

Applications for Embedded Optic Modules in Data Communications



White Paper

Derek Vaughan, Rob Hannah, and Mitchell Fields
Avago Technologies

Introduction

Bandwidth demands of data center, computing, and telecommunication applications are driving networking I/O requirements to levels that are challenging to achieve with edge-mounted copper and optical interfaces; the only solution is embedded parallel optics. In fact, since 2000, core switches, routers, and high-performance computers have employed embedded parallel optics to deliver the I/O performance between racks in multi-chassis configurations. Today, with the ever increasing demand for data and bandwidth, embedded parallel optics from Avago Technologies have become the solution of choice in more systems and applications from enabling the interconnect between nodes of the world's most powerful supercomputers to transmitting data from sensors in challenging medical imaging environments.

Parallel-optical modules are devices that have more than one lane of data, transmitter and/or receiver, in a single package. Common lane counts within one device include 4 transmitter/receiver pairs, 12 transmitters, 12 receivers, and 12 transmitter/receivers. Embedded parallel optics are parallel-optic modules that are not mounted at the board edge but rather directly onto the PCB in order to gain advantages in density, signal integrity, thermal management, I/O consolidation, and EMI/ESD. Each lane can be independent or be part of a high-capacity channel to cost effectively provide the target bandwidth. For example, aggregate bandwidths of 120 Gbps can be achieved with a 12-lane module operating at 10 Gbps per lane. This parallel-optical implementation is analogous to the electronic-signaling approach of achieving a target bandwidth using a parallel electrical interface, for example XAUI and PCIe.

Avago Technologies has been the world leader in parallel optics since 2000 and currently delivers in excess of 200,000 modules per quarter. This white paper describes the latest embedded parallel optics products from Avago: MicroPOD™ and MiniPOD™. We detail the features, benefits, complementary technologies, and example applications.

Avago Parallel Optic Embedded Modules

Avago has introduced two new twelve-channel, parallel, high performance fiber optics modules for short-range multi-lane data communication and interconnect applications: MicroPOD and MiniPOD. Comprised of separate transmitter and receiver modules, each with 12 independent lanes capable of > 10 Gbps signaling, the MicroPOD and MiniPOD products enable links with > 120 Gbps capacity over ~100 meters of OM3 grade multimode fiber. The transmitter modules use an Avago monolithic 850-nm VCSEL array and the receiver modules use Avago monolithic GaAs PIN diode arrays. Both use internally developed ICs for optimal performance in demanding applications. MicroPOD and MiniPOD have similar performance specifications and feature a new low-cost optical turn connector, called Prizm™, for coupling to 12-fiber ribbon cable. MicroPOD modules, shown in Figure 1, offer extremely high density, measuring a mere 7.8 mm x 8.2 mm x 3.9 mm and are electrically connected to the host PCB via a high density 9x9 pin Micro-LGA with 0.7424 mm pitch.




Figure 1. MicroPOD™ a) With and without Prizm™ Connector b) On a Dime and c) Showing the Micro-LGA interface

For applications that do not require the density of MicroPOD, the MiniPOD (shown in Figure 2) offers a convenient solution. Using MicroPOD as an engine, MiniPOD has a footprint of 18.6 mm x 22 mm and a height of either 14.5 mm for a flat fiber cable housing or 15.6 mm for the round cable housing. This footprint is determined by the size of the industry-standard 9x9 pin Meg-Array® connector used for the electrical interface.

Two slightly different versions of the MiniPOD module are offered, one that interfaces to a flat ribbon cable and the other to a round ribbon cable. Advantages of flat ribbon cables include lower profile, and dense tiling of modules, while round cables offer ruggedization and more flexibility in fiber routing.




Figure 2. MiniPOD™ Transmitter and Receiver Modules with a) Round Cable and b) Flat Cable: shown with and without dust covers.

Examples of the dense tiling capability of MicroPOD and MiniPOD are shown in Figure 3. This capability allows system designers to take full advantage of the density of the embedded parallel optics modules as there is no keep-out required for the fiber attachment.




Figure 3. MicroPOD™ and MiniPOD™ Transmitter and Receiver Modules shown in high density tiled configurations.

Complementary Technologies

In this section, we describe technologies that complement the MicroPOD and MiniPOD products: Meg-Array® connectors, Prizm connectors, fiber flex circuits, MT-ferrules (12/24/48/72 fiber), MPO (Multi-fiber Push-On) adapters, and blind-mate connectors. These technologies are all commercially available from multiple industry sources.

MiniPOD modules utilize Meg-Array connectors. These RoHS compliant, Telcordia qualified, Ball-Grid-Array terminated connectors were designed specifically for 10 Gbps high density applications.

The detachable feature of the Prizm connector, developed and licensed by USCONEC and available from multiple cable manufacturers, is illustrated in Figure 5. Prizm connectors enable the flexibility to first populate a system with embedded optical modules and subsequently achieve optical connectivity with a separate cable harness.

MTP connectors are commonly available populated with 12 and 24 fibers, but they are also available with 48 and even 72 fibers. In addition, for applications requiring increased robustness and durability to strain relief larger 48 fiber and 72 fiber trunk cables, high density metallic Circular MT connectors are also commercially available.




Figure 4. MiniPOD™ Round Cable Housing showing Port Plug, Meg-Array® Connector, and Close-Up of Meg-Array Electrical Pins.




Figure 5. Avago MicroPOD™ With Prizm™ Optical Connector



Figure 6. a) 72 Fiber MTP Connector and b) Circular MT and 1x3 Circular MT Ruggedized Connectors¹

Note:

1. Pictures supplied by Molex

For applications in highly contaminated environments or requiring frequent mate / de-mates, where there might be concerns regarding increased insertion loss due to dust, embedded optic modules could be readily deployed with optional expanded beam (lensed) MT ferrule interconnects. An Expanded Beam Lensed MT connector, available from Molex, is shown in Figure 7.




Figure 7. Expanded beam Lensed MT Interconnect for increased mating cycles and reduced loss in contaminated environments*

Some of the largest data communication systems in the industry already use parallel optic modules for backplane connectivity in their multi-chassis switch fabrics. MicroPOD and MiniPOD modules combine perfectly with MT based bind-mate adapters to provide a reliable, robust, and cost effective solution to optical backplane connectivity requirements. Figure 8 is an example of how MiniPOD transmit and receive modules might be deployed with a high density backplane interconnect, such as Molex's HBMT™ Backplane Connector. Each channel supports 10 Gbps, so these four transmit plus four receive 12-channel modules provide around 480 Gbps of duplex communication capacity.




Figure 8. MiniPOD Modules in a High Density Backplane Application [1]

Note 1. Inset picture of HBMT™ provided by Molex

For high fiber-count interconnects in backplanes and cross-connect systems, terminated optical flex assemblies provide a manageable means of fiber routing from card-to-card or shelf-to-shelf.

These flex circuits are available in flame retardant materials and can be combined with high density front panel and blind-mate back-plane MTP interconnects to realize ideal solutions for high-density applications.

The 2x8 ganged MTP adapters shown in Figure 11, available from USCONEC, require a 71.9 x 30.3 mm faceplate cutout. When compared to the horizontal and vertical pitch requirement for a CXP module, which is 27 mm by 16.5 mm respectively, then the density increase for a faceplate populated with MTP connectors is a factor of 3.5. When we consider that an MTP connector could contain as many as 48 or 72 fibers, compared to 24 fiber channels for a CXP optical module, then the full duplex density of an embedded optics implementation is up to $(16 \times 3 \times 120 \text{ Gb}/71.9 \times 30.4 =) 2.6 \text{ Gb}/\text{mm}^2$ compared to $(120 \text{ Gb}/ 27 \times 16.5 \text{ mm} =) 0.27 \text{ Gb}/\text{mm}^2$ for a CXP based design.




Figure 9. Molex optical FlexPlane™ assembly Managed Ribbon Fiber Bundle




Figure 10. Example of optical FlexPlane™ Assembly used with Backplane BMT™ & HBMT™ Adapters



Figure 11. Ganged MTP Adapter [2]

Note 2. Illustration from USCONEC

Advantages of Embedded Optics

Embedded optics allow the flexibility to position the high-speed modules at the optimal position on the host board to address engineering challenges such as signal integrity and thermal management, whilst simultaneously achieving maximum density. These advantages are reviewed and discussed in the following sections.

Signal Integrity

High-speed electrical signals, especially at and beyond 10 Gbps, progressively degrade over increasing trace lengths in PCB material. This phenomenon, due to capacitive skin effects, is demonstrated in figure 12, where a 350-mV single-ended 10-Gbps electrical eye is propagated over progressively longer lengths of stripline. These test results are best case as no layer changes or vias are incorporated in the layout design and the test board features high quality SMA connectors. After just 5" the signal integrity,

as measured by eye opening, is close to the limit of IEEE802.3ba standards compliancy. By 10" the signal integrity is already below acceptable limits. Complete closure of the eye occurs over a length of 20".

To address signal integrity challenges, designers can utilize high performance PCB materials, transmitter de-emphasis, receiver equalization, and signal retiming using CDRs. All of these techniques add cost and power to the overall solution.

Alternatively, the designer can work to minimize trace lengths between the high-speed sources and terminations. For example, instead of using edge-mounted optics and necessarily long trace lengths, embedded optical modules can be easily positioned within five inches of the host ASIC.

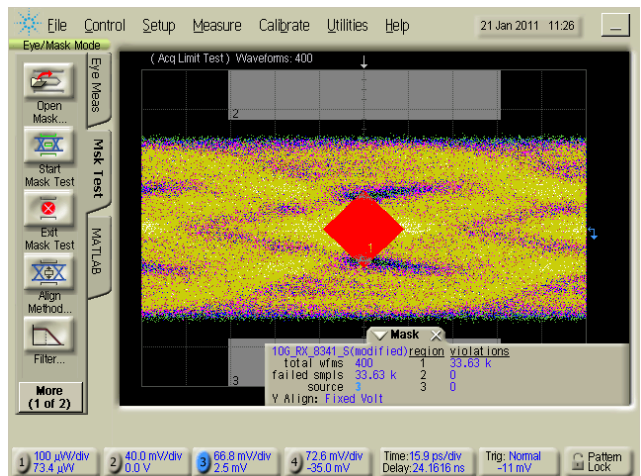

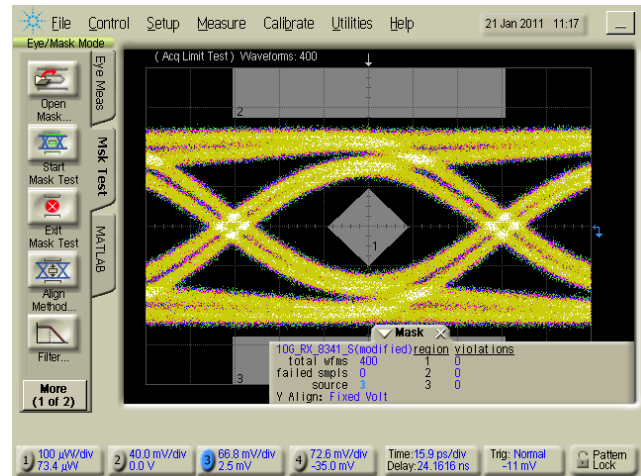
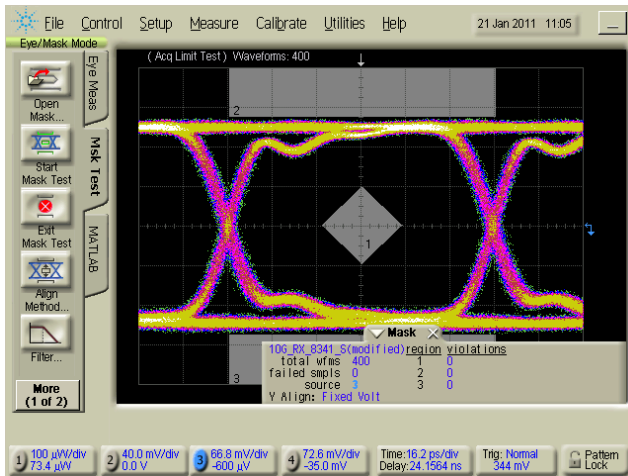


Figure 12. Degradation of signal integrity along a strip-line trace of length a) zero inch b) 5 inches c) 10 inches and d) 20 inches on a IEEE 802.3ba mask.

Dense Tiling of Optical Modules and High-Bandwidth Ports

As previously mentioned, the novel top-attach Prizm connector enables dense tiling of optical modules; thus permitting a maximum I/O density per unit area of host PCB. Furthermore, in conjunction with high fiber count 48 fiber and 72 fiber MPO/MTP connectors, the embedded modules enable extreme I/O density at the card edge. Even considering applications with modest density requirements, one MiniPOD™ transmitter and receiver pair offers the same bandwidth as a CXP module but in a smaller footprint and consumes less front panel area.

Cooling Efficiency

In the same way that host ICs can be positioned for optimal thermal management, embedded optic modules provide flexibility in positioning to facilitate thermal engineering.

A typical layout for a high capacity rack top 1RU switch with edge-mounted optics is shown in Figure 13. Cool air drawn in through the rear of the box is heated as it flows over switch ICs and other circuit components before it passes over the optical modules mounted on the front panel, where air exhausts through small perforated holes in the front face-plate. Pre-heating of the air by high power dissipation ICs presents significant thermal challenges for the edge mounted optics, and can limit the system density.




Figure 13. High Density Datacenter Switch Showing Air Cooling

In comparison, embedded optical modules can be positioned such that they are not subjected to pre-heated air and, since MTP adapters consume less face-plate area than edge-mounted optics, there would be more area on the faceplate for air exhaust. Higher air velocity through the switch and correspondingly improved fan efficiencies are other potential benefits.

To understand and quantify the practical benefits of an embedded optic implementation, a thermal simulation comparison was made between a system populated with CXP modules, similar to the one illustrated in Figure 13, and one populated with MiniPOD modules. Both implementations offered the same optical communication capacity and dissipated the same amount of power, but

the embedded modules were optimally located near the air inlet. These simulations revealed that the MiniPOD modules operated up to 13 degrees Celsius cooler than their front panel mounted CXP equivalents.

Embedded optic implementations readily interface to heat pipes, heat sinks and spreaders, and cooling air or fluid can be concentrated directly onto the heat sink to avoid directing air through the system enclosure. A heat sink can be designed to cover multiple embedded optic modules, and perhaps additional active components, to minimize the number of piece parts and thus reduce cost and manufacturing complexity.

EMI Reduction

High-density systems may require 24, 96, or more individual high speed channels interfacing to optical modules on the front panel. This trend toward high frequency dense signaling will continue as communication rates increase to 14 and 25 Gbps. EMI containment is challenging in these systems and significant engineering resources are consumed mitigating the problem. In contrast, with embedded optics the high-frequency signals are moved far away from port openings in the front of the chassis and are replaced with small non-radiating bundles of optical fibers that could be passed through EMI containment apertures before arriving at the front panels.

Embedded optical modules facilitate a reduction in the required faceplate area and enable EMI containment that would reduce costs and simplify compliance to FCC regulations. Furthermore, it is obvious that Electro-Static Discharge to the front panel mounted MPO connector is much less of a concern than for pluggable modules. The use of diecast adapters, available from Molex and other connector companies, further mitigates EMI issues.




Figure 14. Diecast EMI adapters reduce EMI while providing multiple mounting options for improved cable management¹

Note:

1. Picture supplied by Molex

Application Examples

We conclude by highlighting four applications for Avago MicroPOD and MiniPOD embedded parallel optics. The applications are driven by requirements for density, signal integrity, and I/O consolidation.

1. High-Density Switching

Many data center applications are optimized by having the maximum number of servers included in one switching element. However, the front-panel area of the switching platform can limit the number of ports available to the switch fabric (typically to 48). In figure 15, we illustrate a 192-port switch that is enabled by embedded parallel optics. The 12-channel MicroPOD modules are placed in close proximity to the switch ASIC for optimal signal integrity. At the board edge there are 8 MTP-24 and 4 MTP-48 adapters. This configuration highlights the flexibility of bandwidth aggregation that embedded optics affords. The MTP-24 ports provide 120 Gbps of full-duplex bandwidth and the MTP-48 ports provide 240 Gbps.




Figure 15. Application of tiled MicroPOD modules in a high density switch

2. Blade Servers

Blade server systems are dense modular server systems that provide tight integration amongst multiple (typically up to 16) servers with storage, switching, I/O, cooling, and power sub-systems. They are advantageous compared to rack-mounted servers in terms of power and space efficiency. Increased virtualization, and the steady increase in computational performance of individual CPUs, stresses traditional I/O technologies within bladed systems. Current system deployments adopt a high-speed electrical interface between the server blades and the I/O modules through a complex electrical mid-plane. These architectures present complicated signal integrity and thermal challenges to system designers.

Electrical I/O from the blade server mezzanine card could alternatively be replaced with optical I/O, and the complicated electrical mid-plane replaced with a simple optical equivalent. An optical mid-plane would provide all the high-speed connectivity between the servers and any other modules found in the system, including storage, memory, and I/O, without the complications of electrical signal integrity, crosstalk, EMI, and ESD immunity. In addition, simple optical pass-through modules could interface to the mid-plane to directly connect the server blades with external switches, storage, and memory.




Figure 16. Typical Blade Server Mid-plane Architecture

3. Multi-Chassis Switches and Routers

Many switching and routing applications require switching/routing capacities that exceed the capabilities of individual switching/routing ASICs. To extend the capacity beyond that of a single ASIC, a multistage “fabric” is configured between multiple ASICs; a typical example being a three-stage Clos fabric. Although modestly scalable systems can be achieved in a single chassis using an electrical interface between the ASICs, large-scale systems require multiple chassis. Typically, the system is comprised of (i) “line” chassis, which include the client I/O as well as the first/third stage of the fabric, (ii) “fabric” chassis, which include the second stage of the fabric, (iii) a high-bandwidth optical interconnect “trunks” between line and fabric chassis. This high-bandwidth interconnect is conveniently enabled by embedded parallel optics. Embedded optics preserve the signal integrity required by the high-end systems, simplify EMI containment whilst reducing ESD exposure, and enable flexible options for system cooling, by placement in close proximity to the host ASIC. Furthermore, as the trunks are typically 24 – 36 lanes wide, the embedded optics allow convenient bandwidth aggregation at the card edge for simple fiber management and chassis interconnect.




Figure 17. Multi-Chassis Routers

4. High Density Pipe for Uplink Connectivity

There are applications for high-bandwidth trunks that may not have symmetric data flow. For example, detector arrays in medical imaging and high-energy physics produce many gigabits per second of data. The detectors themselves are often located in challenging environments where long electrical interconnects are not possible. Embedded optics removes the need for complicated electrical signaling whilst overcoming the interference concerns in copper cabling. Another example is a collapsed backbone of high-bandwidth uplinks from sub-networks of an infrastructure to a single device that serves to monitor and control the network. These uplinks are often on space constrained front panels and require high density interfaces. Embedded optical modules conveniently located near the switch chip ASIC, in a location that does not consume precious front panel real estate, are an excellent solution to this challenge.

In Figure 18, five MiniPOD transmitter modules plus one receiver module are combined in a tile arrangement, utilizing flat ribbon fiber cables, to realize a 60x10 Gbps uplink and a 12x10 Gbps downlink.




Figure 18. MiniPOD Modules Interfaced Through a 72 fiber Circular MT Connector

Conclusion

As data communication systems move towards increased density and higher transmission speeds, several design limitations are being simultaneously encountered by systems engineering teams attempting to reduce component temperatures, comply with FCC emissions limits, and maintain extremely low bit error rates.

Avago is established as the world leader in the development of parallel optic technology and remains at the forefront of new product offerings, which now include a variety of embedded optic modules. There is a comprehensive range of commercially available complementary products for parallel optics including electrical and optical connectors, optical adapters, cables, and flex circuits available from multiple suppliers. Systems designers are selecting embedded optic parallel modules for several reasons including;

- Highest possible faceplate I/O density, limited by the physical size of the MTP connector.
- Excellent signal integrity, achieved by locating modules close to the host ASIC.
- Thermal management, provided by the flexibility to place modules at cooler system locations.
- EMI containment, since multiple high speed electrical channels are not routed to ports on the front panel.

Embedded optical modules have been extensively deployed in high performance computer systems, and in core routers and switches, where high density communications with excellent signal integrity are primary requirements. Avago MicroPOD and MiniPOD modules fulfill these requirements, whilst solving critical engineering challenges, and offer practical solutions to EMI, ESD, and thermal management.

For product information and a complete list of distributors, please go to our web site: www.avagotech.com

Avago, Avago Technologies, and the A logo are trademarks of Avago Technologies in the United States and other countries.
Data subject to change. Copyright © 2005-2011 Avago Technologies. All rights reserved.
AV02-2869EN - March 25, 2011

