# ADVANCES IN DWDM ROADM TECHNOLOGY USING PHOTONIC INTEGRATED CIRCUITS

Gaylord Hart
Infinera

*Abstract*

*Modern photonic integrated circuits (PICs) integrate multiple optical subsystems on a single chip, which greatly reduces the traditional cost structure for DWDM ROADMs (reconfigurable optical add/drop multiplexers). This allows ROADMs to be cost-effectively architected for the first time using an Optical-Electrical-Optical (OEO) conversion for every wavelength at every node, thereby allowing the use of digital electronic switches for reconfigurability instead of all-optical, wavelength-only switches. The resulting digital ROADM enables new capabilities and yields significant advantages over analog all-optical ROADMs.*

## INTRODUCTION

Optical transport systems for delivering digital services have evolved significantly over the last several years. At each stage of this evolution, advances have been driven by economics, as well as the need for greater capacity and scalability. Inflection points in this evolution have typically occurred when technological breakthroughs have enabled a paradigm shift that allowed significant cost reductions or new, advanced capabilities, or both.

The first optical transport systems deployed were point-to-point, single-wavelength systems. To create larger networks, these platforms were often placed back-to-back at a node and electrically interconnected for traffic to transit the node. This is an inefficient and expensive method for building optical networks.

To address these shortcomings, add/drop multiplexers were developed which essentially combined the two back-to-back platforms into the same chassis. In these systems, an OEO conversion is performed by the multiplexer at each node so services can be digitally added, dropped, groomed, or switched. Services merely transiting the node still undergo the OEO conversion, but are digitally directed to the next node. This type of system is typified by SONET multiplexers and offers the benefit that any service can be digitally groomed or reconfigured at any time for add/drop or pass-through.

As the demand for increased transport capacity has grown, these systems have not scaled well because they require a fiber pair for each pair of connected multiplexers, resulting in fiber exhaust in many cases. To address this shortcoming, WDM systems were developed which allow multiple wavelengths to be carried on a single fiber pair, connecting multiple multiplexers at either end.

While these WDM systems relieved the fiber exhaust problem, they do not economically scale well because they utilize an expensive OEO conversion based upon discrete optical components for every wavelength at every node, even for wavelengths which are not adding or dropping services at the node. To squeeze additional costs out of these platforms, fixed optical add/drop multiplexers (FOADMs) were developed so that only those wavelengths adding or dropping traffic at a node undergo an OEO conversion. These systems utilize transponders to add or drop a specific service on a specific wavelength at a specific node. All other wavelengths are passed through the node in the optical domain.

FOADMs relieve fiber exhaust and lower CapEx costs by reducing the number of OEO conversions in a network, but they also created new challenges. Since some wavelengths now pass through nodes in the optical domain, optical engineering for the network is no longer a single-span engineering problem, but a multi-node, network-wide challenge. Additionally, analog optical impairments now accumulate for those wavelengths that must transit through multiple nodes before undergoing an OEO conversion. This can limit the size of networks and may require periodic re-gen of these wavelengths to remove the impairments.

FOADMs also present operational challenges. When a typical FOADM network is initially turned-up, it requires manual power balancing of every wavelength at each node in the network to ensure optimal operation. Moreover, these networks typically require rebalancing whenever services are added or deleted. If the initial network design was not carried out guaranteeing any-to-any connectivity, but only based upon initial services, it is possible the entire network may require a redesign and reconfiguration when services are changed. This is inconvenient and disrupts existing services on the network.

In all but the simplest of networks, FOADMs are time-consuming and complex to engineer and operate. To address some of these limitations, reconfigurable optical add/drop multiplexers (ROADMs) have been developed. Optical ROADMs still use transponders for the OEO conversion for adding and dropping services at a node, and they still pass through a wavelength in the optical domain when no services are being added or dropped locally from that wavelength. But a ROADM offers three capabilities a FOADM does not typically provide: auto power balancing, optical wavelength switching, and a communications control plane to automatically and remotely reconfigure the network when necessary. Some ROADMs also support transponders with tunable lasers to provide enhanced re-configurability.

Optical ROADMs are a great improvement over FOADMs. They accelerate and simplify network turn-up and service changes by automatically balancing optical power throughout the network. They also allow wavelengths to be remotely switched for optical add/drop or pass-through operation at each node. However, they still have many of the limitations of FOADMs, including complex optical layer engineering and a dependence on transponders, which tie services to wavelengths and hence make service layer engineering dependent upon optical layer engineering.

One way to address these optical ROADM limitations is to perform an OEO conversion for every wavelength at every node, then process all services digitally at each node. This would reduce the optical network engineering problem to a series of simple single-span designs, even for mesh networks. Using 3R (regenerate, reshape, retime) OEO conversions at each node, optical impairments would not accumulate, and this would allow networks of essentially any arbitrary size to be built. Finally, digital processing at each node would allow replacing the all-optical wavelength selective switch of the traditional ROADM with a digital switch which can support advanced capabilities not possible with an optical layer switch. Such a digital ROADM would allow for the first time the creation of a flexible digital optical network where the service layer is independent of the optical layer

In the past, digital ROADMs were not economically feasible due to the many discrete optical components required to perform an OEO conversion for every wavelength at every node. However, recent advances in PIC technology allow complete optical subsystems to be economically placed on a pair of chips (TX and RX) less than 5 mm square and supporting 100

Gb/s per chip in a 10λ x 10G DWDM configuration. These PICs integrate over 60 discrete optical components (lasers, detectors, mux/demuxes, etc.) on a single pair of chips, eliminating all the discrete optical packaging and the fiber jumpers formerly required to interconnect the devices (see Figure 1, below). The resulting cost-reduction enables digital ROADMs to be built cost-effectively for the first time, and commercial digital ROADMs built on PIC technology have now been available for three years. PIC technology yields other benefits, as well: reduced power consumption, smaller footprint, reduced heat generation, fewer modules, fewer fiber jumpers, and higher reliability.

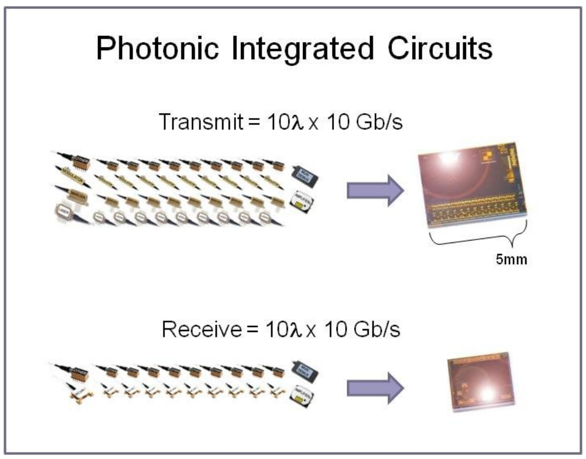


**Figure 1 – Large Scale Photonic Integration Places Several Optical Components on a Single IC**

Just as large-scale integration in digital electronics has enabled radical cost reductions even while increasing processing power, PICs now enable the digital paradigm shift to digital optical networks. Even more exciting, Moore's law indicates significant gains are possible in future generations of PIC technology. Fully functional PICs have now been built in the lab which support 10λ x 40G configurations, and these 400 Gb/s PICs are scheduled to be available in 2009.

## ROADM ARCHITECTURE COMPARISONS

SONET has lost favor with most MSOs as they have migrated to Ethernet, but SONET is still widely deployed in the Telco world and in some CATV commercial services environments where some customers require SONET transport. FOADMs are also widely deployed, but are finding more favor in applications where the networks are smaller and less complex. For metro core, regional, and national DWDM networks, most cable operators are now deploying ROADMs to lower total cost of network ownership and to accelerate new service and bandwidth turn-up. In these networks, digital ROADMs now compete with all-optical ROADMs for market share and technical leadership.

Optical and digital ROADMs typically share many common attributes, including the ability to provide dispersion compensation, amplification, and automated optical power balancing at each node. They both typically implement an intelligent control plane (preferably using GMPLS) that allows remote and/or automated configuration as well as other features. Many support automated topology and inventory discovery and optical layer turn-up. Where optical and digital ROADMs primarily differ is in the way they provide core re-configurability: whether they switch in the optical or digital domain. To understand these differences, it is necessary to examine the way optical and digital ROADMs are architected. These differences have implications in the cost, engineering, and operation of a network, and a digital ROADM, as we shall see, provides significant advantages in each of these areas.

### Optical ROADM Architecture

Optical ROADMs (just as FOADMs do) only perform an OEO conversion at a node for wavelengths being added or dropped at that node. Wavelengths not being added or dropped at the node are "expressed" through the node in the optical domain.

Unique to the optical ROADM is a wavelength selective switch (WSS), or equivalent functionality implemented with wavelength blockers or other technology, which allows individual wavelengths received at the node to be switched for local add/drop or for "express" transit through the node. Figure 2, below, shows the high-level architecture of a typical 2-degree optical ROADM.

Looking at signal flow from west to east, the RX optical signals originating from the network's west-facing fiber interface are presented to an optical demux which breaks out the individual wavelengths for processing by a software configurable optical switch. The switch then individually passes each RX wavelength directly through the ROADM and out the reciprocal east-facing TX optical mux, or it drops it to a west transponder for local service handoff. The TX signals originating from these west transponders are usually passively added to the pass-through channels travelling from east to west and sent out the west-facing TX fiber.

The ROADM also provides reciprocal functionality for the TX and RX optical signals originating from the network's east-facing fiber interfaces. For redundancy purposes, this is implemented in the ROADM with separate east and west optical modules with each containing an optical switch and any other required optical muxing, combining, or splitting components. Some ROADMs also provide optical performance monitoring points here for "express" wavelengths, as well.

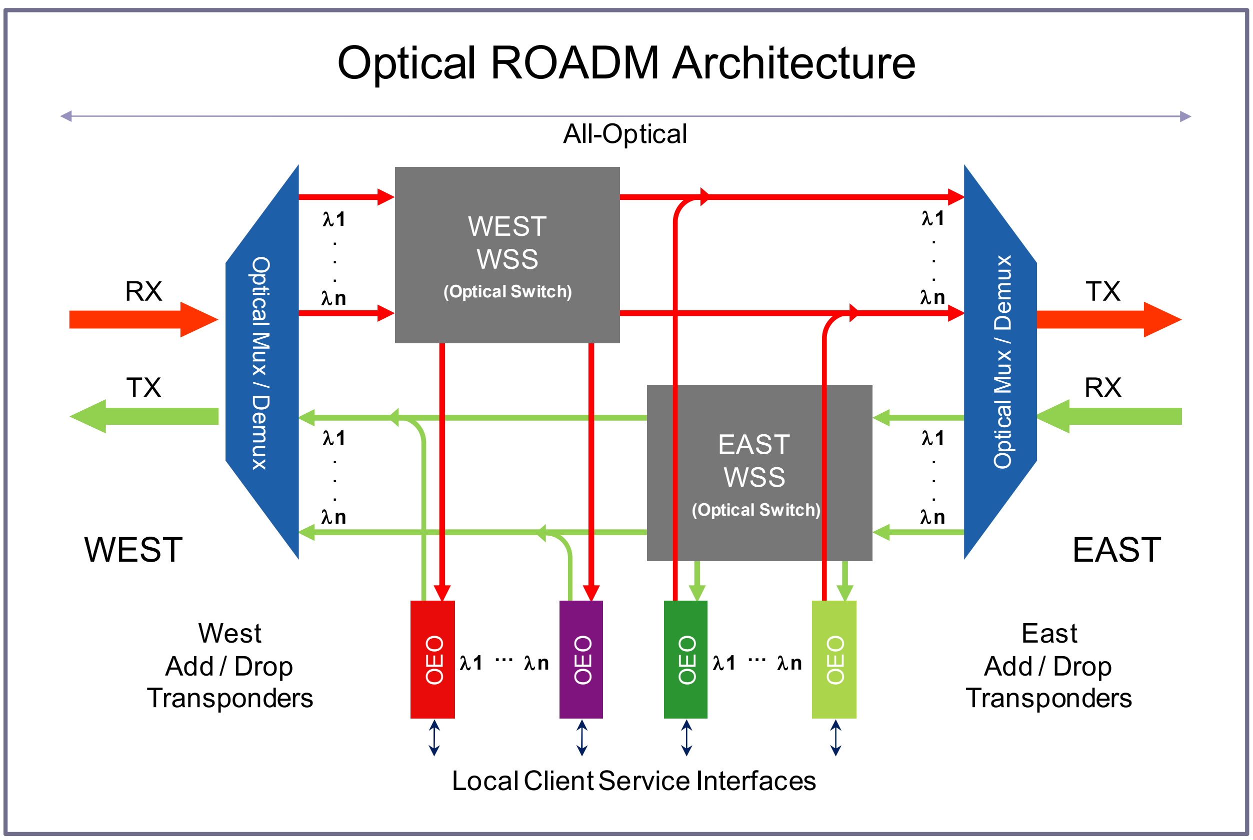


**Figure 2 – Typical 2-Degree Optical ROADM Architecture**

It should be noted that some optical ROADMs are banded, requiring add/drop or express treatment for a group of wavelengths as a whole (banding is usually implemented in wavelength groups that are an integer multiple of four). Because banded ROADMs must treat all wavelengths in the banded group the same, this can result in stranded bandwidth. For example, if only two wavelengths need to be dropped off at a node, and the banding group size is four, the ROADM will strand two wavelengths at the node.

For the wavelengths being dropped off locally, the transponders perform an OEO conversion and provide client optical interfaces for connection to local equipment. Between the client and line-side optical interfaces, services are processed digitally on the transponder, and it is here that bit-error-rate (BER) evaluation, forward error-correction (FEC), digital wrapper processing (usually G.709), and digital performance monitoring takes place. The transponder also typically provides performance monitoring for any service-specific attributes. Optical performance monitoring points for the add/drop signals are frequently provided here as well.

Each transponder at the node links a particular line wavelength to a particular client service interface (if tunable laser transponders are used, this provides a greater degree of reconfigurability, but at a higher cost than fixed wavelength transponders). Because the same wavelength may arrive from the east and west, and for redundancy purposes, the ROADM supports both east and west-facing transponders, and these typically can only receive wavelengths from the direction of the optical switch module they are attached to. When protected services are required, two transponders must be used, one each for east and west. Service growth is implemented by adding transponders at the service source and destination nodes, and these are typically only installed at a node when services are actually required there.

Digital ROADM Architecture

Digital ROADMs perform an OEO conversion for every wavelength arriving at every node, both east and west. Wavelengths not being added or dropped at the node are still "expressed" through the node, but in the digital domain before the conversion back to optical when exiting the ROADM.

Unique to the digital ROADM is a core digital cross-connect switch connected to every digital signal derived from every optical interface, including the line interfaces and the client interfaces. This switch allows any input to be connected to any output. Because this switch resides in the middle of the OEO conversion (unlike with an optical ROADM, where the OEO takes place in the transponder after optical switching), the line-side optical layer is completely segregated from the client side service layer. Figure 3, below, shows the high-level architecture of a typical 2-degree digital ROADM constructed with PICs.

At each line-side interface, both east and west, an optical band multiplexing module segregates groups of incoming wavelengths into Optical Carrier Groups (typically comprised of 10 wavelengths at 10G each) for handoff to the RX PICs. The same band multiplexing module aggregates the reciprocal group of outgoing wavelengths from the TX PICs for handoff to the line-side fiber. In turn, each PIC simultaneously processes multiple line-side wavelengths in parallel, performing O-E conversions on the RX side and E-O conversions on the TX side. Each PIC processes 10 WDM wavelengths at 10 Gb/s each. A typical initial ROADM deployment starts with a pair of TX/RX PICs on each of two digital line cards facing east and west, respectively. This provides an initial transport capacity of 100Gb/s in both directions. Services are then added until all 10 wavelengths on the initial PICs are consumed, and then an additional pair of digital line cards is installed to support further service growth.

The electrical interfaces on each PIC are connected to the core digital cross-connect switch, which is also connected to tributary adapters for local service handoff on the client side. This switch is under software control and may be remotely configured. It is implemented in a redundant configuration to eliminate any single points of failure.

Internal to the ROADM, all services are digitally processed. As the PIC's electrical signals interface with the core digital switch, they are processed for BER monitoring, FEC, digital wrapper manipulation, and other performance monitoring. Once inside the core switch, these signals may be directed in any direction, to any interface (line or client), and to any wavelength (east or west, any color). In this manner, express services are switched directly through the ROADM, and add/drop services are directed to local client interfaces.

Services are added or dropped at the node through a tributary adapter. The tributary adapter converts the digital transport stream from the core switch for native service handoff and provides this service on a client optical interface for local use. The tributary adapter also provides performance monitoring for any service-specific attributes and loopback capabilities for test purposes.

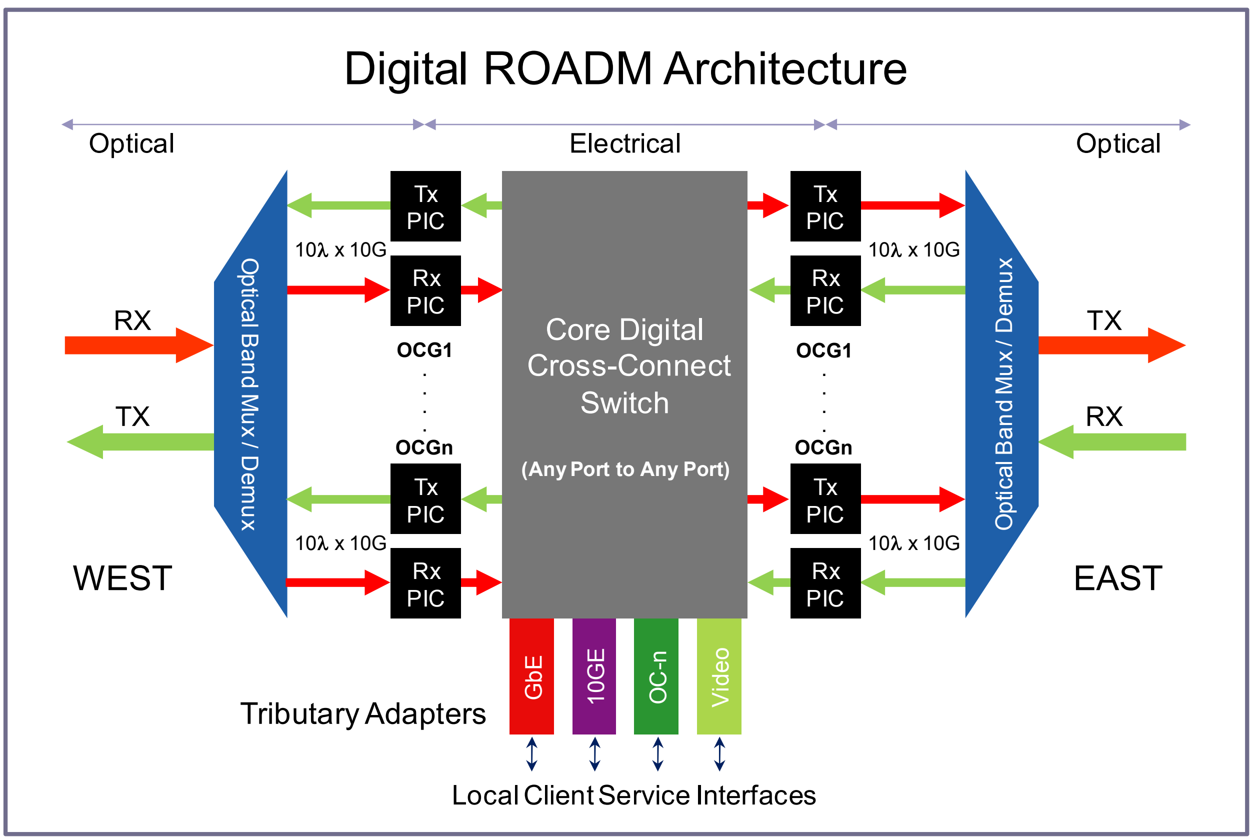


**Figure 3 – Typical 2-Degree Digital ROADM Architecture Using Photonic Integrated Circuits**

OPTICAL LAYER ENGINEERING

The optical layer engineering process and requirements are very different for optical and digital ROADMs. Optical ROADMs only perform an OEO conversion for wavelengths locally adding or dropping services at a node. All other wavelengths are passed through the node in the optical domain. For these pass-through wavelengths, optical impairments such as dispersion and cascaded filter losses accumulate from span to span. Since the source and destination nodes may be anywhere in the network, each wavelength must be individually engineered for its particular path through the network. But since different wavelengths often share common spans in the network, this is not a simple process. The final design must represent a common denominator that works for all wavelengths in the network, regardless of their paths, and this may result in a less than optimal design that may limit node counts or span distances. In some cases, it may be necessary to perform a 3R re-gen in the network, and this requires back-to-back transponders for every wavelength requiring re-gen.

A digital ROADM performs an OEO conversion for every wavelength at every node. In this case, all wavelengths on a span share the identical path regardless of whether they are being added or dropped at a node, and only individual spans need to be engineered between nodes. Moreover, because a 3R OEO conversion takes place at each node, optical impairments do not accumulate from span to span in the network. This essentially allows networks of any arbitrary size to be built.

Ideally, an optical network's initial design should allow any service at any node to be transported to any other node in the network at any time without requiring any re-engineering or re-configuration of the optical layer to do so. To guarantee this any-to-any connectivity in an optical ROADM network, every relevant analog optical transport parameter (OSNR, dispersion, power levels, etc.) has to be calculated east and west for TX and RX for every wavelength for every possible combination of end-nodes (see Figure 4, below). In a ring network, this requires $2(N^2\text{-}N)WP$ calculations (where N is the number of nodes in a network, W is the number of wavelengths supported by the network, and P

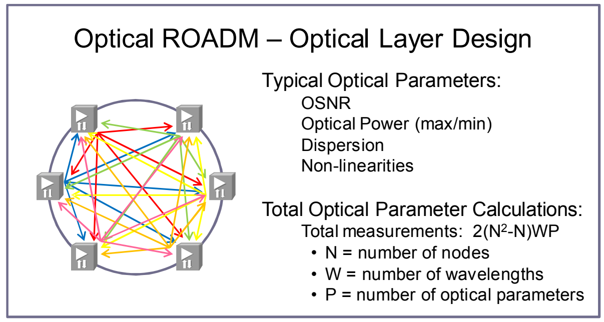


**Figure 4 – Optical Layer Design for an Optical ROADM**

is the number of transport parameters to be verified).

In a large ring network with a large number of wavelengths, thousands of calculations may be required to verify the optical layer design when optical ROADMs are used. If the initial design assumptions prove unworkable, redesign and recalculation may be iteratively required to find a workable combination of all optical layer parameters. For a mesh network with multi-degree nodes, the engineering problem becomes vastly more complex.

Digital ROADMs limit the optical layer design to a series of independent span designs where the design of one span does not impact the design of any other span (see Figure 5, below). In a ring network, this requires a total of 2NWP calculations (where N is the number of nodes in a network, W is the number of wavelengths supported by the network, and P is the number of transport parameters to be verified). To guarantee any-to-any connectivity in the network, one only has to guarantee that each individual span between a pair of nodes has been

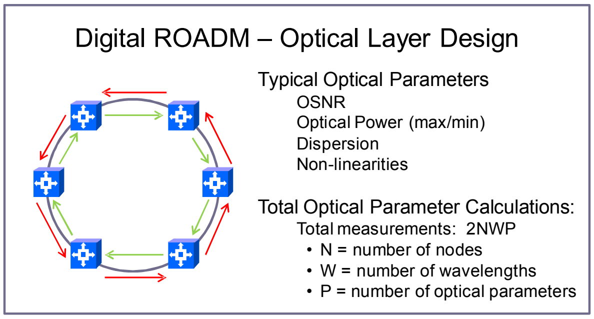


**Figure 5 – Optical Layer Design for a Digital ROADM**

properly engineered. And this is true even for large, complex mesh networks.

## OPTICAL vs. DIGITAL PROCESSING

Aside from the major difference in these ROADMs between switching in the optical or digital domain, another key difference is where the switching is actually performed. An optical ROADM switches above the OEO conversion, and the transponder which actually performs this conversion tightly couples the line-side optical transport layer to the client side service layer. In contrast, a digital ROADM separates the OEO conversion process into distinct line-side optical layer and client-side service layer interfaces, and does its switching in the middle of these.

Having examined the respective architectures for optical and digital ROADMs, we can now examine the specific capabilities supported by each.

### Switching

Optical ROADMs, of course, switch entire wavelengths, which are typically transported today at a nominal 2.5 Gb/s or 10 Gb/s line rate. If any of these wavelengths have lower rate services multiplexed up to the line rate by the transponder (muxponder), as is common to maximize the use of transport capacity, these sub-rate services cannot be individually switched, groomed, or added/dropped at intermediate nodes on the network unless back-to-back transponders are used where this functionality is required. This is because the sub-rate services can only be processed digitally, and an optical ROADM only has access to digital signals at the transponders, which are only located at the source and destination points for a service.

While some transponders provide some integrated switching capability, this is typically limited to within the transponder itself or to an adjacent back-to-back transponder for digital add/drop capabilities. It is not uncommon when service grooming or switching is required at a node to accomplish this with manually installed fiber jumpers between the client interfaces on the back-to-back transponders. In these cases, one is often faced with the economic tradeoff between stranding bandwidth or paying for additional transponders.

In WDM optical networks, especially mesh networks, it is quite common to have wavelength contention (or even blocking) between services being carried over the network. This occurs when two independent services using the same wavelength need to travel over the same fiber on at least one span. Of course, one may simply use a different wavelength for one of the services when this occurs, but this may in turn create contention with another service on another span. To fully utilize available network capacity, it may be necessary to perform a wavelength conversion for a service for transport over the contended path. Because the transponders on an optical ROADM tightly couple a service to a wavelength between the transponders, and because an optical ROADM cannot switch services between wavelengths, the only way to perform this wavelength conversion is to use back-to-back transponders using different wavelengths on the contended path. This is an expensive solution, but may be the only option when blocking occurs.

Digital ROADMs provide simple and inexpensive solutions to these switching challenges, and provide additional capabilities as well. As already noted, these ROADMs have digital access to all services and wavelengths at every node and also have an integral digital switch interconnecting all interfaces at every node. This combination yields very powerful capabilities for switching services and maximizing bandwidth usage in the network. An examination of these capabilities follows, but a more detailed discussion of how a digital ROADM processes services is necessary first.

To maximize transport capacity, advanced digital ROADMs utilize PICs providing 10 x 10 Gb/s wavelengths for transport across the network. A 10 Gb/s digital transport frame (DTF) using a G.709 digital wrapper or an enhanced version of it is used on each wavelength to transport native client services end-to-end. The DTF also provides forward error correction (FEC) and performance monitoring, not only between intermediate transit nodes, but between end-to-end service points as well. The DTF in turn has four 2.5 Gb/s digital signals asynchronously multiplexed into it, thus the DTF may transport one 10G or four 2.5G services.

The DTF and its 2.5G sub-rate signals are transparent to the client signals and may carry Ethernet, SONET or other protocols. A 10G DTF can simultaneously support any combination of client signals mapped into its 2.5G signals, up to the full 10G rate. A 2.5G signal may in turn have 2 GbEs mapped into it. Thus a 10G DTF may carry one OC-192, one 10GE, four OC-48s, 8 GbEs, or some combination of these that does not exceed the 10G line rate. Other mapping possibilities exist, and other protocols can be supported as well. The DTF's flexible mapping capabilities allow each wavelength to be used efficiently by multiplexing any combination of sub-rate services into it until the wavelength is fully utilized.

At the switching level, the ROADM's integrated digital switch works at the 2.5G sub-rate granularity, so sub-rate switching is fully supported along with switching at the 10G line rate. This supports add/drop multiplexing of any service at any node at any time. In fact, even the tributary interfaces at a node may be switched to face east or west in the network. Unlike with transponder based optical ROADMs, digital switching makes service transport completely independent of the transport wavelengths: full grooming and switching of every service at every node (including sub-lambda services) is possible to and from any wavelength, and the service bit rate is independent of the transport line bit rate. Figure 6, below, shows several of the processing capabilities provided at every digital ROADM node.

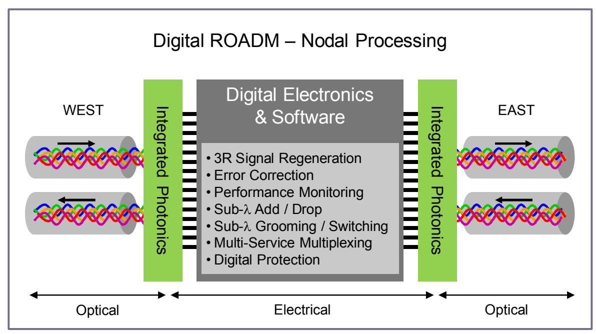


**Figure 6 – Digital ROADMs Process All Services Digitally at Each Node**

Digital switching is particularly effective in creating multi-degree networks, especially more complex mesh networks, where the complexities of optical layer engineering and wavelength planning become great. With a digital ROADM, the optical layers of each path at a junction node are independent, and traffic is simply switched between them digitally. Services may then be routed end-to-end through any available path in the network, and services may be created or torn down on demand without any optical layer engineering. Mesh networks provide enhanced protection and bandwidth management options by providing multiple paths between endpoints in the network.

Digital ROADMs are also quite effective in building hierarchical networks, where a single digital ROADM can serve, for example, as a junction site between a regional and metro core network. In this configuration, a digital ROADM can provide north and south interfaces for one network, and east and west for the other. Traffic may then be switched digitally in any direction between or within the networks, and protection is supported for both networks. Because a digital ROADM provides 3R OEO conversions at every node, there are inherently no distance or node count limitations in a digital network. It is actually possible to build a national backbone with integrated regional networks and metro core rings using a scalable digital ROADM.

Bandwidth Virtualization

Whereas transponders in an all-optical ROADM tightly couple services to wavelengths, a digital ROADM for the first time allows the service layer to be fully segregated from the optical layer. This allows wavelengths (and 2.5G sub-lambdas) to be treated as a pool of virtual bandwidth, that is, as an allocatable resource to be assigned to services when and only as needed. As long as sufficient bandwidth exists between any two nodes in a digital ROADM network, any service may be turned up between these nodes at any time. If sufficient bandwidth does not exist, digital line modules may simply be added in those spans where more bandwidth is needed to support the new service. Since PIC based line modules add bandwidth in groups of 100 Gb/s, any unallocated bandwidth on the new line module is added to the pool for future use. In this way, capacity usage is optimized on individual wavelengths and between individual nodes.

Because a digital ROADM provides 3R OEO conversions for all wavelengths at every node and because optical layer engineering is then limited simply to spans between these OEO nodes, service layer engineering is now made fully independent of optical layer engineering. Once a span has been engineered for the first digital line module, additional modules may be installed without the need for any additional optical layer engineering or reconfiguration. This greatly reduces the time and effort required to turn up new bandwidth and services.

Bandwidth virtualization enables flexible reconfiguration of the network at any node at any time without any optical layer re-engineering and provides significant additional capabilities that cannot be provided with an all-optical ROADM. First, unlimited add/drop capabilities are supported for any service at any node at any time, including sub-lambda services. Second, wavelength or sub-lambda services may be groomed or switched from any wavelength to any other wavelength at any node. This enables wavelength conversion for wavelength services and inter-wavelength grooming for sub-lambda services, so a service may actually be transported on several different wavelengths in the network between its source and destination points. This flexibility allows existing wavelengths in the network bandwidth pool to be fully utilized before having to add more capacity and provides a simple solution to wavelength contention or blocking.

Bandwidth virtualization also supports super-lambda services. For example, a 40G service can be delivered using tributary adapters that concatenate 40G service transport over four 10G wavelengths from the available bandwidth pool. In this manner, 40G services may be delivered over a 10G network without any additional optical layer engineering or reconfiguration. Since the client handoffs at either end of the service are standard 40G interfaces, service delivery is fully transparent.

Digital processing enables other features as well, many of which support advanced maintenance capabilities or unique architectures. A digital ROADM supports hair-pinning services, where a service is brought into the ROADM on one client port and is "hair-pinned" out a different client port at the same node. Bridge-and-roll is also supported. Using bridge-and-roll, an alternate path is created for a service between the two endpoints. A bridge is then created to duplicate the service over both paths. Finally, a roll operation is executed which transfers the end-to-end service over to the new path in under 50 ms. The original path is then free for maintenance operations. Non-obtrusive digital test ports may also be created at any node in the network to observe traffic on any service. In this case, a service is digitally replicated at any node in its path, and this copy is then sent to another client port at any node in the network, where it may then be attached to test equipment or otherwise examined. Finally, digital processing enables a wide range of loopback options to be used for circuit verification and fault-isolation purposes, including the ability to perform loopbacks not only at the service endpoints, but at intermediate nodes in the service path as well. All of these maintenance and test mechanisms are remotely provisionable.

Using digital multicast, digital ROADMs also make unidirectional drop-and-continue digital video broadcast architectures simple and inexpensive to implement at Layer 1. In this application, the ROADM's digital switch replicates the broadcast digital video service at any required node, drops it locally, and then passes it on to the next node. At each node, up to three output multicasts may be created, supporting local drop and multicast branching at junction nodes. Since the multicast is handled digitally, no optical splitting or specialized transponders are required, and there is no limit to the number of drop or branching sites. Digital multicast can also be used to support switched digital video when an edge-switched architecture is used for this application. Digital multicast services may be digitally protected for increased reliability.

Performance Monitoring

All ROADMs provide some degree of performance monitoring, and the information derived from this is used for many purposes, including guaranteeing service level agreements (SLAs) with end customers, establishing base network operating parameters so any observed degradation can be used to solve problems pro-actively before an outage occurs, and to diagnose and sectionalize an outage if it does occur. All ROADMs provide performance monitoring points (PMs) for both optical and digital parameters, but significant differences exist between the number of parameters monitored and where they are monitored.

Optical PMs typically consist of optical power levels measured at various points in the ROADM, but may include other parameters as well, such as laser bias current. Power measurements may be aggregate power (combined power of all wavelengths at the PM point) or individual power (power level of an individual wavelength at the PM point). An aggregate power level PM can indicate a problem exists (level too high or low), but it can't tell you which wavelengths are contributing to the problem and therefore is usually less helpful diagnosing and locating a particular problem.

Optical ROADMs typically provide optical PMs at various TX and RX points in the ROADM, but how and where these are measured varies considerably. Optical PMs should be provided for both TX and RX signals on both the east and west line-side interfaces. Ideally, these should provide aggregate and per-lambda data for express wavelengths as well as those originating from local transponders. This provides the greatest amount of useful information for diagnosing and pinpointing any problems. However, some ROADMs only measure optical power levels on the line-side interfaces in aggregate and only look at individual wavelength levels on the local transponders.

Digital ROADMs, because they perform an OEO conversion for every wavelength at every node, typically provide a full complement of optical PMs, measuring TX and RX power levels for every wavelength on both the east and west line-side interfaces. Aggregate optical power is also usually monitored on the line interfaces. This provides rapid and robust fault diagnosis and location capabilities at the optical layer.

Optical power level monitoring is quite useful, but there are only a limited number of optical layer problems which can be diagnosed with power levels. For example, optical power levels can help diagnose fiber cuts or laser failures, but not dispersion problems. Moreover, optical PMs are of no help in diagnosing digital or service layer problems, and these are usually assessed using digital PMs.

Digital PMs typically provide a large amount of information derived from a large number of monitoring points. Typical information would include loss of frame (LOF), loss of signal (LOS), uncorrected BER, corrected BER, errored seconds, severely errored seconds, and numerous other parameters. This information is usually gathered on the optical transport path for each service and/or wavelength being transported on the network, depending on the particular ROADM. BER data is usually gathered for the optical transport path via a G.709 or enhanced digital wrapper, which provides FEC, as well. Additional service-specific PMs (e.g., Ethernet frame and errored frame counts) are also typically available.

Optical ROADMs, by their nature, only provide digital PMs where signals are processed digitally (i.e., only at the transponders). Since transponders are only used at the service endpoints, digital PMs are only available there, and not at any intermediate nodes (see Figure 7, below). This makes diagnosing and localizing a problem more difficult since intermediate nodes in the path cannot provide PMs which might indicate which section is responsible for the problem.

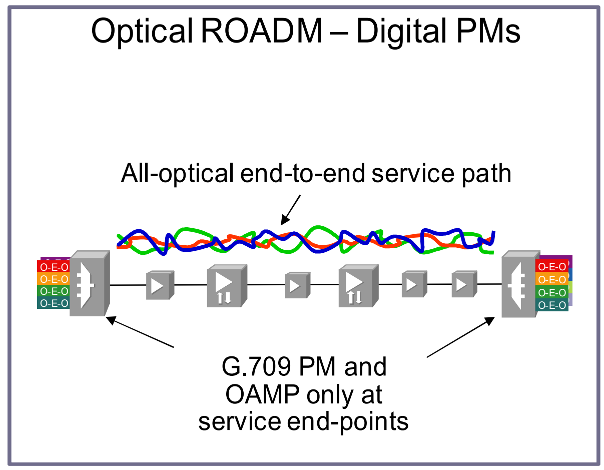


**Figure 7 – Digital PMs on an Optical ROADM**

Digital ROADMs, on the other hand, perform an OEO operation at every node for every wavelength, so digital PMs are available at every node for every signal that transits the node (see Figure 8, below). Digital ROADMs provide PMs for both the transport section (transport between each node) and the transport path (end-to-end service transport). This allows rapid pinpointing of any problems to an individual link between two nodes.

To facilitate testing, some digital ROADMs also directly incorporate a pseudo-random bit stream (PRBS) generator for direct BER testing between any two nodes without the need for any external test equipment. In PIC based digital ROADMs, a PRBS stream is run continuously on those 10G wavelengths which are installed on operational digital line cards but which are not yet carrying services. This provides the MSO

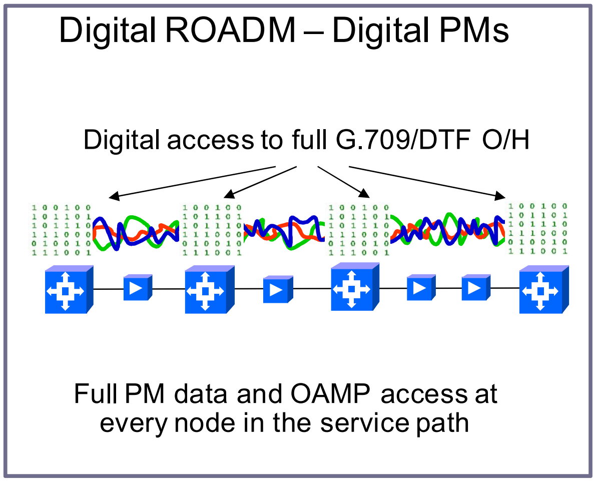


**Figure 8 – Digital PMs on a Digital ROADM**

with an operational history and track record even for those wavelengths not yet turned up with services and provides a higher degree of confidence in successful service turn-up when the time comes to activate these wavelengths.

Protection

Optical ROADMs typically protect a service by using two route-diverse paths between the endpoints, with a pair of transponders on each path. One transponder pair serves for the working path, one for protection. This configuration provides effective equipment and path protection, but the protection wavelength and its transponders are dedicated and cannot be used for other services. A single client handoff at either end is usually provided by a Y-cable connected to the clients on the working and protect transponders. Some optical ROADMs support a digital backplane interface between the transponder protection pairs, which allows a single client interface on one of the transponders to be used for the protected service without requiring a Y-cable. In this case, the best signal from either of the transponders is presented at the single client. For 1+1 protection, where the local end device connected to the ROADM (typically a switch or router) actually performs the protection switching, the clients on each of the two transponders are connected directly to reciprocal clients on the end device.

Digital ROADMs can also be configured for dedicated protection. In this case, both working and protection paths are pre-defined through the network between the service endpoints. These paths are set up through each node using the ROADM's integrated digital switch and any available wavelengths through the network. Route diversity is used for path protection. At the endpoints, a pair of tributary adapters is used to provide the protected clients. As with optical ROADMs, the paths and tributary adapters are dedicated and cannot be used for other services. A Y-cable may be used to provide a single protected client interface, just as with optical ROADMs, or the ROADM's integrated switch may be used to deliver the best signal from either path to a single tributary adapter client. This second method saves the cost of one tributary adapter and the Y-cable. Dedicated 1+1 protection is handled with two tributary adapters just as with an optical ROADM.

Digital ROADMs, however, support shared protection modes which are not supported by most optical ROADMs. With shared protection, the protection wavelengths are left uncommitted throughout the network, and these remain in the pool of allocatable bandwidth until actually needed. In this way, a small shared bandwidth pool can be used to protect many services across the network, resulting in much lower bandwidth consumption when compared to dedicated protection.

When a failure occurs in a shared protection network, the GMPLS control plane finds a path through the network with sufficient bandwidth to restore service, allocates this bandwidth from end-to-end, then switches the service over to the new path. If sufficient paths and bandwidth are available, protection may be provided over multiple failures at multiple points in the network. Since multiple failures in different segments of the network are unlikely to occur simultaneously, shared protection provides an effective means of conserving bandwidth while providing a high degree of confidence in protecting against network failures. Shared protection may be used in ring or mesh networks, but mesh networks typically provide more paths through the network between service points and therefore may provide more options for protection routing than would be available in a ring network.

Unlike dedicated protection, shared protection frees up significant bandwidth in the network for other uses, and the network operator has full control over how much spare bandwidth to provide for shared protection. However, there is a tradeoff in using shared protection. A dedicated protection path, because it is always live, provides protection switching in under 50 ms. A shared protection path must be found and routed after a failure, and this can take a few seconds. In a digital ROADM, protection may be provisioned as dedicated or shared on a service by service basis, so this is typically not a problem since any sensitive services can always be configured with dedicated protection.

## SUMMARY AND CONCLUSIONS

Photonic integrated circuits represent a major inflection point in optical networking evolution, enabling the digital paradigm shift to digital optical networking and delivering the scalability that will be required for next-generation networks.

Major advances in photonic integrated circuits have resulted in commercial production of inexpensive, highly reliable photonic ICs integrating all the components required to deliver ten 10 Gb/s wavelengths, on a pair of chips (TX and RX) no more than 5 mm square. This allows a low-cost OEO conversion to be used for every wavelength at every optical multiplexer in a network. The PIC cost-savings in turn support integrating a full digital cross-connect switch into the core of the multiplexer, creating for the first time a cost-effective digital ROADM.

The architecture of the digital ROADM (and hence its name) is such that all services are processed digitally, rather than optically. Since a digital ROADM performs an OEO operation on every wavelength at every node, its integrated digital switch has unrestricted access to every service entering or leaving the multiplexer, and therefore has unrestricted ability to groom, switch, or add/drop services at the node. This delivers a much wider range of reconfiguration options than an all-optical ROADM can provide, but it also enables a completely new set of features and capabilities.

Digital ROADMs greatly simplify optical layer engineering, and their 3R OEO architecture supports networks of essentially any size or shape to be built and provisioned easily. Because digital ROADMs segregate the optical layer from the service layer, turning up new services is quick and requires no optical layer engineering whatsoever. Digital ROADMs support wavelength conversion, sub-lambda grooming, inter-wavelength switching, and are in general more bandwidth efficient than their optical counterparts.

Digital ROADMs bring significant benefits to all phases of network ownership, from network engineering, to network turn-up, to service growth, and to network evolution. They offer a lower total cost of network ownership while accelerating network operations, bandwidth expansion, and service delivery.