TechTarget SearchTelecom

E-guide

PRO+ Content

Optical Network Design and Transport

In this e-guide

Best practices for optical network design

Fiber-optic technology -- not long ago used only in long-haul networks -- has become the transmission medium of choice not only in the core, but in metro and access networks. The game-changer is global growth in consumer broadband and the need to distribute enormous amounts of content without hauling it halfway around the world. This Telecom Insights guide to best practices for optical network design looks at access, metro and core network issues affecting fiber deployment.

Fiber-optic technology -- not long ago used only in long-haul networks -- has become the transmission medium of choice not only in the core, but in metro and access networks. The game-changer is global growth in consumer broadband and the need to distribute enormous amounts of content without hauling it halfway around the world.

This Telecom Insights guide to best practices for optical network design looks at access, metro and core network issues affecting fiber deployment including:

- Recent developments in FTTX ("fiber to the whatever") deployment and passive optical networking (PON) technology in access networks
- Factors affecting metro network use of SONET, WDM or Ethernet, and the services that will be supported by fiber deployment
- Specific core optical network design considerations including aggregation, geography and reconfigurability.

Fiber-optic networks: Access network design

The rapid growth in consumer broadband seen worldwide today would not be possible without a major shift in the practices for provisioning access infrastructure. Copper loop and CATV cable were once the only means of transporting information from a provider central office or head end to the customer. Today, both these media are being "shortened" or even eliminated by the use of fiber optics.

Fiber is not a new development in access networks. Not only has it been used for almost two decades in the provisioning of high-speed commercial/enterprise customers, service providers in the 1990s found that replacing large bundles of copper by a few fiber strands could improve service reliability and lower craft cost. BellSouth took the lead in deployment of access fiber in that period, and the move was justified completely on cost savings.

The traditional access fiber architecture has been the **fiber remote**, which is a high-speed fiber trunk (SONET or Ethernet) that terminates in an electro-optical multiplexer. In analog phone days, these were called "digital loop carriers" (DLCs), and the term "new generation DLC" was used for a time, but most such devices today deliver DSL services and so are usually called "remote DSLAMs." A remote DSLAM's primary benefit is to shorten the access copper to allow higher DSL speeds and improve reliability. Most providers would counsel against offering premium DSL on loops over 8,000 feet, and the highest DSL speeds may be achievable only on loops 1,000 feet or less in length.

Pushing fiber close to the customer is generically called "deep fiber," and various acronyms are used to indicate just how deep the fiber is. FTTH means "fiber to the home," which is the extreme of giving every user an optical-electrical termination. FTTC takes "fiber to the curb," serving a group of homes, while FTTN means "fiber to the node" or "neighborhood," and allows each fiber remote to serve a larger population.

The problem with all deep fiber strategies, and the reason why providers don't simply run fiber to every home, is cost. If loops are kept to a length of 5,000 feet, a single remote can serve customers in an area of almost 2,000 acres. Shorten the loop to 1,000 feet and it serves only a little over 70 acres.

Since the user population is generally proportional to the geography, this reduction means the cost per user could rise 50 times or more. Shorter loops mean higher speeds, however, and for video over IP, most operators would require at least 24 Mbps (ADSL2) connections. In Asia and some other areas, VDSL is used with speeds of 50 Mbps or more. Both these require much shorter loops (8,000 feet is optimum for ADSL2, according to reports, and 500-800 feet for 50+Mbps VDSL).

Balancing cost and performance is the goal of the various **passive optical networking** (PON) systems. PON creates a "tree" structure of fiber connections using optical splices without electrical termination or handling. PON typically supports 32 branches, and each of these can in theory be a remote or a home. A single PON tree supporting 32 branches has 33 electrical devices, counting the head end. Serving 32 locations with point-to-point fiber would require 64 electrical devices and generate higher costs and greater reliability risk.

PON systems use a common fiber architecture but a variety of opto-electric approaches. The original broadband PON (BPON) and the successor Gigabit PON (GPON) are both based on ATM. The new Ethernet PON (EPON) standard has been ratified, and most operators contemplating major new PON deployments are conducting assessments and procurements of EPON. GPON and EPON have sufficient capacity for video delivery and high-speed Internet. Some providers like the ATM framework of GPON for its ability to create multiple independent service channels to the user via virtual circuits. Others prefer EPON because it matches better with Ethernet-based metro architectures.

Planning for access network fiber deployment demands a careful consideration of the following:

1. The demographics of the area to be served, including household income, family size, and age distribution. This data is critical in establishing the service market opportunity. In general, favorable demographics justify deeper fiber deployment.

2. The geography and topology of the service area, including the household density (average lot size), the rights of way available, and whether cabling is underground or above ground. This data is critical to set the cost points for each approach. Obviously, poor characteristics here will create profit margin challenges if not taken into account. Studies in Japan, where fiber deployment is high, indicate that even whether the ground is flat or hilly has an impact on deployment cost.

3. The service mix to be provided, over at least a five-year period, considering both trends in demand and in competition. The worst possible outcome in an access fiber deployment is a new set of

requirements that the fiber architecture deployed cannot effectively support.

In the installation and maintenance phase, access networks present special problems because of the high cost of rolling a truck to fix a problem. A broadband consumer may require three years to pay back the cost of a single service call. This means that it is absolutely critical that each fiber strand be properly installed and that, in particular, the splicing used in PON installations be carefully done and verified. Fiber should also be tested end-to-end prior to committing it to customers. Unlike copper, whose problems tend to develop over time, operators report that most fiber problems are uncovered shortly after installation and result from improper practices.

Optical networks: Metro network design best practices

The trends in telecommunications today show clearly that the largest incremental amount of fiber deployed in the next decade will not be in the network core but in the access network and metro network. Content, the fuel of consumer broadband traffic growth, is an application that delivers a relatively small number of movies or programs to a large population of users. In most cases, this means that content will be cached at a metro level and that the greatest traffic growth created by content will be in the metro area.

Metro fiber today is based largely on SONET, which is 1310nm single-wavelength deployment. SONET networks are usually constructed as a series of protected rings that allow fast failover to the alternate "rotation" in the event of a fiber cut. Rings are connected via optical add/drop multiplexers (ADMs).

The advent of wavelength division multiplexing (WDM) -- coarse or dense -- deployed in the 1550nm range has added versatility to metro optics by providing multiple lightpaths per fiber and greatly increasing the capacity of a given fiber strand. At the same time, the increased volume of packet traffic, which does not require SONET's synchronous delivery behavior, has changed the traffic profile for the metro network of the future. Today, Ethernet is more likely to be the planned electrical layer of metro networks, and WDM the optical. This shift is changing the balance of tasks between electrical and optical components and the best practices for deployment.

SONET rings can be replicated in metro Ethernet and dense wavelength division multiplexing (DWDM) networks by simply using the same fiber and relying on wavelength separation or by running multiple SONET paths over WDM. Since there are probably no major metro networks worldwide without any traditional synchronous TDM traffic, planners should expect to use a hybrid of SONET and Ethernet technology. Where there is a large installed base of SONET equipment, no plans to eliminate PSTN switches, and major customers with direct SONET access, it may be advisable to plan a

In this e-guide

transition in the metro optical network from SONET-over-1310 to SONET/WDM and then to begin to integrate Ethernet-over-SONET, finally moving portions of the network to Ethernet-over-WDM.

The changing economics of WDM appear to be defusing the "SONET replacement" issue. Most operators now expect to maintain SONET for PSTN transport for as long as those services are offered, moving to non-SONET architectures only as packet voice displaces TDM voice. However, the gradual evolution is most likely to be compromised by exploding consumer broadband use, particularly by IPTV plans. Operators report voice traffic is stable while data traffic is growing at often triple-digit rates. The faster packet traffic grows relative to "circuit" or TDM traffic, the more likely it is that hybrids of SONET and Ethernet (Ethernet over SONET) will have too small a window of value to justify investment. This is probably the reason why more and more optical vendors are offering hybrid products with reconfigurable add-drop multiplexing (ROADM) and Ethernet.

A WDM issue that receives considerable media play is the way in which transit optical connections are handled. Most products have converted between optical and electrical (O-E-O) to perform a wavelength transit connection because pure optical cross-connect (O-O-O) has been expensive. While O-O-O products have been available for five years or so, most vendors still use O-E-O technology.

In this e-guide

The primary issue today is still cost; service providers believe that future ROADM products will provide all-optical transit connections. The key issue for designers, regardless of the mechanism used, is that wavelengths can be "transcoded" to a different wavelength across the switch; a system that fails to provide this is too complex to manage because wavelength assignment on various fibers becomes interdependent, and some reconfiguration modes may not be available because of collisions.

Metro optical deployment is affected by the service mix to be supported, but the service topology has an equal or greater impact. A primary question to be addressed is the amount of intra-metro traffic to be carried relative to the volume of traffic that will simply be connected to a metroPOP for transport outside the metro area.

In areas where consumer broadband traffic makes up the bulk of total traffic, most fiber deployment will focus on linking serving offices to a POP for core network interconnect. While these connections have to provide resiliency, they will rarely require the SONET standard of failover, 50ms, since they will support Ethernet traffic. The introduction of IPTV may change this picture because loss of connectivity will cause pixelization that may produce viewer dissatisfaction, particularly on pay-per-view systems. If user buffering is available, the failover time should be no more than about two-thirds of the buffer interval. Often, consumer broadband failover is best accomplished at the Ethernet level.

Where there is significant synchronous (TDM) traffic and significant corporate packet traffic, it may be necessary to provide optical failover at SONET 50ms levels, in which case SONET or resilient packet ring (RPR) may be required. As noted, WDM may allow metro optical designers to separate traffic according to optical failover requirements and provide improved failover only where needed.

Reconfigurability, meaning the ability to create variable metro optical topologies by interconnecting wavelengths in various ways, is most likely to be needed either to accommodate a large amount of business traffic (metro Ethernet services) or to support alternate routing between serving offices and metroPOPs where the core network connection is made. Where IPTV is delivered, this multi-homing may also be needed for content service points.

At the optical layer, reconfigurability and fast failover are very different things. ROADMs offer a great amount of topology flexibility to adapt to changes in traffic demands, to the point where wavelength services can be offered to metro customers and where even Gigabit Ethernet customers can be quickly accommodated. Adding rapid and multiple spanning trees to Ethernet can provide resiliency at the electrical layer for everything but the most stringent failover requirements.

Many believe that metro optics will, over time, migrate away from the 50ms failover standard of SONET as circuit-switched and TDM traffic become a smaller portion of network load. If this is true, then a pure ROADM-and-

Ethernet solution, particularly one based on one-box optical/electrical approaches, may be the best long-term solution.

Optical networks: Core network design best practices

Optical networking is the only relevant Layer 1 technology today in the network's core except in very unusual markets or geographic conditions where terrestrial microwave may still be deployed. While core optical network deployments in some areas may be literally indistinguishable from metro fiber, there are other networks where the key requirements are totally different. Thus, the first question in optical network design and deployment for core networks is the nature of the network itself, and how core and metro requirements might differ.

The "core" of a network is a place where aggregated traffic moves among on- and off-ramps. Because the traffic is highly aggregated and thus represents thousands or millions of user relationships, core network nodes are likely to have traffic destined for virtually all other core nodes, meaning that the nodes are highly interconnected. This may contrast sharply with metro networks, where "preferred topologies" often involve simply connecting serving or edge offices with points of presence (POPs) for connection to the core -- a star topology instead of a mesh.

The aggregation of traffic onto the core creates another difference in core optical networks: It is unlikely that a single network user will contribute a large portion of total traffic, and thus adding new users will generally not change the core significantly. In contrast, a large metro user may require reconfiguration of network bandwidth to accommodate the traffic. For this reason, and because of the "mesh" factor above, reconfigurability is likely to be less of an issue in core networks.

Core networks are also typically immune from the need for fast failover, the 50 ms optical alternate routing available with SONET. Ring configurations using fully redundant fiber paths are harder to create and more expensive to maintain in core networks, and so resilience is typically left to the electrical layer.

The final difference is that of geographic scope. A large metro network might span 50 to 100 km; a large core optical network can circle the globe. This long reach necessarily means that core fiber may have to span great distances without intermediate repeaters, including submarine environments, deserts, etc. Thus, ultra-long-haul fiber technology is often critical in core networks. The greater geographic scope of the core network also means getting craft personnel to an area to fix a problem may require days or weeks, and so it is critical to have some form of backup plan and to reduce outages as much as possible through design.

One issue that core and metro networks share is the issue of synchronous or circuit-switched traffic. Where PSTN calls and T1/E1 lines are to be supported over the core, it will likely be essential to utilize SONET/SDH transport for at least some of the optical paths to provide for synchronous end-to-end delivery. SONET/SDH services over global distances also require very accurate clocking to insure that bit errors are not created through "clock slips." These SONET/SDH trunks can either use the standard 1310 nm wavelength or one of the 1550 nm WDM wavelengths. Packet traffic does not require SONET/SDH, but many core network operators continue to use some SONET/SDH ADMs and switching in the core to preserve the option of circuit switched services.

A "pure packet core" can be made up of single-wavelength or WDM fiber, and thus it may be possible to create a virtual optical topology that approaches a mesh to avoid electrical handling. However, routers often have "adjacency problems" when installed in a full physical mesh, creating very long convergence times in the case of a failure. This, combined with the fact that reconfigurability in the core is often not a major requirement, means that core networks are more likely to use very high-speed fiber paths (OC-768, or 40 Gbps, for example) if the economies of these single electrical interfaces are better than the sum of the cost of WDM and a larger number of slower interfaces (4x10Gbps Ethernet).

The router adjacency issue is an example of an important point in core fiber design, which is that the needs of the electrical layer and even the service goals must be considered. Current trends in service provider Ethernet, spearheaded by work in the IEEE and the Metro Ethernet Forum are making Ethernet a strong candidate for core network deployment both to provide flexible virtual routes for higher-layer protocols like IP and to serve as the basis for actual customer services. This approach allows operators to create meshed optical networks for resiliency and add packet routing and even multicasting without creating additional router adjacencies.

A major consideration in optical core networking is the location of major points of service interconnection. The larger a provider core network is in terms of geographic scope, the more likely it will interconnect with networks of other providers, especially for local access in other geographies. These interconnection points are obviously both major traffic points requiring special capacity planning and points of major vulnerability. No interconnection point with another operator should be single-homed in fiber connection, nor of course should metro connections with the core provider's own metro infrastructure be single-homed.

The final point in core optical design is the management framework. Core networks carry aggregated traffic from millions of users, and failures will result in a flood of customer complaints. In addition, optical failures will trigger an avalanche of faults at the higher protocol layers, generating so many alerts that the network operations personnel may be overwhelmed. Many operators have insufficient integration between packet and optical layer management, and this increases vulnerability to alert storms and also makes customer support personnel less likely to have ready answers to complaints. The best optical core is no better than the operator's ability to manage it properly.

Next article

In this e-guide

100G DWDM optical networking transport

The telecom industry wants to move quickly to develop 100G DWDM optical network transport. In part 1 of this expert lesson, optical networking expert Eve Griliches looks at the forces driving first-generation adoption. After learning from 40G mistakes, standards organizations want to make sure interoperable high-speed optical transmission products get to market quickly at the right price point for easy deployment. She goes on to provide a short history of 100G DWDM development and looks at the technologies that will enable 100G DWDM. As a bonus, Griliches offers her projections on future 100G DWDM deployments.

Telecom industry prepares for 100G DWDM optical network transport

With the telecom industry giving a resounding "yes" to 100 gigabits per second (100G) dense wavelength division multiplexing (DWDM) optical network transport, the race is on to develop and deliver first-generation 100G DWDM products that will boost optical channel rates to 100G. Having learned from the multimodulation mistakes of 40G, the industry has quickly moved to standardize components so vendors can bring products to market faster and at cost points that will make 100G adoption viable in a shorter time frame.

The first implementations for 100G DWDM network transport will be high-capacity connections between switches, switches and routers, and router-to-router at Internet exchanges and within service provider and carrier points of presence (POPs).Why? One reason is that innovative content providers with huge data centers have put pressure on the industry to develop 100 Gigabit Ethernet on the client side. In addition, the aggregation of 10G links from IP routers has driven the line-side development of 100G DWDM for long-haul optical transmission on service provider and carrier networks.

These innovations should lead to extensive deployment of 100 Gigabit Ethernet in data center networks, which will be required to manage the bandwidth necessary to aggregate the server and storage community within the data center. 100G optical transport links will also be deployed as a requirement for higher bandwidth applications and to aggregate internal service elements within the carrier.

Capacity-constrained data centers and long-haul transmission routes will ultimately benefit from the tenfold expansion of Ethernet to 100G. Ethernet data rates have always increased by a factor of 10, while traditional SONET/SDH and optical data rates have increased by four, keeping

Ethernet and transport on separate migration paths until they converged at 10 GB/s. This convergence brings up the question of 40G Ethernet and DWDM equality and puts 100G DWDM squarely in the sweet spot for flexibility, cost points and ease of deployment.

Yet 100G transport is not a panacea. There are challenges, such as its susceptibility to transmission problems and the need to meet cost points for wide deployment once it has been tested and deployed. Most likely, it will take second- and possibly third-generation products hitting the right cost points to move providers from 10G and 40G to 100G.

System requirements for 100G DWDM network transport: 100G DWDM and Ethernet standards groups

The following groups are working on relevant 100G DWDM optical network transport (optical channel rates) and Gigabit Ethernet (GbE) standards:

- **IEEE 802.3 High Speed Study Group** (HSSG): Client-side-interface focus for up to 40 km transmission. It supports full-duplex operations and bit error rate (BER) better than or equal to 10 (to the minus 12). With different PHY requirements, a host of physical solutions for 100 Gigabit Ethernet is necessary to support data center interconnect as well as interoffice carrier connectivity. Support for 10 copper, 100m MMF, 10km SMF and 40km SMF are currently proposed.

- **ITU Study Group 15; Next Generation Optical and Transport Networks:** Focus on metro and backbone DWDM interfaces. Proposals for the next higher bit rate beyond 43G (OTU-4) have been submitted for three times ODU-3 at 130G, as well as a data rate optimized for 100G at approximately 112G.

- **OIF-CEI 25G Next Generation Backplanes**: Defines common electrical interfaces for high-speed applications.

- **ITU G.709 OTN Standard:** Defines basic bit rates, multiplexing and sub-lambda capabilities, as well as how signals are mapped and assigned overhead within transport networks.

- **OIF DP-QPSK Standard**: Defines the modulation scheme for 100G DWDM to speed integrated optical components to market.

While high-speed links are needed in wide area network (WAN) *and* local area networks (LANs), the requirements vary for each topology.

For the LAN, data center requirements are typically Ethernet-based for 1G, 10G, 40G or 100G Ethernet. In this environment, multivendor interconnectivity is so prevalent that standards are key factors and must be supported in order to interoperate. Often, the interconnectivity is over a single fiber (in most cases, multimode) where spectral efficiency is not important. It is here that the IEEE is working on 802.3 Ethernet LAN standards.

For the WAN and backbone regional networks, spectral efficiency -- a measure of how efficiently a limited frequency spectrum is utilized by the physical layer -- is extremely important. Typically, the International Telecommunications Union (ITU) standard for ITU G.709 is used for the serial rate. The interface is often the proprietary piece. This is where vendors often differ in approach and promote their advantages, which center in the modulation formats that are chosen, the way dispersion is mitigated on the fiber, polarization mode dispersion (PMD) mitigation, and the implemented forward error correction (FEC) technology.

So, as higher-speed 100G DWDM network transmission rates are deployed, the following basic requirements must be met by vendors if they want to be considered by service providers:

- Equipment must be compatible with the existing infrastructure that supports 10G DWDM deployment. Providers are asking vendors to make sure they support the existing regeneration infrastructure in distance and amp length for 10G/40G and 100G. This means equipment *must* be tolerant to current PMD and nonlinearity and meet the existing engineering rules established for 10G.

- Equipment must support 50 GHz channel spacing. Since most providers' grids are deployed on 50 GHz spacing, it would be spectrally inefficient to deploy any system on 100 GHz or anything other than 50 GHz.
- Systems must have a large dispersion tolerance and *must be* easily deployed and managed. Equipment cannot require tweaking and be difficult when adding lambdas.
- Products must be cost effective. The target for third-generation 100G equipment is 6.5 x 10G, or 65G of bandwidth at 10G pricing. This goal will most likely not be achieved until the industry deploys third-generation equipment. The estimates are that the price of early 100G deployments will exceed 10 times that of 10G, and second-generation products are expected to top the cost of two times 40G. Clearly, if the price points are there, providers will move faster to consolidate their networks to 100G.
- 100G DWDM will be deployed first in core and long-haul networks, where distance requirements run from 600 km to 2,000 km and support for fewer than five or six ROADMs (optical filtering) is required. These networks are less cost sensitive than regional or metro networks, where the distance is less challenging and the number of ROADMs increases significantly.

What becomes apparent is that transmission challenges increase with higher bit rates; 40G and 100G DWDM optical networking transport now require advanced modulation and coherent detection techniques to meet network operators' requirements. Most 40G networks use differential phase-shift keying (DPSK) or differential quaternary phase-shift keying (DQPSK) modulation schemes for long-haul networks. For 100G, the expectation is that DP-QPSK will be the standard modulation format used, while individual coherent detection formats will vary from vendor to vendor.

Next article

A short history of 100G DWDM optical network transport development

As long-haul dense wave division multiplexing (DWDM) growth is accelerating, interest in 100G DWDM optical network transport continues to surge, even as operators deploy 40G DWDM systems to increase optical channel rates. To discover the reason for the rush, review the history of increasing optical channel rates and the challenges facing trials and deployment and production networks.

It comes as no surprise that traffic growth is the underlying driver for 100Gigabit (Gbit) Dense Wave Division Multiplexing (DWDM) network transport. This was the case with the transition from 2.5G to 10G and from 10G to 40G and now 100G optical channel rates.

Telecom carriers and independent vendor assessment reports estimate 50% to 60% traffic growth year over year. As the number of Internet users expands, bandwidth per user is increasing, putting enormous pressure on metro networks and core backbone networks.

Routers with 10G and 40G interfaces emerged to hand off to DWDM transmission equipment, driving the need for 40G wavelengths. Router handoffs will ultimately push the need for 100G as access streams into routers become higher data rates (10G and higher) that force the transport backbone link to go to an even higher rate.

Changes in network architecture are another driver for higher-bit-rate networks. In the past, providers had separate networks for each access technology or service. Today, most providers are attempting to put these services and technologies onto one IP backbone, which puts pressure on the core backbone to handle much more traffic. When AT&T merged multiple networks (SBC and BellSouth), for example, it led to integrated networks overflowing onto one large backbone that required an immediate upgrade to 40G and is now 100G capable.

Consolidating traffic onto fewer wavelengths and the associated economics were other factors for the transition. Larger bandwidth wavelengths have always promised better operation efficiencies merely because there are fewer wavelengths to manage and a smaller number of parts in the network that can fail.

In the drive toward faster optical channel rates, some major providers have already moved to 40G networks:

- Comcast has one of the largest IP-over-WDM 40G deployments for its national network, deployed with Cisco CRS-1 integrated optics and Nortel transmission equipment.

- In Japan, NTT is offering a 40G network service over a wide-area Ethernet network.
- Verizon has deployed live traffic on 40G links on high-bandwidth routes with Alcatel-Lucent equipment.
- Deutsche Telekom (DT) also has deployed a 40G network on equipment from Ericsson (Marconi).
- Sprint Nextel started upgrading its 10G network to 40G using Cisco's IP over WDM and Ciena's CoreStream.
- In October 2008, AT&T completed its IP/MPLS backbone network with more than 80,000 fiber optic wavelength miles of 40G running on Nokia Siemens' ULH platform. According to AT&T, every 10G lambda will be a 100G lambda by 2012 because of 60% traffic growth year over year.

100G DWDM network transport performance trials and tribulations

For the last two years, various 100G DWDM network transport technology experiments have come to market, all with different distances, alternative formats and widely disparate margin allocations. It is technically feasible to light one wavelength on great fiber and go a few hundred kilometers. It is a totally different proposition, however, to have preproduction equipment

In this e-guide

running thousands of kilometers with multiple 100G DWDM wavelengths while assuming system operating loss margins.

Vendors must take into account optical signal-to-noise ratio (OSNR) after transmission (the loss incurred), nonlinear transmission penalty, polarization mode dispersion (PMD) and filtering (ROADM) penalty, as well as aging and end-of-life margins. Real-world performances in the field versus hero experiments take into account these penalties before announcing performance metrics.

Some companies are conducting performance metrics. Comcast executed a realistic trial on the same fiber as 10G and 40G wavelengths for 100G with dual-polarization quadrature phase-shift keying (DP-QPSK) on a 300 km span with live traffic. Verizon has tested 100G with Nokia Siemens Networks (NSN) over a 1,000 km in the lab and has also tested Nortel 100G on very high PMD fiber with 10G and 40G wavelengths running on the same fiber. AT&T has also tested NEC over 600 km. The outcome was that these products, which were prototype hardware, met challenging yet forgiving situations.

Next article

Enabling technologies for 100G DWDM network transmission

As 100G DWDM optical network transport takes hold to speed optical channel rates to accommodate traffic growth, DWDM technology vendors, telecom service providers and standards groups have to make hard choices in terms of modulation formats. Find out why the industry is favoring multilevel phase modulation to avoid the problems it encountered with 40G DWDM development.

Previous long-haul transmission systems supported internal modulation and used non-return-to-zero (NRZ) modulation formats. Both 40G and 100G optical channel rates required external modulation formats, of which the most common are amplitude, frequency and phase modulation, which is the most likely to be deployed.

A multilevel phase modulation is most interesting since it has a small sensitivity penalty, provides constant intensity and is relatively simple to generate. Examples of multilevel phase modulation are differential phase-shift keying(DPSK) and quadrature phase-shift keying (QPSK), which has four phases, two bits per symbol at about 25 GBaud rate per symbol.

In this e-guide

Depending on the modulation format, a provider can modulate intensity, phase, frequency and polarization or any combination thereof. Each modulation scheme affects the capability to support 10G, 40G and 100G on the same fiber and same infrastructure, making decisions on modulation format for support critical. Each format has its own complexity and cost; optical signal-to-noise ratio (OSNR) sensitivity; spectral efficiency; tolerance to fiber nonlinearity; and tolerance to chromatic dispersion and polarization mode dispersion (PMD).

The Optical Internetworking Forum (OIF) is working on standardizing the modulation format for 100G at dual-polarization quadrature phase-shift keying (DP-QPSK). Dual polarization means there are two 50G optical signals; one goes up and down, the other side by side. Both signals are on the same frequency, but one is horizontal, while the other is vertical. The signals are thus polarized at 90-degree angles so that they never interact with each other. QPSK provides four phase states in the receiver, which drops the signals down to 25 GBaud signals. While DPSK is certainly a fine modulation technique, if PMD is an issue on the fiber, providers will probably deploy DQPSK where the PMD tolerance is better.

In addition, 100G performance will be improved by adding a coherent receiver, which eliminates the need for dispersion compensation at each amp site and development comes within a digital signal processor ASIC. With coherent detection, the chromatic dispersion and PMD mitigation are done electronically within the chip at the receiver end, cleaning up the signal. This requires high-speed analog-to-digital conversion followed by advanced digital signal processing (DSP).

To date, this improvement has been exciting to the optical community, as the packaging is typical CMOS, which should be a better and more cost-optimized technology that separates dynamic dispersion compensation methods. Polarization demultiplexing can be combined and included in the DSP to improve tolerance against distortion without affecting noise or performance.

Phase modulation, coherent receivers, polarization multiplexing, electronic dispersion compensation and enhanced forward error correction (FEC) all provide promising solutions to deliver 100G transmission. But there will always be the inherent trade-offs of cost, complexity and maturity of the technology and how they affect the performance of the entire transmission system and margin.

What slowed the 40G market were the multiple modulation formats proposed, which led to proprietary deployments at much higher cost points and large packaging. Each vendor chose a different format to deliver 40G, resulting in higher costs across the board for any implementation of 40G. This is not the case today. Inexpensive cost packaging with standard modulation formatting will enable lower-cost components to come to market in volume within smaller packages and with lower power generation.

In this e-guide

100G DWDM technology constraints

Typical transmission issues and impairments arise in the gray area of operations and margin. Linear issues typically are crosstalk or filtering on the line, amplification, OSNR, PMD and dispersion compensation. These often depend on the chosen fiber type, bit rate, modulation scheme deployed and channel spacing. Increasing channel spacing to 50 GHz enables nonlinear impairments such as intrachannel effects where the colors travel at different speeds, resulting in forward mixing frequency (FWM) or pulses overlapping and causing intrachannel effects.

Research is focusing on coherent detection and electronic compensation to correct or reduce the impairments of higher bit rate transmission. However, there are different approaches: Other phases and dispersion compensation may be done in the time domain or the frequency domain. There are trade-offs with these directions, with some providing better performance but lower margin; some offering a better system margin but at higher cost points.

Future outlook for 100G DWDM network transport

ACG Research expects 100 Gigabit Ethernet connections to become available in small quantity by late 2010. By 2011, 100G wavelengths will debut on some production networks testing live traffic. By 2012, we expect first-generation 100G DWDM optical network transport deployments to be commercially available and deployed on the most bandwidth-constrained routes.

100G growth is probably dependent on the continued bandwidth growth and availability of second- and third-generation hardware, which will bring more compelling cost points to 100G. Once those are available, we expect rapid transition of 10G to 100G, but this depends on the cost efficiencies and price points that come to market.

As 100G technology becomes a more efficient use of data rates, we expect this will affect the 40G deployments worldwide, limiting them in scope once 100G is available. The market for 40G will still be viable for the foreseeable future, however, as the industry tests the waters for 100G. Trade-offs in deployment techniques must be seriously considered, but we believe that those which meet the requirements and deliver product at the cost-effective sweet spot for 100G will be the key to success. By then, there will be many networks ready to upgrade to 100G.

Next article

In this e-guide

Getting more PRO+ exclusive content

This e-guide is made available to you, our member, through PRO+ Offers – a collection of free publications, training and special opportunities specifically gathered from our partners and across our network of sites.

PRO+ Offers is a free benefit only available to members of the TechTarget network of sites.

Take full advantage of your membership by visiting http://pro.techtarget.com/ProLP/

Images; Fotalia