000	3500
Cost price/unit/channel	[€/unit/channel]	120	62.5	40
Wavelength range	[nm]	O, S, C and L band		
Reliability/channel	FIT	50	40	30
Life cycle	[years]	15 years		
Bandwidth (per channel)	[Bps]	400G	800G	1600G

Paradigm Shifts: In previous photonics roadmaps, the move to cloud computing was well publicized and a key component of providing the computing and data services for the nearly always available connectedness to on-line computing capability for both personal and business reasons. As portable electronics is growing quickly, the networking and data center infrastructure must grow to meet the demand. This shift is happening around the globe creating demands on the network bandwidth and the availability of cloud computing and data processing.

These trends are continuing and accelerating in some areas as expected. It is still true that the data center cost and power is being held essentially flat, as the capability must increase to meet the cloud-computing demands. For a computer system, this translates into increasing cost pressure as the system capacity grows, driven by more processor cores per socket, and more data bandwidth to memory. The virtualization and networking integration into the computer system tends to reduce the number of physical devices in a subsystem allowing

room for even higher levels of integration. The amount of data being consumed, however, means the number of devices interconnected is increasing rapidly.

OPTICAL TRANSCEIVERS

The packaging technology development for Data Center is currently driven by a need to process the quickly increasing amounts of data (Big Data) in a heterogeneous distributed environment (the Cloud). The optical links in the datacenter and between datacenters all have high speed transceivers with different characteristics, mainly depending on the fiber-length.

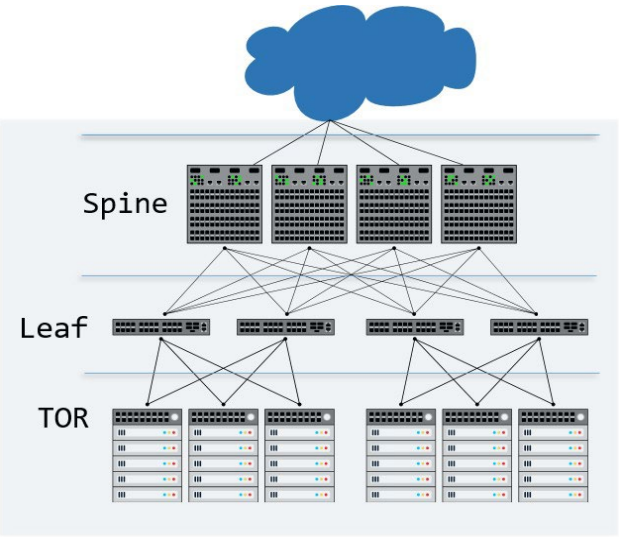

Figure 20: Data center communications architecture

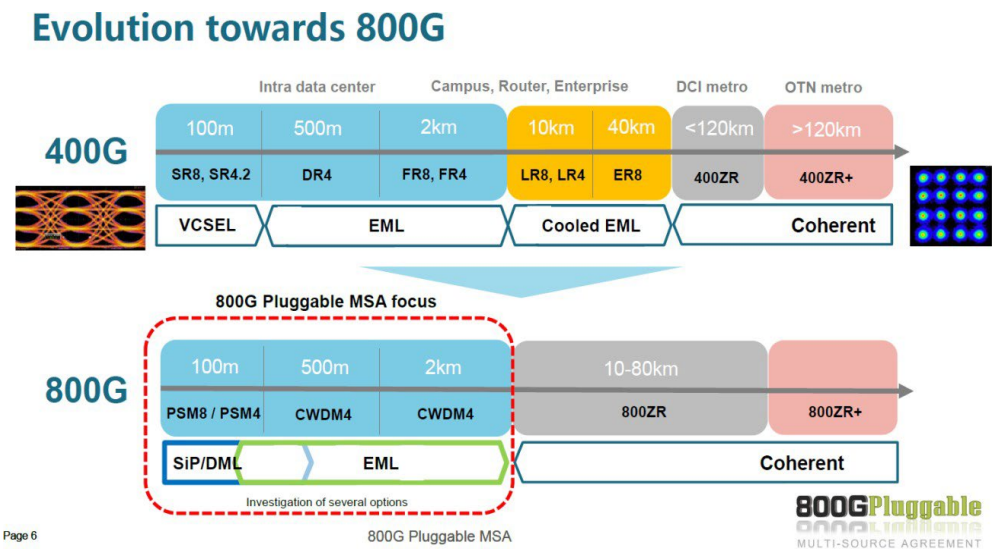

Figure 21: Evolution towards 800G in the datacenter

In Figure 20 we can distinguish intra-rack communication, mainly by AOC, intra-building communication, between TOR-leaf switch mainly supported by QSFP transceivers using 4 or 8 wavelength or 4 or 8 parallel fibers. The bulk of the cabling, using SMF fiber is between leaf and spine switch. Here, the signal speed will soon be 800 Gbps and as it connects over longer distances, like kilometers or more, it favors using a single or duplex fiber

supporting multiple wavelengths as opposed to many parallel fibers. Beyond the spine switch and for DCI (Data center interconnect) transmission will be based on small high speed coherent transceiver modules that can be multiplexed on DWDM grid in the C-band. The 800 Gbps pluggable MSA, Figure 21, has provided some guidance on the next step of speed-upgrades that can be expected. Beyond that even higher speeds require high bandwidth modulation technologies are required to support baud rate of 58 Gbaud and 116 Gbaud, 200Gbaud and above.

The attributes of these systems are discussed in the following paragraphs.

Bandwidth

The interconnect bandwidth demand is growing quickly and is expected to keep a steady growth well into the next decade. The bandwidth increase to now has been primarily provided with rapidly increasing bit rates per channel and incrementally increasing channel density. Today, optical devices are being implemented at 400 Gbps, and the forecast over the next decade is for 800 Gbps and 1600 Gbps solutions. As we go to 2030, the challenge of getting enough reach (i.e. sufficient trace lengths) to interconnect two devices while simultaneously overcoming the loss of the trace will create the need for various signaling technologies and there will be a drive to look at alternative packaging schemes such as co-package solutions. Both binary non-return to zero (NRZ) and pulse amplitude modulation (PAM) signaling will be developed with the appropriate equalization and application. Optical communication will become more broadly used as the cost and power of electrical interfaces at 50 GBaud, 100 Gbaud and above will make optical communication more attractive. Optical component devices currently exhibit bandwidths of 100 GHz which translate to 150 Gbaud using NRZ, and 300 Gbps using PAM4 signaling will be necessary over the next decade. Presently, only a few InP modulator chips can perform at 60 and 70 GHz analog bandwidths (with significant drive power), while new and novel modulators in polymers have demonstrated analog bandwidths at 100 GHz and over 250GHz with polymer plasmonics. There is an expectation that both InP as well as SiPh baseline integrated photonics platforms may look to alternative materials for the modulator function over the next decade. Silicon photonics is expected to become available in a broader range of applications because of the advantages of cost with silicon processing and packaging in addition to the opportunity to create high-density interconnections.

Power

The challenge of electrical power defines the limits of how components are integrated, that is, the number of cores on a processor chip, the density of interconnect with the trade-off of the signaling options of speed and optical vs. electrical. The footprint and content of rack electronics are constrained, and the availability of power from the utilities is capped in many installations. Therefore, the amount of power each rack and component can consume is constrained. Concern about overall power usage in datacenters becoming a significant proportion of global power consumption (estimates of > 5%) has led to a focus in design of components with higher efficiency and requiring no external cooling. This should form an important part in the design of PICs.

Thermal

The thermal capability, that is, the ability to remove heat from components and also from the frames that make up the system, is at the limit of existing capability. Because the cooling of the data center also uses a sizable amount of power, the temperatures in the data center are rising while the chip temperatures must be maintained at a constant limit from generation to generation to maintain reliability at acceptable levels. At the component level, lower thermal resistant interfaces will be developed with advanced technology. The introduction of silicon photonic components will present a new challenge to maintain optical alignment and stability in this environment.

Environment

As more large data centers are established globally, there is a broader range of environmental conditions that are encountered. Corrosive environments are encountered more frequently, and the electronics must be able to

withstand those elements. To maintain the power used for air handling in the data center, higher temperatures and higher humidity are being allowed by new ASHRAE standards. As packaged components are subjected to more moisture, the loss of interconnections will be increased which must be factored into the designs or materials chosen to minimize the impact.

Latency

Reducing the latency in communications in data centers is a paradigm shift that drives the need for new architectures such as disaggregation and the need to change from disk storage to solid-state storage to reduce latency.

Sustainability

It is expected in the next few years that datacenter product design decisions will align with good environmental practices during the next decade with few exceptions. It further suggests that stakeholders need to develop Sustainability Metrics to quantify the impact of using electronic products to replace energy intensive processes. In terms of end-of-life and recycling potential:

- Data Centers
  Owners will reuse whenever possible because of modularity
  Recycling is economically viable because of large mass in single location
- Internet of Things
  Re-use is possible for certain products in certain countries
  Because individual products are distributed, low mass, extremely heterogeneous, and contain little valuable material they are not economically viable for recycling

ASSEMBLY AND PACKAGING

The new technologies that are becoming available must meet the challenges of the previous section – bandwidth, power, thermal and environmental. Key new processor packaging technologies are being developed with some fundamental changes in the rest of the electronics industry and will impact the technology that can be leveraged. Most recently, office and desktop computing hardware could be used for memory DRAM development, CPU and MPU core development, signaling protocols, and cooling hardware.

The packaging and component technology that will be developed and integrated into Data Centers will be those that successfully developed with acceptable cost and risk of adoption that includes hybrid flexibility for the customer. Thus, the packaging for integrated silicon photonic components must utilize as much common technology as possible from the technology developed during the next decade for conventional electronic packaging. Examples would be chip-scale silicon based hermetic packaging, co-package solutions based on hybrid silicon photonics platforms, and increased density of electrical interconnects through techniques such as flip-chip bumping.

The following paragraphs discuss the packaging technology challenges for both electronics and silicon photonics.

TSV

Through Silicon Vias are enabling 2.5D silicon interposers and 3D chip stacking providing high-density interconnect, and therefore, high bandwidth capability between components. Also, glass interposers may be a factor for some applications with Through Glass Vias (TGV) providing advanced connectivity. Memory modules have already been introduced and applications will expand. The introduction of TSV has lagged expectations due to yield and cost issues which still need to be addressed.

Advanced Packaging - SiP and PoP

System-in-Package (SiP) and Package-on-Package (PoP) technologies provide the capability of optimizing cost and function in a package. Integrating voltage regulation and silicon photonics with processor chips or bridge chips

will increase. The mobile systems are where the current growth driver in this technology segment originates. However, the Data Center will adopt these advanced package technologies because the increased interconnect pins, greater memory, and additional cores when placed in close proximity enables high-bandwidth interconnection in the existing power envelope. These tradeoffs will make the appropriate technology aspects economically scalable from mobile platforms to Data Centers.

Electrical connectors for packages and cards

Electrical interconnection will continue to be the dominant interconnection for short-reach (< 3m) communications. The developing signaling standards are under discussion to go beyond 50 Gbps per channel. Electrical connectors for printed circuit board and cable communication delivering low insertion loss, flat impedance profiles and minimal crosstalk will maximize the reach of the copper interconnect at an acceptable BER. The rate of adoption of the higher speeds will depend on the ability to equalize the channels in the existing power envelope while the channel cost-performance as measured in $/Gb/s is reduced over time. The cost-performance is strongly impacted by bandwidth density. Use of Photonics signaling reduces many of these concerns and will be particularly effective when single-mode interconnects and cabling are employed. The degree to which embedded waveguides will be used depends on a number of factors but can be alleviated via on-board fly-over interconnects that use currently developed optical receptacles.

Optical Interconnects

Optical interconnect will be used more broadly. Transceivers and active optical cables (AOC) will be used for in-frame communication, potentially replacing copper interconnects in backplanes or cables when the cost, power and bandwidth tradeoffs justify the switch to optical. Integrating optical devices into packaging (co-packaging) to reduce trace length and, thus, power demand for high-bandwidth interfaces will demand advanced heterogeneous packaging and leverage the SiP and PoP technology components for increasing integration at the package level. Low-cost single-mode optical connectors will be needed to support pluggable electro-optical modules.

Silicon Photonics

The desire for higher levels of integration of optics will favor the adoption of silicon photonics. The system-level cost management, integration density, and power limit trade-offs must be carefully considered as the development of silicon photonics is pursued. The technology selected must leverage the existing silicon technology and infrastructure wherever possible to reduce both risk and cost. The use of silicon photonics will be universal as it will be fully integrated with other attractive, high performance optical technologies such as dielectric photonics and polymer photonics.

Datacentre and HPC	[unit]	2025	2030	2040
Annual revenue	[€/year]	50B	80B	150B
Cost price	[€/unit]	1000	450	200
Energy consumption	[W]	10	4	1
Wavelength range	[nm]	850 (Oband)	O/C band	O/E/S/C/L
Reliability	[%]	99	99.9	99.99
Footprint	[mm²]	2000	1000	600
Output power	[W]	5	10	15
Lifecycle	[years]	3	2.5	2
Bandwidth	[Bps]	800	1600	3200
Swap time	[seconds]	60	30	10

Speed	[material]	Optical	Optical	Optical

Table 5. Functional requirements for Datacenter and HPC optical components

PRIORITIZED DEVELOPMENT & IMPLEMENTATION NEEDS (< 5 YEARS RESULT)

Critical Milestones		
CM1	Higher speed modulation >200 Gbaud, <1 V operation (100 GHz+ analog optical bandwidth)	
CM2	Integration of electronic drivers and photonics	
CM3	Co-packaged photonics and electronics	
CM4	Low-cost tunable laser	
CM5	Operation 0 °C to 85 °C uncooled (0 °C to 70 °C in datacenter/datacom)	
CM6	Packaging (simplicity/cost)	
CM7	Non-hermetic packaging and chip scale packaging	
CM8	High responsivity high speed detector (>200 Gb/s)	
CM9	Lower power consumption (50% reduction)	

Table 6. Prioritize development and implementation needs (critical milestones)

PRIORITIZED RESEARCH NEEDS (> 5 YEARS RESULT)

Critical Milestones		
CM10	High operating temperature laser (>150 °C)	
CM11	Wavelength stable lasers with respect to temperature	
CM12	Low noise integrated optical amplifiers with NF <5 dB	
CM13	Modulators (E/O polymer/SiPh, InP etc.) >200 Gbaud	
CM14	Low loss Hollow Core Fiber with wide spectral capability (<0.2 dB/km)	
CM15	Laser, modulator and detector components compatible with hollow core fiber	
CM16	Transceiver operation -40 to +85 °C uncooled	
CM17	Uncooled laser for coherent with better than 2 GHz stability	

Table 7. Prioritized research needs (>5 years result)

Table 8. 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, e.g., coherent quadrature amplitude modulation), 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 I/Os 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 μms, 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 patterns 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, a 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 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 multi-fiber 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. 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 22 and 23 below show schematic configurations for mid-board modules with expanded beam connectors although there may be other ways to address precision requirements and contamination sensitivity. In Figure 22, the optoelectronic transceiver and the ASIC are both mounted on a system PCB. In Figure 23 the optoelectronic transceiver and the ASIC are mounted on a separate substrate that forms the base of a package. This second configuration has the 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

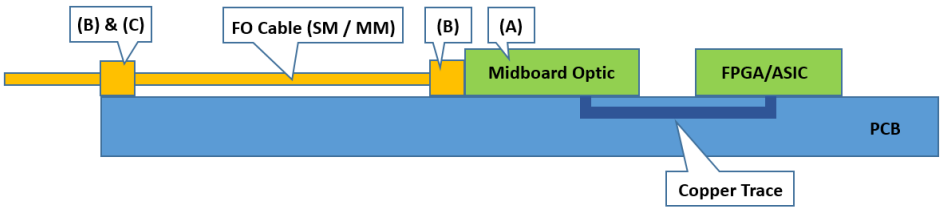

Figure 22. Technology approach for next generation systems in Application Areas 1 and 2.

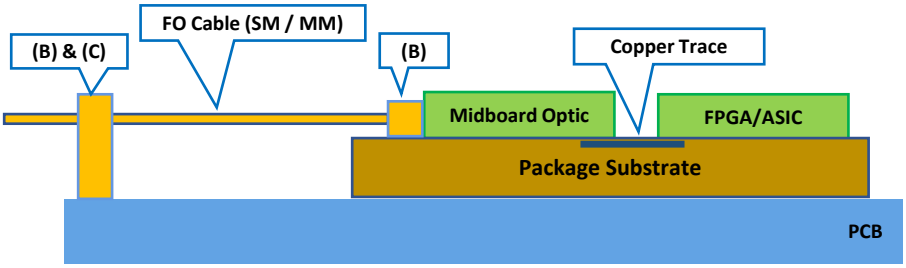

Figure 23. Co-packaged version for Application Areas 1 and 2.

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

As in the case of rack-to-world interconnects, pluggable optical 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 9: 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 for reach up to 100 meters. Because of the relative ease of packaging MM VCSELs vs. SM integrated photonic transceivers due to the relatively large, well defined emission area, relative cost benefits are possible but off-set by higher fiber costs. 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 has driven the implementation of SM, WDM-based silicon photonic links particularly in hyperscale data centers desiring future-proof structured cable.

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 an open area for the 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. Longer-term solutions for high-density inter-blade interconnections will require the use of multicore optical fibers or 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. Pluggable board-edge transceivers 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-µm 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 pluggables, 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 10. 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 is 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 24 below. [55]

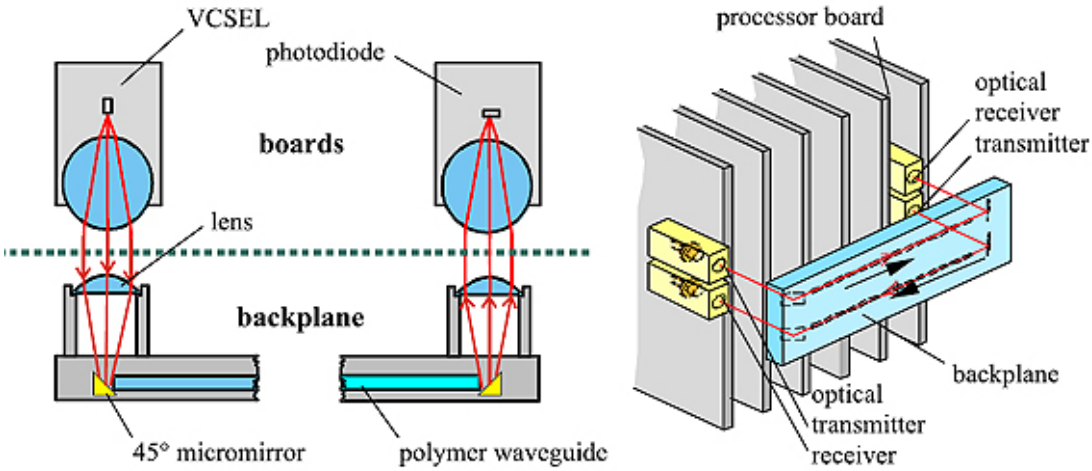

Figure 24. Approach for optical backplane connector technology.

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 [56] 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. For achieving lowest loss, glass waveguides fabricated either by ion-exchange or laser-writing have shown promise showing values <0.1 dB/cm. This is a key area where significant technology advances in materials and manufacturing processes are required. If losses cannot be consistently reduced, then the use of embedded waveguides will be limited to small (~10 cm) interposers.

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.

(A) Pluggable midboard mount transceivers with fly-over fiber-based media
(B) Optical embedded waveguides including optical interfacing to optical front panel / backplane
(C) Optical backplanes simplifying PCB to PCB optical routing
(D) Expanded-beam front panel, backplane, and midplane optical connectors

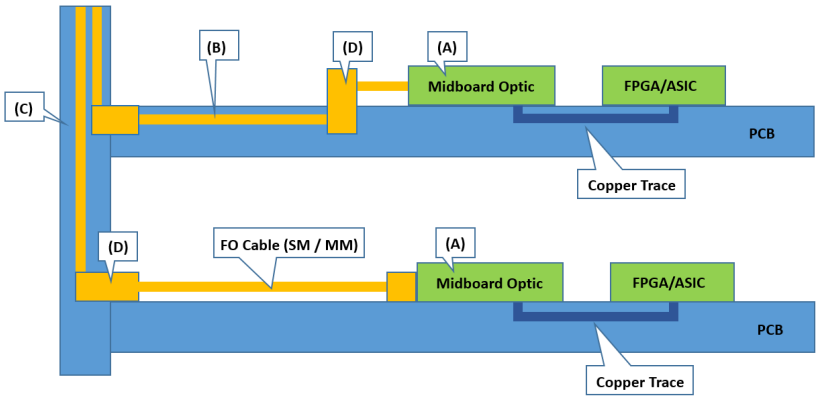

Figure 25: Example expanded beam approach for next-generation systems in Application Area 3.

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 11. 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 200 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 SiP modules. Materials with loss less than ~ 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. [57, 58] It has also been embedded within or on top of PCBs. Fraunhofer/TTM EU funded project(s). 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 wafer technology (e.g., SiP or silica-on-silicon wafers etc.) compared to 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 needed. 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 flyovers 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.

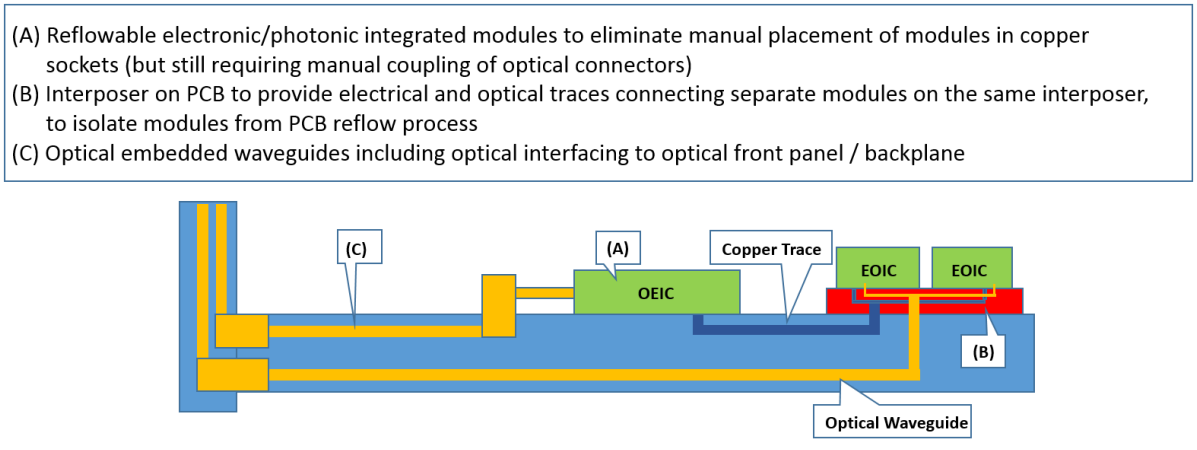

Figure 26: Technology approach for next generation systems in Application Area 4.

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).

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 12. 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 pitch in high-index-contrast semiconductor waveguides can be small (e.g., less than a few tens of μms), whereas fibers are large (80- or 125-μm diameter), and low-index-contrast waveguides (photorefractive polymers or ion-exchanged glass) require pitches > 50 μms 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 μms) connector interface.

Technology Needs

Tables 13 and 14 below summarize the near and longer-term technology requirements.

Table 13. Prioritized Development and Implementation Milestone **(≤ 2030)**	**Relative Priority**
Low-cost connector termination technology	Critical
Low-cost PIC packaging technology (high I/O 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
Table 14. Prioritized Research Milestones (>2030)	**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

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 15: 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	#/mm²	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 µms	Pigtail, pitch 250, 127 or 84 µms	Less than 84 µms, 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 and 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 tolerance requirements. For coupling of conventional single-mode fibers (mode size ~ 9 μms), sub-micron alignments must be held. For coupling of tightly confined waveguides on PICs, where mode sizes can be < 1 μm, tolerances are hundreds of nanometers. Such tolerances are very difficult 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 Chinese 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 versions of expanded-beam products are in the early stages of development and have proven to be much more difficult to design and manufacture. Other more promising approaches are described further in the Expanded Beam Technology section.

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. One approach utilizes polymer photolithographically-defined mirror and waveguides to address this problem.

There is 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 has 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 μm. 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 μm 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 be 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 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 (e.g., glass), 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 on 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, for 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.

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.

Four stages are foreseen over the next two decades:

- 2023-2025: Heterogeneous photonic solutions with advanced 3D EIC packaging; wide adoption of wavelength division multiplexing (WDM); persistence of board-edge mounting of transceiver modules but beginning of transition to co-packaging (hybrid solution of multiple switch I/O types); commercial introduction of expanded beam and physical contact co-packaged connectors
- 2025-2030: Next generation, lower power VSR SERDES in co-packaging; SM expanded-beam and physical contact connectors implemented; coarse WDM (CWDM) to 4 wavelengths will be common; possibly connector-on-board solutions
- 2030-2035: 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; fiber-to-the server begins; next generation of co-packaged connectors improve manufacturing scale.
- 2035-2040: Intra-module interconnect via waveguide interposers will begin; 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; Mid-board modules will be reflow compatible, and optically self-aligning during the

assembly process.

Further, more detailed information on the evolution of some key attribute needs of optical interconnects is provided in Table 16.

Table 16. Evolution of key interconnect application attributes.

Roadmap of Quantified Key Attribute Needs	*[unit]*	*2025*	*2030*	*2035*	*2040*
Optical Connector Cost/Bit	$/Gbps	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)	13.6 x 4.8 (XB)	13.6 x 4.8 (XB)	13.6 x 4.8 (XB)
Gbps speed per Wavelength Channel	Gbps	100	400	400	400
Wavelengths per Fiber Core or Waveguide	#	4	4	8	16
Cores/Fiber	#	1	3, 4	7, 16	7, 16
Fiber Pitch	µm	127	127	84	84
Fibers/ Connector	#	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	3db	3db	3db	3db
Key Connector Types	----	Chip-X[3], PCIe (SM)	PCI-X, Chip-X, Co-packaged	Other Co-packaged, TBD	Other Co-packaged, TBD
Sockets/ Interposers	(status)	Connector or Chip OEM/OSAT?	Chip-Integrated	Chip-Integrated	Chip-Integrated
Cables	----	100s of Racks,	10s of Racks, WGs[4]	10s of Racks, WGs	10s of Racks, WGs
System	----	3D-E0-SiP	Monolithic E0-SoC	Monolithic E0-SoC	Monolithic E0-SoC
Roadblocks	----	SiP[5] Integration (SoC)	SiP/SoC Modules; multicore fiber coupling	SiP/SoC Modules; multicore fiber coupling; T-stable muxes	SiP/SoC Modules; multicore fiber coupling; T-stable muxes

Comment	----	Mid-Board, I/O Standards	Modularization < Connectors[6]		

In Production Modest Difficulties Significant Roadblocks Major Technology Challenges

* Connector costs 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 µms in size, lateral alignment to around 1 µm is adequate. When coupling components with smaller waveguide modes, say 2 µms, for a SiP PIC, lateral alignment to around 0.2 µms is needed for the same level of loss.

Similarly, the spectral bandwidth required for components is dependent not only on 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, the presence of the legacy infrastructure may be an obstacle to the development of the much needed, new infrastructure. This is because the legacy optical interconnect infrastructure was built around products for long-haul or specialty communications applications where connector cost is negligible compared to the total system cost. 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. 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 µm clearance in the connector ferrule; and must be bonded in place. Finding an economical path to automating this process is a challenging proposition.

Nevertheless, the 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-µm 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 is not yet fully perfected is the use of solder surface tension and etched stops to position elements relative to each other with sub-µm accuracy. This precise positioning technique requires sub-µm 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 µms) 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-µm 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.

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)
- SiP 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. 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 00has 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 they are slow. Also, testing at the module level wastes resources when defective chips are packaged
- Wafer-level testing needs 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

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 interconnects, 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 produced 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 sub-micron 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. Resolution is also limited to ~1 μm.

- 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-μm 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 interconnects:

- Subminiaturization Barriers to Conventional Fiber Optic Connector Technology: With connector housings at several millimeters and the optical fibers they encase at 125 µms, but waveguide cores less than 10 µms in size, there appears to be room for miniaturization.

- Subminiaturization Barriers of the Electronic Packaging Platform (e.g., HDI, 3D, Printed Electronics): Requirements for pitch below 200 µm 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 interconnects, 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.

- 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 including: 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) has been 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 of 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 SiP 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 SiP package may replace the outboard PCB assemblies with highly-integrated 3D PIC packages

REFERENCES

1. *Silicon Photonics in Pluggable Optics White Paper*. 2021 December 6th, 2021 [cited 2023 April 4th, 2023]; Available from: https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/transceiver-modules/silicon-photonics-wp.html.

2. Nezami, M.S., et al., *Packaging and Interconnect Considerations in Neuromorphic Photonic Accelerators.* IEEE Journal of Selected Topics in Quantum Electronics, 2023. **29**(2).

3. Piehler, D. *Optical interconnects in enterprise and hyperscale datacenters*. in *Optical Interconnects XX 2020, February 4, 2020 - February 6, 2020*. 2020. San Francisco, CA, United states: SPIE.

4. Hardy, S. *New OSFP-XD form factor could challenge co-packaged optics*. 2021 June 10th, 2021 [cited 2023 April 4th, 2023]; Available from: https://www.lightwaveonline.com/optical-tech/transmission/article/14205029/new-osfpxd-form-factor-could-challenge-copackaged-optics.

5. *The migration to 400G/800G: Part II*. 2023 [cited 2023 April 4th, 2023]; Available from: https://www.commscope.com/insights/the-enterprise-source/migration-to-400g800g-the-fact-file-part-2/.

6. Fathololoumi, S., et al. *Silicon photonic integrated circuit for co-packaging with switch ASIC*. in *Silicon Photonics XVI 2021, March 6, 2021 - March 11, 2021*. 2021. Virtual, Online, United states: SPIE.

7. Dykes, P. *OFC 2023: Broadcom launches 51.2T CPO optical engine*. 2023 March 1st, 2023 [cited 2023 April 4th, 2023]; Available from: https://opticalconnectionsnews.com/2023/03/ofc-2023-broadcom-launches-51-2t-cpo-optical-engine/?utm_source=OpCons&utm_campaign=2023&utm_medium=Email&utm_term=TechnologiesProducts&utm_content=Main.

8. *Broadcom Ships Tomahawk 5, Industry's Highest Bandwidth Switch Chip to Accelerate AI/ML Workloads*. 2022 August, 16th, 2022 April 4th, 2023]; Available from: https://investors.broadcom.com/news-releases/news-release-details/broadcom-ships-tomahawk-5-industrys-highest-bandwidth-switch.

9. Mahajan, R., et al., *Co-Packaged Photonics For High Performance Computing: Status, Challenges And Opportunities.* Journal of Lightwave Technology, 2022. **40**(2): p. 379-392.

10. *Nubis Communications Announces Industry's Highest Density and Lowest Power Optical Interconnect*. 2023 February 22nd, 2023 [cited 2023 April 4th, 2023]; Available from:

https://www.businesswire.com/news/home/20230222005042/en/Nubis-Communications-Announces-Industry%E2%80%99s-Highest-Density-and-Lowest-Power-Optical-Interconnect.

11. *Design Considerations of Optical Connectivity in a Co-Packaged or On-Board Optics Switch*. 2022 2022 [cited 2023 April 20th, 2023]; Available from: https://www.onboardoptics.org/white-papers.

12. Sabi, A.K.a.B. *A Breakthrough Decades in the Making: Intel's Co-packaged Optics Demonstration*. 2022 September 27th, 2022 [cited 2023 April 6th, 2023]; Available from: https://medium.com/intel-tech/a-breakthrough-decades-in-the-making-intels-co-packaged-optics-demonstration-ca1349c112c9.

13. Gradkowski, K., et al. *Demonstration of a Single-Mode Expanded-Beam Connectorized Module for Photonic Integrated Circuits*. in *European Conference on Optical Communication (ECOC)*. 2022. Electr Network: Ieee.

14. Israel, A., et al. *Photonic plug for scalable silicon photonics packaging*. in *Conference on Optical Interconnects XX*. 2020. San Francisco, CA: Spie-Int Soc Optical Engineering.

15. Pezeshki, B., et al. *MicroLED Array-based Optical Links Using Imaging Fiber for Chip-to-chip Communications*. in *Optical Fiber Communications Conference and Exhibition (OFC)*. 2022. San Diego, CA: Ieee.

16. Brusberg, L., et al., *Glass Platform for Co-Packaged Optics.* IEEE Journal of Selected Topics in Quantum Electronics, 2023. **29**(3: Photon. Elec. Co-Inte. and Adv. Trans. Print.): p. 1-10.

17. Janta-Polczynski, A., M. Robitaille, and Ieee. *Optical fiber pigtails integration in co-package*. in *72nd IEEE Electronic Components and Technology Conference (ECTC)*. 2022. San Diego, CA: Ieee.

18. Buscaino, B., et al., *External vs. Integrated Light Sources for Intra-Data Center Co-Packaged Optical Interfaces.* Journal of Lightwave Technology, 2021. **39**(7): p. 1984-1996.

19. *Cisco Annual Internet Report (2018-2023) White Paper*. 2020 March 9, 2020 11/14/2020]; Available from: https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html.

20. *Photonic integrated circuit*. 2023 [cited 2023 April 6th, 2023]; Available from: https://en.wikipedia.org/wiki/Photonic_integrated_circuit.

21. Butler, D., *High-density FAUs and optical interconnection devices and related methods*, USPTO, Editor. 2023, Corning, Incorporated: USA. p. 101.

22. Cho, J.K., et al. *Optical performance and reliability assessment from self-aligned single mode fiber attach for O-band silicon photonics platform*. in *72nd IEEE Electronic Components and Technology Conference, ECTC 2022, May 31, 2022 - June 3, 2022*. 2022. San Diego, CA, United states: Institute of Electrical and Electronics Engineers Inc.

23. Nauriyal, J., et al. *Low-loss, Single-shot Fiber-Array to Chip Attach Using Laser Fusion Splicing*. in *2022 IEEE Photonics Conference (IPC)*. 2022.

24. Guan, L.T., et al. *FOWLP and Si-Interposer for High-Speed Photonic Packaging*. in *IEEE 71st Electronic Components and Technology Conference (ECTC)*. 2021. Electr Network: Ieee Computer Soc.

25. Mounier, E.M., Jean-Louis, *Silicon Photonics and Photonic Integrated Circuits 2019. Si photonics, beyond the tipping point!* . 2019.

26. Schroder, H., et al., *Hybrid photonic system integration using thin glass platform technology.* Journal of Optical Microsystems, 2021. **1**(3): p. 18.

27. Meany, T., et al., *Laser written circuits for quantum photonics.* Laser & Photonics Reviews, 2015. **9**(4): p. 363-384.

28. Izawa, T. and H. Nakagome, *OPTICAL WAVEGUIDE FORMED BY ELECTRICALLY INDUCED MIGRATION OF IONS IN GLASS PLATES.* Applied Physics Letters, 1972. **21**(12): p. 584-&.

29. Tervonen, A., B.R. West, and S. Honkanen, *Ion-exchanged glass waveguide technology: a review.* Optical Engineering, 2011. **50**(7): p. 15.

30. Schwietering, J., et al. *Integrated optical single-mode waveguide structures in thin glass for flip-chip PIC assembly and fiber coupling*. in *70th IEEE Electronic Components and Technology Conference (ECTC)*. 2020. Electr Network: Ieee Computer Soc.

31. Brusberg, L., et al. *Study of Temperature and Time Dependent Performance Alteration of Silver Ion-exchanged Waveguides in Glass*. in *8th IEEE Electronics System-Integration Technology Conference (ESTC)*. 2020. Vestfold, NORWAY: Ieee.

32. Nisar, S., L. Li, and M.A. Sheikh, *Laser glass cutting techniques-A review.* Journal of Laser Applications, 2013. **25**(4): p. 11.

33. Terbrueggen, R. *From Proof of Principle to 98.5% Yield of a high-speed laser processing tool*. in *Conference on Laser Applications in Microelectronic and Optoelectronic Manufacturing (LAMOM) XXV*. 2020. San Francisco, CA: Spie-Int Soc Optical Engineering.

34. Grenier, J.R., et al. *Ultrafast laser singulation of optical circuits with optical quality endfacets*. in *Conference on Frontiers in Ultrafast Optics - Biomedical, Scientific, and Industrial Applications XXII at SPIE LASE Conference*. 2022. Electr Network: Spie-Int Soc Optical Engineering.

35. Kopp, C., et al., *Silicon photonic circuits: On-CMOS integration, fiber optical coupling, and packaging.* IEEE Journal on Selected Topics in Quantum Electronics, 2011. **17**(3): p. 498-509.

36. Marchetti, R., et al., *Coupling strategies for silicon photonics integrated chips Invited.* Photonics Research, 2019. **7**(2): p. 201-239.

37. Weninger, D., et al., *High density vertical optical interconnects for passive assembly.* Optics Express, 2023. **31**(2): p. 2816-2832.

38. Mahajan, R., et al. *Embedded Multi-Die Interconnect Bridge (EMIB) - A High Density, High Bandwidth Packaging Interconnect*. in *66th IEEE Electronic Components and Technology Conference (ECTC)*. 2016. Las Vegas, NV: Ieee Computer Soc.

39. Pitwon, R.C.A., et al., *Pluggable electro-optical circuit board interconnect based on embedded graded-index planar glass waveguides.* Journal of Lightwave Technology, 2015. **33**(4): p. 741-754.

40. Jiang, L., et al., *SmartPrint Single-Mode Flexible Polymer Optical Interconnect for High Density Integrated Photonics.* Journal of Lightwave Technology, 2022. **40**(12): p. 3839-3844.

41. Luo, H., et al., *Low-loss and broadband fiber-to-chip coupler by 3D fabrication on a silicon photonic platform.* Optics Letters, 2020. **45**(5): p. 1236-1239.

42. Rhee, H.W., et al. *Chip-to-chip optical interconnect using direct optical wire bonding*. in *Conference on Optical Interconnects XXII at SPIE OPTO Conference*. 2022. Electr Network: Spie-Int Soc Optical Engineering.

43. Hoose, T., et al. *Multi-Chip Integration by Photonic Wire Bonding: Connecting Surface and Edge Emitting Lasers to Silicon Chips*. in *Optical Fiber Communications Conference and Exhibition (OFC)*. 2016. Anaheim, CA: Ieee.

44. Ishigure, T. and Ieee. *GI-core Multimode and Single-mode Polymer Waveguides for High-density Co-packaging*. in *Optical Fiber Communications Conference and Exhibition (OFC)*. 2022. San Diego, CA: Ieee.

45. Billah, M.R., et al., *Hybrid integration of silicon photonics circuits and inp lasers by photonic wire bonding.* Optica, 2018. **5**(7): p. 876-883.

46. Lindenmann, N., et al., *Connecting Silicon Photonic Circuits to Multicore Fibers by Photonic Wire Bonding.* Journal of Lightwave Technology, 2015. **33**(4): p. 755-60.

47. Flory, N., et al. *Highly reliable polymer waveguide platform for multi-port photonic chip-packaging*. in *IEEE 71st Electronic Components and Technology Conference (ECTC)*. 2021. Electr Network: Ieee Computer Soc.

48. Flory, N., *On-Board Optics Expand Bandwidth and Reduce Power Use in the Data Center*, in *Photonics Spectra*. 2023, Photonics Media.

49. Takenobu, S. and T. Okazoe. *Heat resistant and low-loss fluorinated polymer optical waveguides at 1310/1550 nm for optical interconnects*. in *European Conference and Exposition on Optical Communications, ECOC 2011, September 18, 2011 - September 22, 2011*. 2011. Geneva, Switzerland: Optical Society of America (OSA).

50. Wei, W., et al. *Loss Characteristics of Polymer Optical Waveguide at 1310 nm Wavelength on An Optical Printed Circuit Board*. in *19th International Conference on Optical Communications and Networks, ICOCN 2021, August 23, 2021 - August 27, 2021*. 2021. Qufu, China: Institute of Electrical and Electronics Engineers Inc.

51. Yasuhara, K., F. Yu, and T. Ishigure, *Circular core single-mode polymer optical waveguide fabricated using the Mosquito method with low loss at 1310/1550 nm.* Optics Express, 2017. **25**(8): p. 8524-8533.

52. *Multimode Waveguide Interconnect System Design for Photonic Circuit Integration*. 2022 April 2022 [cited 2023 April 17th, 2023]; Available from: https://www.onboardoptics.org/white-papers.

53. Merget, F., et al. *Glass Molded Optical Interposers for Wafer Scale Datacom Component Packaging*. in *European Conference on Optical Communication (ECOC)*. 2021. Bordeaux, FRANCE: Ieee.

54. Minkenberg, C. and Ieee. *Viability of Fiber-to-the-Server: Are We There Yet?* in *International Conference on Photonics in Switching and Computing (PSC)*. 2018. Limassol, CYPRUS: Ieee.

55. Savage, N., *Linking with Light: High-speed optical Interconnects*, in *IEEE Spectrum*. 2002.

56. Jou, J.-J., T.-T. Shih, and C.-L. Chiu, *400-Gb/s optical transmitter and receiver modules for on-board interconnects using polymer waveguide arrays.* OSA Continuum, 2018. **1**(2): p. 658-667.

57. Brusberg, L., et al. *Single-mode glass waveguide substrate for PIC packaging*. in *2019 IEEE CPMT Symposium Japan (ICSJ)*. 2019.

58. Neitz, M., et al. *Insertion Loss Study for Panel-Level Single-Mode Glass Waveguides*. in *Conference on Optical Interconnects XVII*. 2017. San Francisco, CA: Spie-Int Soc Optical Engineering.

59. Yeniay, A., et al. "Ultra-low-loss polymer waveguides." *Journal of Lightwave Technology*, vol. 22, no. 1, Jan. 2004, pp. 154–158, https://doi.org/10.1109/jlt.2003.822206.