

Cabling

Installation & Maintenance

PoE

TO THE DESKTOP PAGE 15

STANDARDS PAGE 11

Field-terminating modular plugs

INSTALLATION PAGE 21

Making the case for microtrenching

DATA CENTER PAGE 25

Where OM5 fits

OM5: Hip or jive?

Introduced with fanfare, OM5 has seen few deployments. Why? And why was it created in the first place? What does the future hold for OM5?

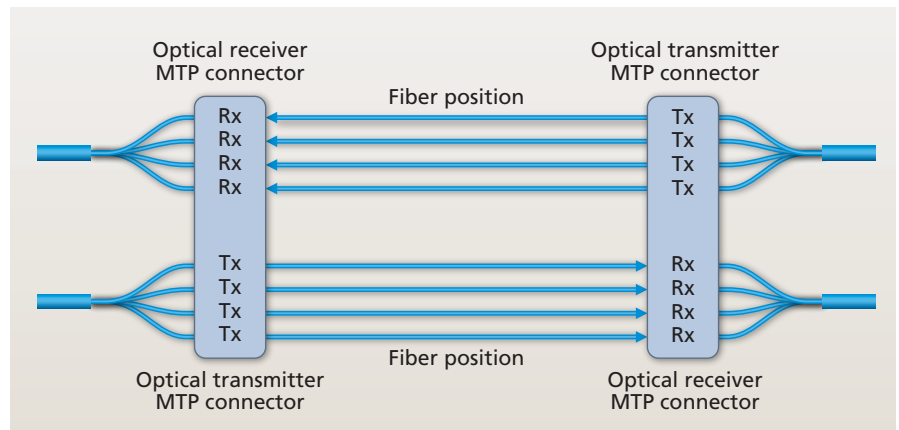
BY SCOTT GREGG, Corning Optical Communications

Multimode fiber (MMF) remains the dominant fiber type in local area networks (LANs) and data centers because it almost always results in the lowest link cost (defined as the cost of the fiber, connectivity, and optical transceivers) for short distances. But gone are the days of OM1 and OM2 enjoying most of the deployments; higher-bandwidth fibers such as OM3 and OM4 now rule the day. And recently, OM5 joined the party. To understand the circumstances around why OM5 was born, it's necessary to understand a few things about optical transceivers, and the standards which govern them.

First, we need to understand the distinction between standards-compliant transceivers, and proprietary transceivers. When we talk about standards-compliant transceivers in an Ethernet context, we're talking about optical transmit and receive guidance which has been ratified as part of an IEEE 802.3 Ethernet standard. When we talk about proprietary transceivers, these are transceivers whose guidance is not part of the IEEE standard, either because the proposed physical media dependent (PMD) technology did not get enough member votes to be included in the standard, or because the transceiver

utilizes a technology that was never offered up to become part of an open industry standard. The last few years have seen a proliferation of transceiver types available in the market, many of which are pro-

duce light at 850 nm, which was one of the multimode fiber specified operating wavelengths. Because the VCSELs were operating at 850 nm, subsequent enhancements in MMF fiber design and manufacturing were focused on optimizing the fiber bandwidth at 850 nm. For example, when OM4 fiber was introduced, it offered significant bandwidth improvements over OM3 at 850 nm, with OM4 offering 4700 MHz-km of effective modal bandwidth (EMB) at 850 nm, compared to the 2000 MHz-km provided by OM3.



40GBase-SR4 employs eight-fiber parallel transmission. In this setup, four fibers each transmit 10G, and a different four fibers each receive 10G. The MTP is the defined connector interface into the transceiver.

proprietary designs.

In the Ethernet realm from 1G up to 100G, all standards-compliant multimode transceivers have one thing in common: they utilize vertical-cavity surface-emitting lasers (VCSELs) that operate at the 850 nm wavelength. When VCSELs were first becoming commercially available, they were designed to

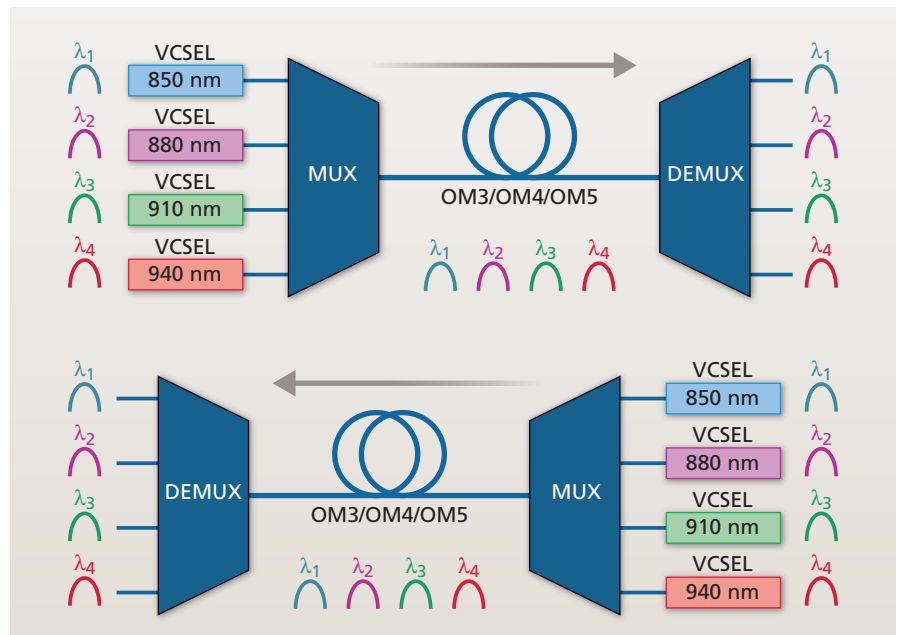
The second thing to know about multimode transceivers is the concept of parallel transmission, which some people refer to as parallel optics. For Ethernet speeds of 1G, 10G, and 25G, the multimode transceivers utilize two fibers, with one fiber carrying the transmit signal and one fiber carrying the receive signal. This is often known as serial transmission, and

OM5: Hip or jive?

since these are two-fiber devices, the connector interface into the transceiver is the LC duplex connector. However, with the adoption of the 40G 802.3ba Ethernet standard in 2010, the concept of parallel optics was introduced. In the case of the 40GBase-SR4 transceiver, we have four fibers in parallel, each transmitting 10G per fiber, and another four fibers receiving 10G per fiber. So these transceivers require eight fibers for a single channel, and as a result the multifiber MTP connector is the defined connector interface into the transceiver. A major feature of parallel optics transceivers such as the 40GBase-SR4 is that because individual fibers are each carrying a 10G signal, a single 40G MTP port on a switch can be broken out to four LC duplex 10GBase-SR ports, which typically results in significant per-port power cost savings and higher switch port density. With this type of breakout, a line card with 32x40G ports can be broken out to 128x10G channels. For network managers who need port breakout functionality and need 40G distances beyond 150 meters, a proprietary extended-reach transceiver, the eSR4, also has been introduced.

Enter SWDM

As we said, the last few years have seen a large number of proprietary transceiver types come onto the market, starting with the 40G BiDi transceiver. The BiDi transceiver is a two-fiber device, with bidirectional transmission over each fiber; each fiber carries both a transmit and receive signal, operating at different wavelengths (850 nm and 900 nm). Because the BiDi transceiver requires only two fibers, it was designed to provide a migration path up to 40G where OM3 or OM4 duplex fiber connectivity was already installed, such that additional MTP connectivity would not be required. The BiDi transceivers have proven to be a good solution for 40G switch uplinks. Importantly, because they operate by having each fiber carry both a



A two-fiber, 40G short wave division multiplexed (SWDM) transmission system operates over four wavelengths per fiber across the range from 850 to 940 nm. One fiber is dedicated for transmit and the other is dedicated for receive.

transmit and receive signal, port breakout functionality is not supported.

Into the fray we've had another transceiver transmission technology join the fight—the short wavelength division multiplexing (SWDM) transceiver. Similar to BiDi in that for a 40G circuit an SWDM transceiver only requires a two-fiber LC duplex connection, SWDM differs in that it operates over four wavelengths per fiber across the range from 850 to 940 nm, with one fiber dedicated for transmit and one fiber dedicated for receive.

As with BiDi, the SWDM transceiver is designed to give network managers who have an installed base of OM3/OM4 duplex connectivity another way to migrate to 40G without having to deploy additional fiber. However, the fact that there are four transmission wavelengths presents an interesting question for the industry: Given that OM3/OM4 fiber bandwidth is typically only specified at 850 nm, how to quantify the peak performance of this transceiver with an operating wavelength up to 940 nm? The answer: The

Telecommunications Industry Association (TIA) initiated a working group in 2014 to develop guidance for so-called “wideband multimode fiber (WBMMF)” to support SWDM transmission. The TIA-492AAAE standard for WBMMF was published in June 2016. WBMMF is effectively a type of OM4 fiber, as the WBMMF still has to meet the OM4 bandwidth criteria of EMB ≥ 4700 MHz·km at 850 nm, however the WBMMF also has bandwidth specified at 953 nm. The EMB specification at 953 nm is ≥ 2470 MHz·km. An international vote in October 2016 gave WBMMF a three-digit designation, and OM5 fiber was born.

Because OM5 fiber exists and is being priced at a premium over OM4, it must provide some benefit, right? Well, yes, otherwise there would have been no reason for the industry to create a standard for it. This is where we need to consider transmission distance capability for the various MMF/transceiver combinations. Even though 100G is mainly the purview of hyperscale data center operators, given that we know higher speeds are coming

TABLE 1: 40G transceiver summary

	Transceiver Type			
	40GBase-SR4	40G eSR4	40G BiDi	40G SWDM
IEEE standard compliance	Yes	No	No	No
Number of fibers required	8	8	2	2
Connector interface	MTP	MTP	LC duplex	LC duplex
Number of wavelengths	1	1	2	4
Port breakout capability	Yes	Yes	No	No

TABLE 2: OM3/OM4/OM5 fiber bandwidth summary

	Fiber type		
	OM3	OM4	OM5
Specified bandwidth at 850 nm	≥2000 MHz·km	≥4700 MHz·km	≥4700 MHz·km
Specified bandwidth at 953 nm	Not specified	Not specified	≥2470 MHz·km

to enterprises in the future, let’s evaluate the distance capabilities of both 40G and 100G, based on published transceiver-manufacturer specifications with standard connectivity.

- In Table 3, we observe the following.
- Using either the SR4 or eSR4 transceivers, because these transceivers operate solely at 850 nm, there are distance benefits for OM4 over OM3, but no distance benefit for OM5 over OM4, because both OM4 and OM5 meet the same bandwidth spec at 850 nm.
- At 40G, both BiDi and SWDM have a distance benefit for OM5 over OM4, given that these are multiple-wavelength transceivers. However, the OM4 distances of 150 meters for BiDi and 350 meters for SWDM are already quite far, and are sufficient for the vast majority of MMF applications. For example, published industry data reveals that in the data center, 90 to 95 percent of OM3/OM4 links are 100 meters or less.
- At 100G, an OM5 distance benefit exists for both the BiDi and SWDM transceivers, as OM5 provides up to 150 meters of reach, as compared to the 100 meter reach provided by OM4. The longest overall reach of 300 meters

is provided by the eSR4 transceiver with either OM4 or OM5 fiber.

Where OM5 fits

So given all this, where is the right place to use OM5? It depends. To make this decision we need to understand a number of

TABLE 3: Transmission distance (in meters) per fiber type and transceiver type

Fiber type	40G transceivers				100G transceivers			
	40 GBase-SR4	eSR4	BiDi	SWDM	100GBase-SR4	eSR4	BiDi	SWDM
OM3	100	300*	100	240	70	200	70	75
OM4	150	400*	150*	350	100	300	100	100
OM5	150	400*	200	440	100	300	150	150

Note: Distances represent guidance published by the transceiver manufacturers; some switch vendors could provide different guidance.

**Longer supported distances are possible, using some connectivity solutions available on the market.*

factors related to network speed, required transmission distance, and the transceiver technology being used.

For example, if you intend to use standards-compliant transceivers, you will be using an SR4-type transceiver and OM5 provides no value over OM4. Or if you know that you will need port breakout capability, then you will be using an SR4 or eSR4-type transceiver, and again OM5 provides no value over OM4.

If you intend to use BiDi or SWDM transceivers, then the network speed and the required transmission distance become deciding factors. Again, most network managers will not have many, if any,

links beyond 150 meters. So OM4 will accommodate most 40G needs, and the 440 meter reach possible with the OM5/SWDM combination is of value to few users. However, if you plan to migrate to 100G, and will have a significant number of links beyond 100 meters, then OM5 does have a use case, as it provides an additional 50 meters of reach over OM4. Given that few network managers have MMF links beyond 100 meters, and there have been extremely few deployments of 100G in enterprise LAN or data center networks, this explains the very slow adoption of OM5 so far.

When 100G enterprise deployments become more prevalent, OM5 could become attractive if a reach up to 150 meters is needed. Deployments of OM5 provide some value for network managers who will deploy 100G using BiDi or SWDM transceivers and have links be-

tween 100 and 150 meters. OM5 currently is not included as a physical media option within any of the published Ethernet or Fibre Channel standards, however if SWDM transceiver types are adopted within Ethernet or Fibre Channel in the future, it would be logical at that time for OM5 fiber to be included as an available option in the standards, along with OM3/OM4 fiber. One thing is certain: The lowest link cost for the distance needed will win. ♦

Scott Gregg is marketing director for Asia Pacific with Corning Optical Communications (www.corning.com).