CD ROADM – Fact or Fiction

Part 1: Applications

shaping tomorrow with you

FUJITSU

Colorless Directionless ROADM (CD-ROADM) architectures have recently generated considerable interest in the optical transport industry. CD ROADM is a relatively new type of WDM node designed to offer additional flexibility and simplicity. Industry technical conferences are frequently populated with papers discussing the applications and perceived benefits of CD ROADMs. Vendors are aggressively promoting CD ROADM in an attempt to outshine their competitors and impress optical transport technologists. Unfortunately, this increased exposure has led to a great deal of inaccurate and misleading information regarding the true benefits and costs of CD ROADM architectures.

This paper is the first of a two-part series that provides an unbiased look at CD ROADM architectures and their network applications, benefits, and costs. Part one examines the primary applications enabled by CD ROADMs. Part two examines the technologies required to implement a CD ROADM architecture and the costs compared to existing ROADM systems.

As a leading provider of optical transport solutions, Fujitsu is committed to supporting the industry regardless of the type of WDM architecture chosen, whether based on current ROADM technologies or CD ROADM architectures. The objective of this two-part paper is to separate fact from fiction, enabling carriers to make informed decisions about their optical transport networks.

ROADM Architectures

WDM networks are primarily designed by utilizing ROADMs, which incorporate the optical amplifiers, optical switching, multiplexer/demultiplexer, transponder, and muxponder cards to form a complete optical transport node. ROADMs are highly flexible and typically support 1–88 wavelengths, span optical reaches from 600 to 2500 Km, and allow channels to be added, dropped, or passed through on a per-wavelength basis at each node. This flexibility allows carriers to provision A–Z optical wavelengths across metro, regional, or long-haul routes. A typical ROADM implementation is shown in Figure 1.

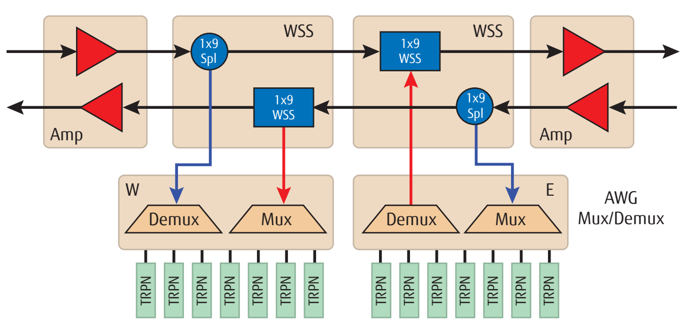

Figure 1: Typical 2D ROADM

The example in Figure 1 is a two-degree (2D) ROADM, with one WDM network fiber from the West direction and one from the East direction. Transponders provide client-side interfaces and can be connected either to the West or East directions via the respective mux/demux units. A mux/demux combines and separates individual wavelengths into different physical input and output ports, based on Arrayed Wave Guide (AWG) technology. An AWG is essentially a prism that separates the different wavelength "colors." This architecture is the most common ROADM implementation based on current generation technology.

This type of ROADM is sometimes referred to as a colored, directional ROADM. Each input/output port supports a specific ITU WDM wavelength, such as 1545.32 nm (channel 22). While the transponder and muxponder themselves are full-band tunable and can be provisioned to any transmit wavelength, they must be connected to a specific port on the mux/demux, which is known as "colored" add/drop technology. Likewise, each degree or direction (West or East) has their own mux/demux. A transponder connected to the West mux/demux only supports services connected to the West direction.

To re-assign wavelengths, either to a new channel or to re-route them to a different direction, requires technician involvement to physically unplug the transponder from one port on the mux/demux and plug it into a different physical port.

CD ROADMs

CD ROADMs allow wavelength re-assignment without the need for manual intervention, but with the tradeoff of higher complexity and cost. From a conceptual view, a CD ROADM architecture is shown in Figure 2.

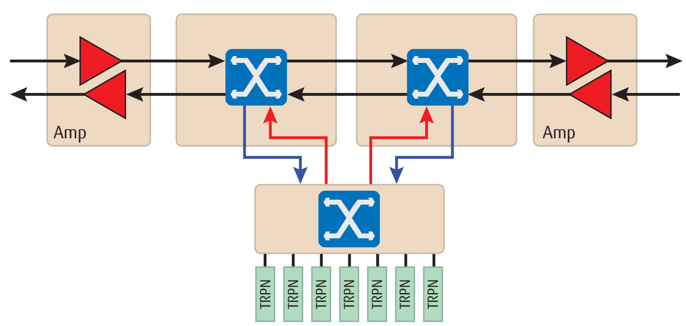

Figure 2: Colorless/Directionless ROADM

Figure 2 is a greatly simplified representation of the architecture and technologies used to implement CD ROADMs. In this diagram, the blue box with crossed arrows represents a colorless/directional optical switch. The exact details of the CD ROADM architecture, optical components used to implement the colorless/directionless switch, and a cost analysis of the architecture are provided in Part 2 of this series. With a CD ROADM, the transponders may be connected to any add/drop port and can be routed to any degree or direction. Wavelength re-assignment or re-routing can be implemented automatically from a network management system, without the need for manual technician involvement.

CD ROADM Applications & Benefits

While CD ROADMs offer additional architectural flexibility, the important question is whether the applications enabled justify the additional cost and complexity of CD ROADM nodes. The remainder of this paper examines CD ROADM applications, benefits, issues and limitations, with Part 2 providing an in-depth cost analysis.

Three fundamental applications are enabled by CD ROADM architectures:

- Bandwidth pre-positioning
- Bandwidth on demand
- Optical layer mesh restoration/re-optimization

Bandwidth Pre-Positioning

With bandwidth pre-positioning, carriers deploy "pools" of transponders and regenerators at major network locations, pre-connected to the CD ROADM nodes, as shown in Figure 3. As additional bandwidth and services are required, an operator simply establishes A–Z wavelength connections across the network. New services can be established in minutes as opposed to the normal weeks required to engineer, order, receive, deploy, install, and provision new equipment. Service velocity and minimizing technician time are the key advantages of bandwidth pre-positioning, especially for carriers with national backbone networks.

For carriers with large national core networks that transport high volumes of traffic between major hub cities, this approach may offer some advantages. Since these networks carry aggregated traffic from access, metro, and regional networks, the overall amount of traffic can be very large and the quarterly growth can be fairly predictable based on historical patterns. For example, if a carrier adds six 10G channels to a node in Chicago every six months, it may make more sense to deploy and connect all six transponders at one time and provision the service "as needed," as opposed to incurring the technician cost and time to deploy one transponder at a time every three or four weeks. Assuming the growth patterns at the core nodes are predictable, pre-positioning the transponders once a half can improve service velocity and reduce technician time and costs.

There are a couple of issues with bandwidth pre-positioning, including the cost of "pooled" transponders and technician involvement with client-side service activation. First, if the growth patterns are not predictable, this can lead to pools of unused transponders or regenerators scattered throughout the network. These unused or underutilized assets would be a network expense, expensive if they are 100G transponders, but they may ultimately be consumed by network growth over time. Second, a CD ROADM enables the pre-connection of the network side of a transponder, but it still requires a technician to connect and activate the client-side optical interface. If a technician is required to connect and provision client-side connections, there's little time or cost savings from pre-connection of the network-side interface.

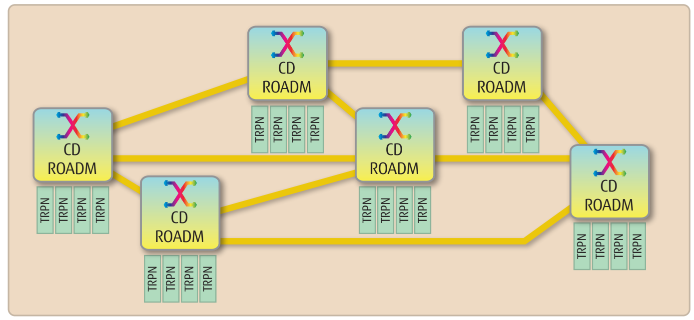

Figure 3: Bandwidth pre-positioning

Benefits	Issues
Improved service velocity • 10G/100G service turn-up in minutes	Cost of "pre-deployed" pooled units • Stranded resources if not fully utilized
Lower operations costs • Fewer truck rolls • Fewer technician "touches"	Client-side service activation • Client-side optical connection and activation • Still requires technician involvement
ROI may be attractive on large backbone/national networks	ROI may not be viable on metro or regional networks
Works best with predictable growth patterns	Strands resources if growth is not predictable

Bandwidth on Demand

Bandwidth on demand is another commonly referenced application and rationale for implementing CD ROADM architectures. Most telecom services are provided on a fixed basis between two or more customer locations and remain operational 24 x 7, whether or not customers are actually using their bandwidth. Over the years, the telecom industry has debated the feasibility of selling transport services "on demand" for fractional periods of time. For example, an enterprise customer may want a 10G wavelength service between primary and backup data centers, but only require the service between 1 and 3 am to perform daily backups and offsite archiving.

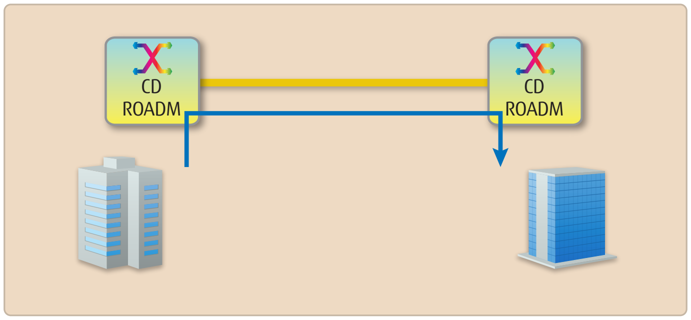

Figure 4: Bandwidth on demand

Bandwidth on demand services come with a whole host of their own issues, including whether a viable market exists, and whether there is a positive ROI for carriers. Several carriers have implemented bandwidth on demand trials in major markets over the years, utilizing SONET networks. While providing valuable experience and feedback, most of these trials were terminated due to lack of customer demand, OSS/billing complexity, or simple financial performance reasons.

Benefits	Issues
New, flexible service model	Previous carrier bandwidth "on-demand" trials had limited success
Competitive differentiation for carriers	Lack of clear end-user market for on demand services
Possibility of new end-user applications	Requires "idle" assets just to support "on-demand" peak intervals
Could benefit end users by lowering network bandwidth costs vs. dedicated private line wavelength service	May reduce carrier revenue, if customers switch to lower priced "on-demand" bandwidth
	Complex OSS, control and billing infrastructure

Optical Layer Mesh Restoration/Re-optimization

The most frequently cited justification for CD ROADMs is optical layer restoration or re-optimization. Mesh restoration is an important protection mechanism for backbone networks. In a mesh protected network, failures are automatically re-routed over other available paths. The protection paths can be pre-defined, calculated dynamically at the time of the failure, or configured using a combination of both methods. Currently, mesh restoration is implemented at the electrical layer, either SONET (STS-1) or OTN.

The optical layer consists of physical fibers, lasers, and modulated signals running over those fibers. A whole host of optical impairments, such as optical loss, CD, and PMD must be managed on each optical span. At the physical layers, these optical impairments vary by fiber type, by span distance, and by overall optical path distance (i.e. OSNR). WDM vendors incorporate techniques to compensate for these optical impairments on each span and on the overall optical path. However, if the optical path changes, for example during optical layer protection switching, then the new path will have different optical parameters (loss, CD, PMD, OSNR). The coherent transponders have to "re-tune" and compensate for these new path parameters. In addition, optical layer control plane has to keep track of these optical parameters, for every route, to ensure that it can dynamically calculate a new path through the network when a failure occurs. This combination of optical layer path computation and control plane signaling can typically require 500ms – 2s to complete.

Optical layer re-optimization enables carriers to periodically clean up their network connections, recovering up to 20% of additional capacity. Large optical networks will incur many service additions and deletions. Over time, service churn results in fragments of stranded bandwidth scattered throughout a network. By re-optimizing their network connections, for example once every six months, the network can be "de-fragmented" and the stranded capacity recovered. Re-optimization is typically performed during a maintenance window, due to the slightly longer optical control plane path computation and signaling switch times.

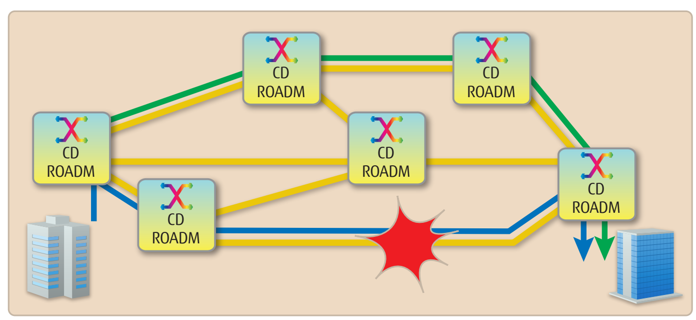

Figure 5: Optical Layer Mesh Restoration

Benefits	Issues
Network dynamically adjusts to faults	Optical layer restoration has many technical challenges: • Dissimilar optical impairments on protection path • Loss, CD, PMD, OSNR • Longer switch times (> 50 msec) • May require regeneration on protect path
Requires less protection bandwidth than 1+1	No standards for optical layer mesh protection • Control plane would have to pass optical layer impairment info • None currently under ITU study
Optical layer protection independent of any protocol technology	Most carriers implementing mesh restoration based on OTN layer • OTN layer restoration <50 msec • OTN is protocol-agnostic

Summary

CD ROADMs provide additional network flexibility by offering wavelength re-assignment and re-routing without the need for manual technician involvement. However, the complexity and cost of CD ROADMs are substantially higher than current generation ROADMs. Part 1 of this two-part series provides an overview of CD ROADMs and their three principal network applications: bandwidth pre-positioning, bandwidth on demand, and optical layer mesh restoration. While there are benefits to each application, there are also issues that may limit their use or return on investment. Part 2 of the series will provide an in-depth look at the architecture of CD ROADMs and their cost relative to existing ROADM technology. As an industry leader, Fujitsu is committed to supporting carriers whichever technology and architectural choices they make, however it is important for those decisions to be based on factual information and not industry hyperbole.

Acronyms	
2D	Two-Degree ROADM
8D	Eight-Degree ROADM
AWG	Arrayed Wave Guide
CD	Chromatic Dispersion
CD ROADM	Colorless Directionless ROADM
DWDM	Dense Wavelength Division Multiplexing
GUI	Graphical User Interface
ITU	International Telecommunication Union
OSNR	Optical Signal-to-Noise Ratio
OSS	Operations Support System
OTN	Optical Transport Network
PMD	Polarization Mode Dispersion
ROADM	Reconfigurable Optical Add Drop Multiplexer
ROI	Return on Investment
SONET	Synchronous Optical Network
WDM	Wavelength Division Multiplexing

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom