THE INTELLIGENT NETWORK: DYNAMICALLY MANAGING BANDWIDTH AT THE OPTICAL LEVEL

Gaylord A. Hart and Steven Robinson
Mahi Networks, Inc.

Abstract

CATV network bandwidth requirements are evolving at a rapid pace, driven by deployment of new services, increasing penetration of existing services, and the ongoing transition from analog to digital services. As services migrate to IP based and on-demand content delivery, bandwidth requirements vary dynamically in real time, and a minimum QoS is required to ensure adequate service performance. All of these place stringent requirements on the network.

Managing evolving & dynamic bandwidth requirements is complex, but emerging reconfigurable optical add/drop multiplexers (ROADMs) allow transport bandwidth to be managed effectively at the optical layer. Moreover, ROADMs enable fully automated optical transport systems which eliminate design and cut-over errors, accelerate service delivery, and lower network costs. This paper explores evolving services, optical transport technologies, and bandwidth management mechanisms in the context of migrating to an all-digital, IP based network that lowers both CapEx and OpEx costs.

INTRODUCTION

Bandwidth requirements in CATV networks are rapidly changing, driven by deployment of new services (VoIP and HDTV), increasing penetration of existing services (DTV and cable modems), and the ongoing transition from analog to digital and from circuits to packets. This picture is further complicated by the transition to on-demand services (with large variations between peak and average usage) and the migration to IP based services (which are connectionless, and therefore sometimes unpredictable in traffic requirements). Layer onto this the need to maintain QoS across a wide range of services, and bandwidth management becomes even more critical to tomorrow's cable network.

There are two key components to managing bandwidth in the future: the naturally changing bandwidth requirements as the network and service penetration evolve over a period of time and the real-time bandwidth management requirements imposed by everything-on-demand (EOD) and the converged IP network. To minimize long-term CapEx and OpEx costs (and hence remain competitive), MSOs must plan for these future requirements today and deploy systems that evolve with the network without forklift upgrades and without compromising service quality.

The basic HFC architecture is a long way from running out of bandwidth: increasing bandwidth in the future is relatively easy and inexpensive to accomplish simply by subdividing optical nodes (i.e., reducing the number of homes served per node and thereby increasing the bandwidth available per home), and the equipment for doing so exists today and is relatively inexpensive. Alternatively, a significant amount of network bandwidth can be made available by converting the broadcast analog TV services on the network to digital services (roughly a ten to one savings in bandwidth).

Because the HFC access architecture largely relies on transparent transport pipes, deploying new services on the network frequently does not require any upgrades to the HFC network itself, thus greatly lowering new service deployment costs and shortening time to market. These new services do, however, place new requirements on the transport and switching components in front of the HFC plant.

Over the next few years, the greatest changes in the CATV network will occur in front of the HFC access plant at those points in the network where services are aggregated, switched, and transported in the purely digital domain. For the MSO, this includes ISP and telephony POPs, content storage and origination points, regional and metro headends, primary and secondary hubs, and large businesses where services will be delivered directly via fiber.

A network's architecture is bounded by the optical transport paths connecting it, both physically and from a bandwidth perspective. These paths can be a bottleneck to delivering sufficient bandwidth and QoS as the network evolves. While it is common to engineer these paths to provide sufficient bandwidth under a pre-defined set of conditions, this can be costly, either requiring constant re-engineering to add incremental bandwidth as needed or resulting in stranded bandwidth if current traffic loads are far less than the deployed transport bandwidth.

This problem only worsens as services migrate to content on demand and IP based delivery. Both of these natural evolutionary steps, though very efficient in only requiring bandwidth when it is actually needed, result in bursty traffic, which makes traffic engineering even more complex. Needless to say, network complexity and traffic variability are only increasing, and new approaches are required for network design and operation in the future if we are to ensure the lowest network costs while maintaining an acceptable QoS.

Network engineers are used to designing optical transport paths as static circuits providing dedicated bandwidth to highly predictable traffic. As new CATV services are deployed, as penetration increases for existing services, and as services migrate to connectionless delivery via IP, new demands are being placed on the transport network and transport engineering. At the same time, rising customer expectations for service reliability are requiring more redundancy in the network. To keep up with these demands and to address emerging requirements for real-time bandwidth management (brought on by EOD and IP services), it is necessary the detailed optical transport network design, configuration, migration, and operational processes be automated, essentially masking much of the network complexity from the engineering and operational processes while enabling accelerated network evolution.

In its fullest implementation, such an optically reconfigurable network would dynamically modify itself in real time to respond to changing network conditions and service requirements. This would be analogous to the way a router automatically discovers paths through the network and dynamically routes packets to their destination based upon changing network conditions and traffic requirements. It now begins to make more sense to treat the optical transport layer (Layer 1 in the OSI model) as an extension of the switched layers residing above it, supporting integration of bandwidth management and QoS across the optical, Ethernet, and IP layers.

All of these objectives for optical transport automation and a dynamic optical layer can be accomplished with ROADMs using DWDM and GMPLS. When these technologies are coupled, an underlying optically reconfigurable network is possible which flexibly and dynamically supports future network evolution as well as bandwidth and content on demand, and without having to manually re-engineer and re-configure each component of the network when changes are made.

Presently, dense wave division optical add/drop multiplexers (DWDM OADMs) allow MSOs to collapse multiple parallel service transport networks onto a common optical network that supports multiple protocols and services over a single fiber. With modern network planning tools, the design of both the optical and service layers may be fully automated. This simplifies and speeds up network design, eliminates the need to memorize and understand complex design rules, and reduces design errors.

OADMs, however, typically rely on fixed configurations and components which are determined at the time of the initial design. This restricts the degree of automation which may be applied in the future to network evolution and slows down the upgrade process itself. ROADMs, because they allow their configuration to be controlled remotely via software, enable provisioning and operations to also be fully automated. This includes network topology discovery and service turn-up, as well as the dynamic monitoring and setting of network parameters for optimized operation. Modern ROADMs already support these capabilities today.

But as optical switching speeds increase and the GMPLS control plane becomes more content-aware and tightly coupled to Layers 2 and 3 above it, ROADM based CATV optical transport networks will become even more powerful, being able to dynamically re-configure transport paths and bandwidth in real time based upon changing service and content requirements, as well as subscriber on-demand service usage.

ROADM TECHNOLOGY

OADMs have been deployed widely in optical transport networks over the last few years. Most of these rely on fixed wavelength components (lasers, multiplexers, filters, etc.) which require a significant amount of manual network design as well as manual configuration and provisioning. A new class of fully reconfigurable OADMs has recently emerged which enables automation of these processes as well as automated and dynamic network operation. These ROADMs are built around flexible optical components which can be controlled via software and an intelligent control plane which supports process automation. The primary underlying technologies which define ROADMs are outlined below.

ROADM Architectures

Several architectural alternatives, based on a variety of markedly different optical technologies, exist today. Early ROADM technology, based on discrete optical-mechanical switches, filters and variable optical attenuators (VOAs), is shown in Figure 1, below.

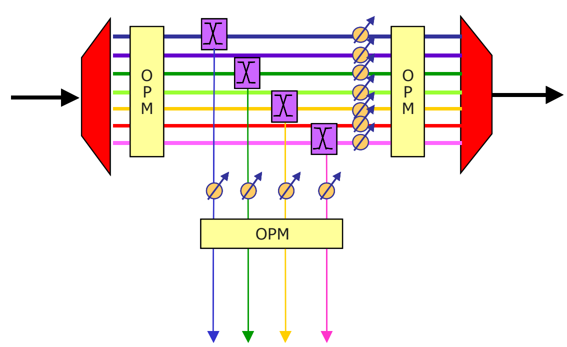

Figure 1. Discrete ROADM Architecture

While simple to implement, since most of the technology is commercially available, this approach utilizes many discrete optical components. The result is very high insertion loss, very high cost, and large size—thus preventing its widespread acceptance.

The second architecture to be used for ROADMs, the wavelength blocker architecture, is shown in Figure 2, below. Essentially, this design splits the incoming DWDM signal into a drop and through path. An integrated DWDM demultiplexer (Demux), VOA and multiplexer (Mux) form the core of the wavelength blocker. Typically blockers are implemented using Micro-Electrical Mechanical Systems (MEMS) or Liquid Crystal Display (LCD) technologies.

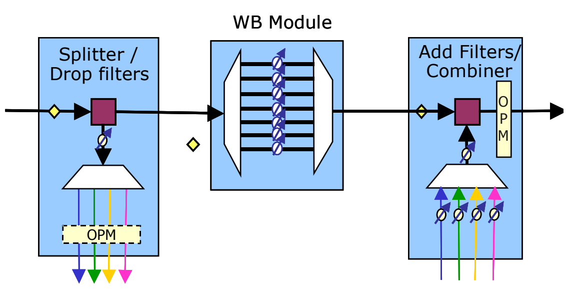

Figure 2. Wavelength Blocker Architecture

While this architecture reduces the number of discrete components, it forces the MSO to pay for all wavelengths at each node on day one. Furthermore, this design only manages the through wavelengths—not the add or drop wavelengths. In fact, most implementations of this architecture use inflexible fixed filters for the add and drop wavelengths. Hence this design is actually a "semi-reconfigurable" ROADM.

A variant of the wavelength blocker architecture is the combined blocker / adder architecture shown in Figure 3, below. This design integrates the add Mux with the wavelength blocker. This design eliminates the extra add Mux, but at the expense of requiring additional optical switches.

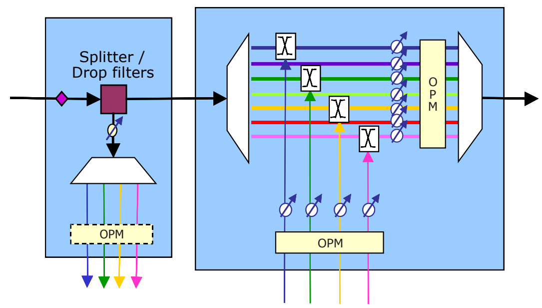

Figure 3. Combined Blocker / Adder ROADM

To be cost effective, this design typically integrates all of the filters and switches into a single module (similar to the wavelength blocker architecture). However, this again forces the MSO to pay for all wavelengths on day one. While this design now manages add wavelengths, it still does not manage drop wavelengths. Furthermore, this design now permanently locks the add wavelengths to a fixed wavelength design—thereby making fully tunable lasers only good for sparing (no dynamic wavelength management). Thus, this design is also a "semi-reconfigurable" ROADM.

Planar Lightwave Circuits (PLCs) are one technology used to implement this combined blocker / adder design. While they have the potential to integrate complex optical components onto one or more substrates (such as silicon), manufacturing yield and power management are still a challenge.

More recently, a new architecture based on Multi-Port Wavelength Selectable Switches (MP-WSS) allows for completely reconfigurable ROADM functionality. This architecture is shown in Figure 4, below. Wavelength switches have the ability to direct one or more wavelengths from an incoming DWDM signal to one or more output ports (usually with individual VOA-like power control for each wavelength).

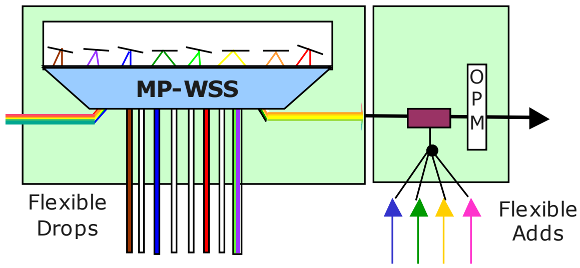

Figure 4. Fully Flexible WSS Based ROADM

Note that in this architecture, the MP-WSS is located on the drop side. In this configuration, the ROADM can manage both the drop and through wavelengths. It can also direct multiple wavelengths to a single drop port for low cost optical ring interconnection or physical mesh nodes. This design is inherently more efficient because the same demultiplexer is used for both through and drop wavelengths.

Another interesting aspect of this design is the use of broadband (wavelength independent) optical add ports. In this design, dynamic wavelength provisioning with full C-band tunable lasers is now possible. MP-WSS technology enables a "fully reconfigurable" ROADM architecture.

Optical Backplane

With the advent of the fully flexible ROADM (using the WSS), practical optical backplanes can now be implemented on an OADM chassis. Because wavelength selective switching eliminates the need for manually configured optical jumpers, an optical backplane permits the 100 or more manually placed front panel optical interconnect jumpers to be moved inside to the rear of the chassis. Reciprocal optical connectors on the back of each transponder allow these optical interconnects to be made when the transponder is plugged into the chassis. Without the WSS, however, an optical backplane would require all slots to use pre-defined wavelengths and thereby waste precious rack space.

An optical backplane frees MSOs from the tangled mess of jumpers interconnecting the lasers & receivers to the DWDM filters and amplifiers and from the need to carefully map and record these interconnects. Operators simply connect the client-side service fibers and walk away.

Tunable Lasers

The recent advent of widely tunable lasers has enabled an equivalent capability for full reconfigurability on the transmit side of the ROADM. Initially, MSOs were interested in tunables primarily for sparing purposes. This typically yields a 32:1 savings in spares and provides tremendous cost savings. More recently, MSOs have realized that tunables also substantially reduce delivery lead times and allow for significant equipment reuse.

More significantly, however, widely tunable lasers, when coupled with wavelength selective switches with single lambda add/drop granularity, enable greatly simplified provisioning as well as dynamic wavelength management. Combined with an optical backplane, remote provisioning and management are also possible. Operators no longer need to match fixed wavelength lasers with fixed wavelength filters using manually configured jumpers during installation. Instead, operators can simply point and click to remotely provision or re-provision wavelengths. This capability delivers tremendous OpEx and CapEx savings.

Optical Power Monitoring and Control

Initial OADM implementations provided only broadband, non-real-time optical power monitoring of the aggregate DWDM signal. However, to properly set up, manage, equalize, and optimize wavelengths, real-time optical power monitoring (OPM) and power level control are required for each lambda at each node. Real-time, direct measurements are more reliable and provide for rapid fault identification as well as support for newer protection schemes such as shared optical protection switching. Ideally, and for a fully robust system, all optical inputs, outputs, adds, and drops would be monitored in real time at each node for system optimization and full fault isolation.

Robust Variable Gain Amplifiers

Early DWDM amplifiers required operators to operate them in fixed gain mode. To ensure a flat gain profile, operators had to manually "pad-out" every span. This unnecessarily added noise to the DWDM signal—limiting reach. As new wavelengths were added or span losses changed, the system ran the risk of becoming unbalanced. Furthermore, any upstream wavelength changes (such as fiber cuts) could cause downstream errors (on both working and protection wavelengths).

Newer DWDM systems are now being implemented with variable gain amplifiers with transient control. Variable gain amplifiers can provide continuous, automated broadband gain adjustments to correct for slow changes in span losses or a change in the number of DWDM channels.

For rapid changes (such as a fiber cut), ultra-fast amplifiers can now provide robust transient control by quickly clamping the gain change to less than 0.3 dB in less than 1us. This accuracy is needed since the transient is additive with cascaded amplifiers. Combined with the ROADM capabilities already outlined above, both per-channel & broadband optical power level management can be fully automated and dynamically controlled in real time, thereby simplifying optical layer design and network management and operation.

GMPLS Control Plane

One final technology is required to fully automate the provisioning—an integrated control plane that supplies the intelligence and communication needed to control the ROADMs in the network. More recently, DWDM vendors have begun implementing GMPLS (Generalized Multi-Protocol Label Switching) control planes. Typically a dedicated wavelength carrying an Optical Supervisor Channel (OSC) is used to communicate the GMPLS messages between nodes. This wavelength is completely independent of the service bearing wavelengths and also supports provisioning communications from the Craft Interface and Element Management System (EMS).

GMPLS enables DWDM nodes to seamlessly work together in a network to provide resource and inventory auto-discovery, topology information, service setup, signaling, path computation (including Routing and Wavelength Assignments), light path setup, and management. Following industry standards, a GMPLS control plane enables true cross-network, A to Z provisioning.

BANDWIDTH MANAGEMENT

When wavelength selective switching, tunable lasers, and an optical backplane are combined in a ROADM along with a GMPLS control plane, a powerful platform is created with an extensive set of capabilities. Because of these capabilities, ROADMs are much more flexible in supporting service migration and network evolution and in automating these. For the following discussion, this paper assumes a ROADM has all the attributes listed above (i.e., genuinely is reconfigurable), though this is not always the case.

Network Design

Bandwidth management in the optical transport network begins with the initial network design. Limitations built-in to the initial design can limit future growth and network evolution, or at least make them much more painful and costly to achieve. Ideally, MSOs should plan for one-time network engineering, which allows unrestricted and non-disruptive addition or removal of services, wavelengths, and nodes to the network at any time.

The initial system design, for example, should readily support the migration (or mix and match capability) from 2.5G to 10G to 40G wavelengths without re-engineering and upgrading the network at each stage. Similarly, the ability to add or delete nodes in the network without re-engineering amplification or dispersion compensation is a significant advantage. ROADMs, because they are reconfigurable and can dynamically adapt to new network configurations, make this possible. ROADMs also further simplify and lower the cost of network evolution because they are remotely configurable, which in most cases eliminates the need to roll a truck to each node when an upgrade is performed.

Many modern software-based network planning tools allow the optical layer design process largely to be automated. This is a critical component in enabling flexible bandwidth management in the future. A good planning tool also allows the MSO to flexibly control design options (if necessary) and to optimize the design based upon selectable criteria such as cost or wavelength conservation.

However, if an OADM uses fixed filters or other components which are not dynamically compensated for by the network itself, these must manually be entered into the network design tool during the initial design and may require subsequent re-engineering of the network when upgrades are performed. By eliminating many of these fixed components or by being able to dynamically compensate for them, ROADMs simplify the initial network design process and any subsequent upgrades.

Network planning tools should also support automated design of the service layer as well, allowing an MSO to select protection options for common equipment and individual wavelengths, service and transponder types, client interfaces, and any other number of options. Once these are selected, the planning tool will automatically design the optical and service layers. At the optical layer, OSNR and dispersion compensation budgets are created, and the planning tool selects appropriate optical amplifiers and dispersion compensation modules to ensure one-time network engineering. At the service layer, common equipment and transponders are selected and assigned to specific chassis slots.

A well-designed planning tool, once it has provided a network design, will also provide detailed network drawings, span information, chassis drawings with slot assignments for each shelf component required, a bill of materials with part numbers, and even costs. The planning tool should also support exporting these items to an external spreadsheet for subsequent analysis, manipulation, or record keeping. Because the planning tool can reduce the time required to design a network to a matter of minutes, it is easy to create different network scenarios for comparative evaluation, allowing the MSO to optimize the network design based upon other service or architectural factors.

ROADMs, providing their underlying components have been designed to provide fast enough switching times, also enable more flexible protection options. Because wavelength selective switches and tunable lasers allow receive and transmit wavelengths to be assigned on the fly, protection wavelengths need not be assigned until an actual failure occurs. Such protection allows more than one path to share a common protection wavelength, which means more protected services may be placed on a given fiber. This conserves fiber while still providing complete service protection.

Good bandwidth management also includes conservation of network resources and performance optimization of these resources. In other words, you should use as few resources as necessary to get the job done, and you should get the most out of them. A well-designed ROADM conserves fiber and wavelengths on a fiber by eliminating stranded bandwidth caused by wavelength banding and by enabling unrestricted wavelength assignment and reallocation (single lambda granularity).

ROADMs which can dynamically optimize performance and optical parameters at each node also allow the typical engineering design rules to be relaxed, which enables transport over longer distances and through more nodes. All of these provide the MSO better bandwidth utilization and at a lower cost.

<u>Collapsed Service Transport Networks</u>

Most CATV networks and the services delivered over them have evolved so rapidly over the last few years that many MSOs have deployed parallel networks optimized for each service. While this made economic and technical sense at the time, it is not a sustainable model for a competitive market because it is costly to maintain, manage, operate, and upgrade these networks and because it makes inefficient use of network resources.

For most MSOs, the first stage in network evolution is collapsing these parallel networks onto common network elements and infrastructure wherever possible. Multi-service, multi-protocol ROADMs are ideal for collapsing the transport sections of these parallel networks onto a single, unified DWDM network (typically a ring) which provides fiber relief and service protection. Because ROADMs enable non-disruptive upgrades, follow-on network evolution can be transparently accomplished as new services are added and old services removed. Because ROADMs also enable flexible add/drop wavelength assignment, lambdas may constantly be recycled with service changes, allowing full utilization of existing fiber resources.

ROADMs with a GMPLS control plane offer another significant benefit, as well. Because GMPLS based ROADMs combine reconfiguration capabilities, intelligence, and communication between nodes, they also support automated network topology and inventory discovery. This eliminates the need to manually provision these into each network node and the EMS. At the same time, this capability allows the network to automatically turn up the individual wavelengths and services to be carried over it.

Another key aspect of bandwidth management for these unified networks is the ability to deploy only that bandwidth which is needed to support today's services, but at the same time to provide migration capacity and capability for tomorrow's bandwidth requirements. MSOs should not only expect a ROADM to support mixing and matching 2.5G and 10G wavelengths, but 40G wavelengths as well. While demand for 40G wavelengths may not be strong today, it will be in the future. In keeping with the one-time network engineering philosophy outlined above, this should be accomplished transparently, that is, without any additional upgrade costs or new restrictions on span budgets, node counts, or ring circumference. This pay-as-you-grow approach optimizes CapEx by tying network costs to service revenue, but it also eliminates or puts off into the future costly forklift upgrades and the deployment of new fiber.

Transport Network Evolution

Once any parallel single-service transport networks have been collapsed onto a common multi-service DWDM optical network, the next stage of network evolution becomes a matter of supporting bandwidth growth, service conversion, and new service deployment, as well as ensuring sufficient QoS is provided for all these services.

Bandwidth growth and new service deployment are relatively easy to provide with an OADM simply by adding transponders, providing sufficient lambdas are available to support the new transponders, or by migrating from 2.5G to 10G to 40G transponders, providing the network and the OADM have been designed to support these. ROADMs configured for one-time network engineering are ideal for supporting this type of growth. New transponders may be installed without any additional network engineering, and the ROADMs will dynamically adjust to the new traffic load. And unlike many OADMs with fixed serial filters or multiplexers, the ROADM permits transponders to be added without any service disruptions.

Seamless service conversion or migration is also readily supported by ROADMs. A good example of this is the migration from TDM voice to VoIP. In this case, new wavelengths are deployed to support the emerging IP service, while existing wavelengths continue to be used to support the legacy TDM service. As customers migrate to the new service, the legacy DWDM wavelengths may be torn down and recycled on the transport ring for new services. ROADMs enable flexible and unrestricted wavelength recycling and non-disruptive migration for existing services and customers.

As part of the bandwidth management role, OADMs may also be used to help manage QoS. Some MSOs have elected to segregate at various points in their networks, at least initially, services which have different QoS requirements and traffic characteristics. For example, IP data services and VoIP services, though both are IP services, have very different QoS requirements with respect to latency, jitter, and lost packets and very different traffic characteristics with respect to average to peak bit rates. It is possible to segregate these services in the network to maintain better control over QoS for each.

OADMs can be used effectively here for parallel transport of each service on a separate lambda over a single fiber, allowing unique bandwidth and protection options to be applied to each service. OADMs also typically support multiplexing transponders, which carry many multiplexed tributaries over a single wavelength and can therefore support this type of service segregation, as well. Of course, in the long run, the efficiencies of common transport and switching of all these services will require convergence. ROADMs can be of particular value here by enabling flexible and non-disruptive changes in the network to support dynamic service requirements and also the future transport convergence of these separate services when suitable QoS mechanisms (such as MPLS) are deployed in the network to guarantee QoS under worst-case conditions.

Converged IP Transport Network

A significant amount of equipment has already been deployed in CATV networks to support existing services and will not go away in the near future. Nevertheless, voice, video, and data services are destined for IP delivery. Indeed, data services in CATV networks are already there, with VoIP rapidly emerging, and IP video to follow.

As this transition occurs, MSOs will continue to operate parallel legacy service networks, as already outlined above, until such time comes when all customers are transitioned to IP based services and the legacy equipment is taken out of service. Multi-service, multi-protocol ROADMs can provide significant help in easing this transition while still allowing all services to be delivered over a common transport network. As IP services begin to dominate, legacy wavelengths can simply be recycled as no longer needed, freeing up additional transport capacity for IP services. However, significant challenges remain in deploying a fully unified IP network.

IP is a Layer 3 protocol and still requires transport by a Layer 2 protocol riding over a Layer 1 Physical Layer in the network. While most MSOs have deployed some ATM and SONET in there networks, it has typically not often been a significant amount, and MSOs are now focused on Ethernet for Layer 2 and Optical Ethernet for Layer 1. OADMs can readily and simultaneously transport all these protocols and services and can ease the transition as MSOs migrate their networks.

Ethernet is a wise choice for building CATV networks. Ethernet now scales effectively and inexpensively across LAN and WAN environments, is very easy to deploy and manage, and as a connectionless service is well-suited for IP, which is also connectionless. Recent Ethernet standards for virtual LANs (VLANs) make Ethernet even more powerful and support QoS and traffic prioritization capabilities in Ethernet for the first time.

When coupled with multi-protocol label switching (MPLS), which allows even tighter coupling between the Ethernet and IP layers, Ethernet clearly offers advantages that cannot be found with ATM or other protocols. Similarly, optical Ethernet, when coupled with GMPLS, offers advantages that cannot be found in SONET, which though widely deployed and quite capable of transporting Ethernet, is still a TDM transport technology.

As legacy services fade away and Ethernet and IP dominate, more wavelengths on the optical transport network will be dedicated to

Ethernet. While services may initially be segregated for transport for QoS reasons, in the long run this will not scale effectively, especially in light of modern VLAN Ethernet switches which also support MPLS and can prioritize traffic. So the network will not only migrate to Ethernet transport, but to multi-service transport in each pipe with each service provided sufficient transport bandwidth to ensure reliable performance.

Many benefits are still to be derived from the over-subscription that statistical multiplexing allows, especially when larger pipes can be used for transport of a larger number of services. So one can also expect a migration to wavelengths with ever more transport capacity (2.5G to 10G to 40G per lambda). Ultimately, it is also more cost effective to use fewer (but higher capacity) switch and router interfaces. MSOs should therefore, as they make transport decisions today, look for solutions that provide one-time network engineering and transport capability for 40G wavelengths.

As CATV networks migrate to IP based services, a parallel migration to on-demand services is also taking place. Broadband data access and VoIP services are already demand based services, for the most part only consuming bandwidth when the services are actively being used. Video on demand (VOD) is also widely deployed, but most CATV subscribers still receive their video services via broadcast analog or digital TV.

On-demand services have the benefit of only consuming network bandwidth when a service is requested by a subscriber, thus allowing statistical multiplexing at the service level. However, on-demand services can also result in a very high peak to average bandwidth usage ratio, which requires careful network planning to prevent network congestion and poor service delivery. This can be particularly true for video, which not only has relatively large bandwidth requirements, but which is also sensitive to latency and lost packets.

Because on-demand usage is also tied to customer preferences and other uncontrolled variables, the peak to average ratio may also vary significantly over time. For example, television viewing typically peaks in the evening, but there can also be significant traffic peaks created by special or unpredictable events. And this can be true across voice, video, and data services. While multicasting in an on-demand environment can mitigate the effects of these peaks for video, it cannot eliminate them.

Traffic planners are faced with a difficult choice: provide sufficient bandwidth for peak demand and let a significant amount of the bandwidth go unused most of the time, or provide less bandwidth and let the network congest under peak traffic loads. The first option is expensive, and the second results in poor service. Of course, most traffic planners try to hit a reasonable compromise, but with service usage varying dynamically over time (both in real-time and as a result of demographic trends), this is not always possible. Clearly, it would be desirable to have the network intelligently monitor its own loading and dynamically apply network resources where and when needed. Such an intelligent network would allow sharing a smaller number of network resources across a wider range of applications and conditions.

ROADMs, given their flexibility and on-the-fly reconfigurability, are entirely capable of supporting such a dynamically configured network in the transport domain. Under such a scenario, network planners could deploy a sufficient amount of transport bandwidth to cover average bandwidth requirements and provide additional uncommitted wavelengths which could be brought to bear when and where needed to prevent congestion as traffic demand builds. Load-sharing switches or routers would provide complementary bandwidth delivery functions at the Layer 2 and/or 3 levels.

Since network bottlenecks can occur in both the transport layer (insufficient bandwidth to carry the traffic) and switching layer (insufficient capacity to switch traffic), a dynamic transport network could provide additional wavelengths when transport capacity is exhausted or route existing wavelengths around congested switches when switching capacity is exhausted.

In this intelligent network, coordination between Layers 1, 2, and 3 would be implemented and automated by a GMPLS and MPLS control plane, which allows label-switched paths to be created and torn down as needed. Because the overall network status would then be known at all layers, optimization could be applied intelligently where it makes the most sense.

At the ROADM level, as congestion begins to build in the network, the GMPLS control plane would assign and route wavelengths between nodes on the ring, providing additional bandwidth in real time where and when it is needed. This mechanism could also be used to route traffic around a congested network segment or element, relieving traffic from congested or near-congested areas and applying it where sufficient resource are available to handle it. This approach can also be used to provide protection switching within the network by dynamically reassigning wavelengths and or transponders to compensate for equipment failures or fiber cuts.

Just as routing protocols dynamically allow routers to discover the most appropriate paths for moving datagrams through a network (and to compensate for failed routers), GMPLS and its associated routing protocols can route wavelengths through the optical transport network as needed to optimize network resources and capabilities. While current GMPLS and MPLS standards do not support this degree of integration across layers 1 2, and 3, this capability will no doubt be implemented at some point.

An intelligent network built upon these principles would require less equipment since resources are optimally applied only when and where needed and would simplify operations since network reconfiguration is done automatically and dynamically, and not manually by traffic planners. Overall, this would result in greater CapEx and OpEx savings for the MSO while providing greatly increased service reliability for subscribers.

SUMMARY AND CONCLUSIONS

CATV optical transport bandwidth requirements are constantly changing as a result of increasing service penetration, deployment of new services, and the transition to digital and IP based services. At the same time, on-demand services are rapidly being deployed and will likely displace broadcast services at some time. On-demand services significantly add to the complexity of traffic planning and bandwidth management. New demands are also being placed on the network to ensure adequate QoS is provided for each digital service to ensure reliable delivery. Service and network evolution is occurring at an ever-increasing rate, and new real-time constraints are being placed on the network which require dynamic optimization.

MSOs also face an increasingly competitive service environment which requires reliable and low-cost service delivery, rapid roll-out of new services, and the lowest possible CapEx and OpEx costs. To achieve these objectives, MSOs must deploy equipment today that optimizes performance and efficiency on the existing network infrastructure and which offers low first-in costs, yet can evolve with the network without forklift upgrades. This equipment must also provide a high degree of automation for network design, service turn-up, and operation, which in turn will result in lower costs, more reliable services, and accelerated service deployment and network migration.

ROADMs deliver the best approach to managing changing bandwidth requirements and evolution in the optical transport network. ROADMs support automated optical and service layer design, as well as one-time network engineering. They provide automated topology self-discovery and service turn-up, unrestricted wavelength usage and reconfigurability, and automated and dynamic monitoring and adjustment of network operational parameters for optimized performance. ROADMs are ideally suited for MSOs to meet their existing and future network transport requirements and also offer the most cost-effective solutions for optical transport, significantly lowering overall lifecycle CapEx and OpEx costs.

REFERENCES

"The Dynamic Network: Managing Bandwidth and Content on Demand at the Optical Level," Gaylord Hart and Zouheir Mansourati, Proceedings Manual: Collected Technical Papers of the 2004 SCTE Conference on Emerging Technologies.

Author Contact Information

Gaylord A. Hart
ghart@mahinetworks.com
303.910.7743

Steven Robinson
srobinson@mahinetworks.com
732.465.1000 ext. 1248