



Parallel Optics Changes the Direction of Network Cabling

Parallel optics technology has been around for some time in high-performance network cabling, but until recently its role has been strictly behind the scenes. In the past two years, however, this technology has made a resurgence and is about to take center stage.

What is Parallel Optics?

Parallel Optics is a term used to represent both a type of optical communication and the devices on either end of the link that transmit and receive information. In traditional (serial) optical communication, a transceiver on each end of the link contains one transmitter and one receiver. The transmitter on End A communicates to the receiver on End B, sending a single stream of data over a single optical fiber. A separate fiber is connected between the transmitter on End B and the receiver on End A, so the link contains two fibers, commonly known as a duplex channel.

In parallel optical communication, the devices on either end of the link contain multiple transmitters and receivers, e.g. four transmitters on End A communicate with four receivers on End B, spreading a single stream of data over four optical fibers. With this configuration, a parallel optics transceiver can use four 2.5 Gb/s transmitters to send one 10 Gb/s signal from A to B. Essentially, parallel optical communication is using multiple paths to transmit a signal at a greater data rate than the individual electronics can support. Parallel transmission can either lower the cost of a given data rate (by using slower, less expensive optoelectronics) or enable data rates that are unattainable with traditional serial transmission.

Just as parallel transmission is fundamentally different than serial transmission, parallel optical devices are fundamentally different in construction than serial optical devices. Serial optical devices employ discrete components, such as transmitting optical subassemblies (TOSAs) or receiving optical subassemblies (ROSAs), and discrete optical connectors, such as the SC or LC connectors, which are almost always grouped together into a duplex pair. These discrete components are not suitable for parallel devices. Two complementary technologies have enabled the development and deployment of parallel optics devices: Vertical Cavity Surface Emission Lasers (VCSELs) and the MPO connector.

Development and Deployment

VCSELs are semiconductor laser diodes that are fabricated in two-dimensional arrays and emit energy perpendicular from the top surface, as opposed to conventional, edge-emitting

semiconductor lasers. VCSELs have a better formed optical output than most edge-emitting lasers, enabling them to couple that energy into optical fibers more efficiently. Also, since VCSELs emit from the top surface, they may be tested while they are part of a large production batch (wafer), before they are cut into individual devices. This in-process testing dramatically lowers the cost of the lasers, while array production makes them ideal for forming the linear arrays used in parallel optics.

The MPO is a standardized, multi-fiber optical connector capable of bringing an array of fibers together in a form factor similar to that of a single-fiber SC connector. While larger fiber-counts are possible, the most common application of the MPO is on 12-fiber optical cables, coupled with parallel optical devices utilizing the 12-fiber MPO similar to the module shown in Figure 1. Industry standardization provided parallel optics vendors with reliable and affordable connectivity with a linear array of fibers that matched their VCSEL arrays. In recent years, the MPO connector has gained wide market acceptance by enabling pre-terminated modular fiber systems for data center and backbone deployments.

Parallel optical devices emerged in 1999, and rapidly grew to encompass several packaging form factors available from multiple vendors. As this expansion and initial deployment of parallel optical devices coincided with the telecom bubble, the drive for more bandwidth led switch vendors, such as Cisco and Juniper, to deploy parallel optics in their carrier class telecom routers. The devices were used to interconnect line cards at very high speeds through an all-optical backplane. Cisco's CRS-1 router helped lead the activity, and still uses parallel optics today.

During this time, several application committees, such as Optical Internetworking Forum (OIF), Infiniband and Fibre Channel, anticipated the need for higher speed interconnects and wrote specifications for 10 Gb/s, 20 Gb/s and 40 Gb/s applications using parallel optics. When the massive bandwidth predictions for 2001 failed to materialize, parallel optics survived only in the internal workings of equipment from vendors like Cisco, Juniper and Alcatel, which were using the very short reach (VSR) protocol from OIF.

A Parallel Optics Resurgence

In the past few years, high bandwidth applications have become more commonly deployed in end-user sites such as data centers and high-speed computing labs. Market-leading companies are now shipping platforms that utilize parallel optical devices on the front of the chassis. Parallel optics-based Fibre Channel and Infiniband cards are now commercially available and being used for high-speed storage and server connectivity. In addition, a number of media-conversion applications using parallel optics have emerged. The parallel optical device has

come out from the backplane and is rapidly gaining acceptance as a useful interconnection technology.

Deployment of these devices in the data center consists of tying high-performance equipment together such as linking core routers together in a kind of external backplane. Cisco's 12000 series routers are capable of this kind of system-to-system link. Similarly, parallel optics may be utilized in a processing application, linking many high-performance servers together in a computing grid. Platforms such as HP's NonStop series and IBM's system p5 may use parallel optical devices to create high-performance computing (HPC) environments spanning multiple computing platforms.

Parallel Optics Connects with Infiniband

The use of Infiniband as an interconnection technology has increased dramatically in the past two years, driven by the protocol's higher bandwidth and lower latency. It is common for Infiniband switches to provide port latency as low as 0.2-0.4 μ s, with total link latency in the range of 1.0-2.0 μ s, compared with typical Ethernet link latency around 350 μ s. This interconnect technology is being used as both a networking and storage transport within the data center. While Infiniband is especially popular in cluster or grid computing, it is also on many standard rack-mounted and blade servers.

As Infiniband becomes more widely accepted, it is driving the use of parallel optics devices, particularly because traditional copper Infiniband connectivity has several disadvantages. The traditional copper cables, like the ones shown in Figure 2, are large (up to 0.45 inches in diameter, three times larger than the smallest MPO cable) and are notoriously difficult to handle due to their large bend radius. Moreover, these cables are not capable of being interconnected, meaning that all devices on an Infiniband network must be directly connected to each other. Therefore, it is impossible to implement structured cabling utilizing patch panels to enable system provisioning or administration. Finally, copper Infiniband has severe distance limitations, with typical maximum channel distances being less than 20 meters. Fortunately, the parallel optical versions of Infiniband have none of these shortcomings, rather using small 12-fiber cables that can span hundreds of meters and may easily be interconnected in a managed environment.

Two distinct versions of optical Infiniband (4x and 12x) have been adopted by the market. The 4x Infiniband defines a 2 x 4-lane parallel optical transceiver within a 12-fiber wide optical path. The full 12-fiber width of the MPO interface is consumed by four transmit lanes on the left and four receive lanes on the right. The 12x Infiniband defines two 12-lane parallel optical devices operating over distinct 12-fiber optical paths. By utilizing all 12 fibers, the 12x application

provides a higher data rate, by aggregating 12 ea. 2.5Gbps transmitter/receiver pairs for a total data rate of 30Gbps.

Media Converters Now Going the Distance

Perhaps the most unexpected application of parallel optics to emerge has been its use in media converters. A media converter is any device that converts a signal from one media (e.g. copper cable) to another (e.g. optical fiber) without modifying the underlying signal. Two applications for parallel optics devices as media converters have reached the market, and at least one shows great promise.

The first application addresses the drawbacks of traditional Infiniband cabling while supporting the installed base of copper Infiniband network cards. This small media converter, shown in Figure 3, has a pinned copper Infiniband connector on one end and an MPO receptacle on the other. The device is plugged into the copper port on a traditional Infiniband network card and converts the electrical signal into a parallel optical signal. Twelve-fiber MPO cables are then used to interconnect this first media converter to the second media converter, which is attached to the traditional Infiniband copper port on the destination network card. The electrical Infiniband signal is converted to an optical signal, transported across the fiber cabling and converted back to an electrical signal at the destination. Two media converters are required for each link.

This application of media conversion has several compelling advantages. First, it significantly extends the range of the Infiniband link, from 20 meters to 400 meters. While 400 may be overkill for most installations, the extension beyond 20 meters is a powerful driver for the use of these devices. Second, there is a significant reduction in cable volume by converting the rigid 12 millimeter- (0.45-inch) diameter Infiniband cable to a small, flexible 3 mm (0.11-inch) fiber cable, resulting in a 16X volume reduction. Finally, the device converts a specialized, multi-pin connector that cannot be interconnected into a standard fiber connector, eliminating the need for direct connection and allowing the link to be interconnected and cross-connected, a critical hurdle to managing the infrastructure via a structured cabling approach.

The second media converter application, as shown in Figure 4, converts 12ea. 1Gbps Ethernet signals from RJ-45 (copper) connectivity to two MPO fiber connections. Similar to the Infiniband converter, the fiber media converter provides both link extension and cable volume reduction. However, it provides extension where it is unnecessary. The device increases the link distance from 100 meters to 300 meters-plus, but 100 is an adequate distance for most applications, particularly in the data center. The extension would be useful for applications like remote IP cameras, but they require much less port density than the 4U, rack-mounted device provides. Converting the 12ea. 6.5 mm (0.255-inch-) diameter UTP cables to 2 ea. 3 mm (0.11-inch) fiber cables reduces cable volume more than 14 times, but this can be misleading. The media

converter requires that one fiber link be present for each copper link. If a small access switch were used instead, the 12 UTP channels would share just two duplex fiber cords. In this case, the media converter seems to waste fiber bandwidth and require more fiber cabling than needed.

Staying Power

The acceptance of parallel optical systems has been assisted by the development of a structured cabling system designed to support this traffic. Cabling systems designed to support serial duplex connectivity are not guaranteed to support parallel optics systems, as the polarity requirements of parallel systems are significantly different than those for duplex systems. In 2005, the Telecommunications Industry Association addressed the problem by publishing TIA-568B.1-AD7, providing guidance for implementing MPO connectivity in structured cabling. While this connectivity (shown in Figure 5) is primarily utilized today to support duplex connectivity, Method B within this standard was specifically engineered to enable the cable plant to support parallel optics applications as well with no change to the backbone cabling infrastructure.

While it has taken some time to realize the bandwidth predictions of the late 1990s, network speeds continue to increase, and soon 10Gbps will be commonplace, with 40Gbps and 100Gbps defining new high-speed applications. Most recently, in February of 2008, Cisco announced a data center class switch called the Nexus 7000, populated with up to 256ea. 10Gbps copper ports. The internal fabric of this switch, however, is geared toward 40Gbps and 100Gbps applications. The IEEE is currently developing 40Gbps and 100Gbps Ethernet specifications, with a projected release in November, 2009. Designated within IEEE-802.3ba task force, in the enterprise space these applications will depend on parallel optical transmission, fully supported by Method B. Over the next five to 10 years, as network speeds gain another zero, and Ethernet specifies the use of parallel optics technology, we can expect to see these devices become truly mainstream.

Author Simon Cowley, CommScope Enterprise Solutions Technical Director, North America