Angela Chiu
Celion Networks

John Strand
AT&T

Internet Draft
Document: draft-chiu-strand-unique-olcp-02.txt
Expiration Date: August 2001

Robert Tkach
Celion Networks

James Luciani
Tollbridge

Features and Requirements for The Optical Layer Control Plane

Status of this Memo

Abstract

Advances in the Optical Layer control plane are critical to ensure that the tremendous amount of bandwidth generated by DWDM technology be provided to upper layer services in a timely, reliable, and cost effective fashion. This document describes some unique features and requirements for the Optical Layer control plane that protocol designers need to take into consideration.

1. Introduction

The confluence of technical advances and service needs has focused intense interest on optical networking. Dense Wave Division Multiplexing (DWDM) is allowing unprecedented growth in raw optical bandwidth; new cross-connect technologies promise the ability to establish very high bandwidth connections within milliseconds; and the insatiable appetite of the Internet for high capacity ``pipesÆÆ has caused transport network operators to tear up their forecasts and add optical capacity as fast as they can.

Critical to these advances are improvements to the "Optical Layer Control Plane" -the software used to determine routings and establish and maintain connections. Traditional centralized transport operations systems (OSÆs) are widely acknowledged to be incapable of scaling to meet exploding demand or establishing connections as rapidly as needed. Consequently much attention has been paid recently to new control plane architectures based on data networking protocols such as MPLS and OSPF/IS-IS), under the Generalized MPLS (GMPLS) umbrella. These architectures feature distributed routing and control logic, auto discovery and self inventorying, and many other advantages. OSPF/IS-IS provides a constraint-based routing capability that takes bandwidth availability into account.

The potential of these new architectures for optical networking is enormous; however, to be successful they need to be adapted to the specific technological, service, and business context characteristic of optical networking. In order to identify some of the enhancements necessary in GMPLS to make it applicable to optical networks, this document attempts to describe several aspects of optical networking which differ from those in the data networking environment inspiring these new architectures:

- Section 2 describes some distinctive technological and networking aspects of optical networking that will constrain routing in an optical network, and

- Section 3 gives a transport network operatorÆs perspective on business and operational realities that optical networks are likely to face which are unlike those in data networking.

Particular emphasis is placed on the multihop optical network problem. We most definitely are not claiming that these differences are fatal to these new architectures, only that the new architectures must be built upon a detailed appreciation of the unique characteristics of the optical world.

2. Constraints On Routing

Optical Layer routing is less insulated from details of physical implementation than routing in higher layers. In this section we give examples of constraints arising from the design of network elements, from the accumulation of signal impairments, and from the need to guarantee the physical diversity of some circuits.

2.1 Reconfigurable Network Elements

Control plane architectural discussions (e.g., [Awduche99]) usually assume that the only software reconfigurable network element is an optical layer cross-connect (OLXC). There are however other software reconfigurable elements on the horizon, specifically tunable lasers and receivers and reconfigurable optical add-drop multiplexers (OADMÆs). These elements are illustrated in the following simple example, which is modeled on announced Optical Transport System (OTS) products:

```
            +                                       +
---+---+    |\                                     /|    +---+---
---| A |----|D|          X              Y         |D|----| A |---
---+---+    |W|     +---------+    +---------+    |W|    +---+---
     :      |D|-----|  OADM   |----|  OADM   |----|D|      :
---+---+    |M|     +---------+    +---------+    |M|    +---+---
---| A |----| |      |      |       |      |      | |----| A |---
---+---+    |/       |      |       |      |       \|    +---+---
            +      +---+  +---+   +---+  +---+      +
             D     | A |  | A |   | A |  | A |      E
                   +---+  +---+   +---+  +---+
                    | |    | |     | |    | |
```

Figure 2-1: An OTS With OADM's - Functional Architecture

In Fig.2-1, the part that is on the inner side of all boxes labeled "A" defines an all-optical subnetwork. From a routing perspective two aspects are critical:

- Adaptation: These are the functions done at the edges of the subnetwork that transform the incoming optical channel into the physical wavelength to be transported through the subnetwork.
- Connectivity: This defines which pairs of edge Adaptation functions can be interconnected through the subnetwork.

In Fig. 2-1, D and E are DWDMÆs and X and Y are OADMÆs. The boxes labeled "A" are adaptation functions. They map one or more input optical channels assumed to be standard short reach signals into a long reach (LR) wavelength or wavelength group which will pass transparently to a distant adaptation function. Adaptation functionality which affects routing includes:

- Multiplexing: Either electrical or optical TDM may be used to combine the input channels into a single wavelength. This is done to increase effective capacity: A typical DWDM might be able to handle 100 2.5 Gb/sec signals (250 Gb/sec total) or 50 10 Gb/sec (500 Gb/sec total); combining the 2.5 Gb/sec signals together thus effectively doubles capacity. After multiplexing the combined signal must be routed as a group to the distant adaptation function.
- Adaptation Grouping: In this technique, groups of k (e.g., 4) wavelengths are managed as a group within the system and must be added/dropped as a group. We will call such a group an "adaptation grouping". Another term frequently used is "wave group".
- Laser Tunability: The lasers producing the LR wavelengths may have a fixed frequency, may be tunable over a limited range, or be tunable over the entire range of wavelengths supported by the DWDM. Tunability speeds may also vary. Note that tunable receivers are becoming a reality and for certain applications.

Connectivity between adaptation functions may also be limited:

- As pointed out above, TDM multiplexing and/or adaptation grouping by the adaptation function forces groups of input channels to be delivered together to the same distant adaptation function.
- Only adaptation functions whose lasers/receivers are tunable to compatible frequencies can be connected.
- The switching capability of the OADMÆs may also be constrained. For example:
    - o There may be some wavelengths that can not be dropped at all.
    - o There may be a fixed relationship between the frequency dropped and the physical port on the OADM to which it is dropped.
    - o OADM physical design may put an upper bound on the number of adaptation groupings dropped at any single OADM.

For a fixed configuration of the OADMÆs and adaptation functions connectivity will be fixed: Each input port will essentially be hard-wired to some specific distant port. However this connectivity can be changed by changing the configurations of the OADMÆs and adaptation functions. For example, an additional adaptation grouping might be dropped at an OADM or a tunable laser retuned. In each case the port-to-port connectivity is changed.

This capability can be expected to be under software control. Today the control would rest in the vendor-supplied Element Management system (EMS), which in turn would be controlled by the operatorÆs OSÆs. However in principle the EMS could participate in the routing process. The constraints on reconfiguration are likely to be quite complex, dependent on the vendor design and also on exactly what line cards have been deployed. Thus the state information needed for routing is likely to be voluminous and possibly vendor specific. However it is very desirable to solve these issues, possibly by advertising only an abstraction of the complex configuration options to the external world via the control plane.

2.2 Wavelength Routed All-Optical Networks

The optical networks presently being deployed may be called "opaque" ([Tkach98]): each link is optically isolated by transponders doing O/E/O conversions. They provide regeneration with retiming and reshaping, also called 3R, which eliminated transparency to bit rates and frame format. These transponders are quite expensive and they also constrain the rapid evolution to new services - for example, they tend to be bit rate and format specific. Thus there are strong motivators to introduce "domains of transparency" - all-optical subnetworks - larger than an OTS.

The routing of lightpaths through an all-optical network has received extensive attention. (See [Yates99] or [Ramaswami98]). When discussing routing in an all-optical network it is usually assumed that all routes have adequate signal quality. This may be ensured by limiting all-optical networks to subnetworks of limited geographic size which are optically isolated from other parts of the optical layer by transponders. This approach is very practical and has been applied to date, e.g. when determining the maximum length of an Optical Transport System (OTS). Furthermore operational considerations like fault isolation also make limiting the size of domains of transparency attractive.

There are however reasons to consider contained domains of transparency in which not all routes have adequate signal quality. From a demand perspective, maximum bit rates have rapidly increased from DS3 to OC-192 and soon OC-768 (40 Gb/sec). As bit rates increase it is necessary to increase power. This makes impairments and nonlinearities more troublesome. From a supply perspective, optical technology is advancing very rapidly, making ever-larger domains possible. In this section we assume that these considerations will lead to the deployment of a domain of transparency that is too large to ensure that all potential routes have adequate signal quality for all circuits. Our goal is to

understand the impacts of the various types of impairments in this environment.

2.2.1 Problem Formulation

We consider a single domain of transparency. We wish to route a unidirectional circuit from ingress client node X to egress client node Y. At both X and Y, the circuit goes through an O/E/O conversion which optically isolates the portion within our domain. We assume that we know the bit rate of the circuit. Also, we assume that the adaptation function at X applies some Forward Error Correction (FEC) method to the circuit. We also assume we know the launch power of the laser at X.

Impairments can be classified into two categories, linear and nonlinear (See [Tkach98] for more on impairment constraints). Linear effects are independent of signal power and affect wavelengths individually. Amplifier spontaneous emission (ASE), polarization mode dispersion (PMD), and chromatic dispersion are examples. Nonlinearities are significantly more complex: they generate not only dispersion on each channel, but also crosstalk between channels.

In the remainder of this section we first outline how two key linear impairments (PMD and ASE) might be handled by a set of analytical formulae as additional constraints on routing. We next discuss how the remaining constraints might be approached. Finally we take a broader perspective and discuss the implications of such constraints on control plane architecture and also on broader constrained domain of transparency architecture issues.

2.2.2 Polarization Mode Dispersion

For a transparent fiber segment, the general rule for the PMD requirement is that the time-average differential time delay between two orthogonal state of polarizations should be less than a% of the bit duration. (A typical value for a is 10 [ITU]. More aggressive designs to compensate for PMD may allow higher than 10%. This would be a system parameter known to the routing process.) This results in a upper bound on the maximum length of an M-fiber-span transparent segment, which is inverse proportion to the square of bit rate and fiber PMD parameter where a fiber span in a transparent network refers to a segment between two optical amplifiers (The detailed equation is omitted due to the format constraint). For typical fibers with PMD parameter of 0.5 picosecond per square root of km, based on the constraint, the maximum length of the transparent segment should not exceed 400km and 25km for bit rates of 10Gb/s and 40Gb/s, respectively. With newer fibers assuming PMD parameter equals to 0.1 picosecond per square root of km, the maximum length of the transparent segment should not exceed 10000km and 625km for bit rates of 10Gb/s and 40Gb/, respectively. In general, the PMD requirement is not an issue for most types of fibers at 10Gb/s or lower bit rate. But it will become an issue at bit rates of 40Gb/s and higher.

2.2.3 Amplifier Spontaneous Emission

ASE degrades the signal to noise ratio. An acceptable optical SNR level (SNRmin) which depends on the bit rate and transmitter-receiver technology (e.g., FEC) needs to be maintained at the receiver. It also includes all the margins for those impairments that are not treated explicitly as described in the following subsection. In order to satisfy this requirement, vendors often provide some general engineering rule in terms of maximum length of the transparent segment and number of spans. For example, current transmission systems are often limited to up to 6 spans with 80km long in each. Startups have announced ultra long haul systems that are claimed to be able to support up to thousands of km. Although these general rules are helpful in network planning, more detailed information on the SNR reduction in each component should be used to determine whether the SNR level through a given transparent segment is within the required value. This would provide flexibility in provisioning or restoring a lightpath through a transparent subnetwork. Here, we assume that the average optical power launched at the transmitter is known as P. The lightpath from the transmitter to the receiver goes through M optical amplifiers, with each introducing some noise power. Unity gain can be used at all amplifier sites to maintain constant signal power at the input of each span to minimize noise power and nonlineararity. A constraint on the maximum number of spans can be obtained [Kaminow97] which is proportional to P and inverse proportional to SNRmin, optical bandwidth B, amplifier gain G-1 and spontaneous emission factor n of the optical amplifier. (Again, the detailed equation is omitted due to the format constraint.) LetÆs take a typical example. Assuming P=4dBm, SNRmin=20dB with FEC, B=12.5GHz, n=2.5, G=25dB, based on the constraint, the maximum number of spans is at most 10. However, if FEC is not used and the requirement on SNRmin becomes 25dB, the maximum number of spans drops down to 3.

2.2.4 Other Impairments

Other Polarization Dependent Impairments Other polarization-dependent effects besides PMD influence system performance. For example, many components have polarization-dependent loss (PDL) [Ramaswami98] which accumulates in a system with many components on the transmission path. The state of polarization fluctuates with time, and it is generally required to maintain the total PDL on the path to be within some acceptable limit, typically 1dB margin in OSNR.

Chromatic Dispersion In general this impairment can be adequately (but not optimally) compensated for on a per-link basis, and/or at system initial setup time.

Crosstalk Since crosstalk in the system affects Q which is a measure of the electrical signal-to-noise ratio assuming Gaussian noise statistics, it can be factored in with some margin in Q. As a result, one can increase the OSNR requirement by some modified margin.

Nonlinear Impairments It seems unlikely that these can be dealt with explicitly in a routing algorithm because they lead to constraints that can couple routes together and lead to complex dependencies, e.g. on the order in which specific fiber types are traversed. Note that different fiber types (standard single mode fiber, dispersion shifted fiber, dispersion compensated fiber, etc.) have very different effects from nonlinear impairments. A full treatment of the nonlinear constraints would likely require very detailed knowledge of the physical infrastructure, including measured dispersion values for each span, fiber core area and composition, as well as knowledge of subsystem details such as dispersion compensation technology. This information would need to be combined with knowledge of the current loading of optical signals on the links of interest to determine the level of nonlinear impairment. Alternatively, one could assume that nonlinear impairments are bounded and result in X dB margin in the required OSNR level for a given bit rate, where X for performance reasons would be limited to 1 or 2 dB, consequently setting a limit on the maximum number of spans. For the approach described here to be useful, it is desirable for this span length limit to be longer than that imposed by the constraints which can be treated explicitly. When designing a DWDM transport system, there are tradeoffs between signal power launched at the transmitter, span length, and nonlinear effects on BER that need to be considered jointly. Here, we assume that an X dB margin is obtained after the transport system has been designed with a fixed signal power and maximum span length for a given bit rate. Further work is required to determine the validity of this approach. However, it is possible that there could be an advantage in designing systems which are less aggressive with respect to nonlinearities, and therefore somewhat sub-optimal, in exchange for improved scalability, simplicity and flexibility in routing and control plane design.

2.2.5 Implications For Routing and Control Plane Design

- Additional state information will be required by the routing algorithm for each type of impairment that has the potential of being limiting for some routes.

- It is likely that the physical layer parameters do not change value rapidly and could be stored in some database; however these are physical layer parameters that today are frequently not known at the granularity required. If the ingress node of a lightpath does path selection these parameters would need to be available at this node.

- The specific constraints required in a given situation will depend on the design and engineering of the domain of transparency; for example it will be important to know whether chromatic dispersion has been dealt with on per-link basis, and whether the domain is operating in a linear or nonlinear regime.

- In situations where only PMD and/or ASE impairments are potentially binding the optimal routing problem as two constraints OSPF algorithm enhancements will be needed. However, it is likely that relatively simple heuristics could be used in practice.

Additionally, routing in an all-optical network without wavelength conversion raises several additional issues:

- Since the route selected must have the chosen wavelength available on all links, this information needs to be considered in the routing process. This is discussed in [Chaudhuri00], where it is concluded that advertising detailed wavelength availabilities on each link is not likely to scale. Instead they propose an alternative method which probes along a chosen path to determine which wavelengths (if any) are available. This would require a significant addition to the routing logic normally used in OSPF.

- Choosing a path first and then a wavelength along the path is known to give adequate results in simple topologies such as rings and trees ([Yates99]). This does not appear to be true in large mesh networks under realistic provisioning scenarios, however. Instead significantly better results are achieved if wavelength and route are chosen simultaneously. This approach would however also have a significant affect on OSPF.

2.3 Diversity

"Diversity" is a relationship between lightpaths. Two lightpaths are said to be diverse if they have no single point of failure. In traditional telephony the dominant transport failure mode is a failure in the interoffice plant, such as a fiber cut inflicted by a backhoe.

Why is diversity a unique problem that needs to be considered for optical networks? So far, data network operators have relied on their private line providers to ensure diversity and so have not had to deal directly with the problem. GMPLS makes the complexities handled by the private line provisioning process, including diversity, part of the common control plane and so visible to all.

To determine whether two lightpath routings are diverse it is necessary to identify single points of failure in the interoffice plant. To do so we will use the following terms: A fiber cable is a uniform group of fibers contained in a sheath. An Optical Transport System will occupy fibers in a sequence of fiber cables. Each fiber cable will be placed in a sequence of conduits - buried honeycomb structures through which fiber cables may be pulled - or buried in a right of way (ROW). A ROW is land in which the network operator has the right to install his conduit or fiber cable. It is worth noting that for economic reasons, ROWÆs are frequently obtained from railroads, pipeline companies, or thruways. It is frequently the case that several carriers may lease ROW from the same source; this makes it common to have a number of carriersÆ fiber cables in close proximity to each other. Similarly, in a metropolitan network, several carriers might be leasing duct space in the same RBOC conduit. There are also "carrier's carriers" - optical networks which provide fibers to multiple carriers, all of whom could be affected by a single failure in the "carrier's carrier" network.

In a typical intercity facility network there might be on the order of 100 offices that are candidates for OLXCÆs. To represent the inter-office fiber network accurately a network with an order of magnitude more nodes is required. In addition to Optical Amplifier (OA) sites, these additional nodes include:

- Places where fiber cables enter/leave a conduit or right of way;
- Locations where fiber cables cross;
- Locations where fiber splices are used to interchange fibers between fiber cables.

An example of the first might be:

```
                             A                 B
A-------------B                     \             /
                                      \         /
                                        X-----Y
                                      /         \
                                    /             \
                                  C                 D
C-------------D
```

```
(a) Fiber Cable Topology      (b) Right-Of-Way/Conduit Topology
```

Figure 2-2: Fiber Cable vs. ROW Topologies

Here the A-B fiber cable would be physically routed A-X-Y-B and the C-D cable would be physically routed C-X-Y-D.   This topology might arise because of some physical bottleneck: X-Y might be the Lincoln Tunnel, for example, or the Bay Bridge.

Fiber route crossing (the second case) is really a special case of this, where X and Y coincide. In this case the crossing point may not even be a manhole; the fiber routes might just be buried at different depths.

Fiber splicing (the third case) often occurs when a major fiber route passes near to a small office. To avoid the expense and additional transmission loss only a small number of fibers are spliced out of the major route into a smaller route going to the small office. This might well occur in a manhole or hut. An example is shown in Fig. 2-3(a), where A-X-B is the major route, X the manhole, and C the smaller office. The actual fiber topology would then look like Fig. 2-3(b), where there would typically be many more A-B fibers than A-C or C-B fibers, and where A-C and C-B might have different numbers of fibers. (One of the latter might even be missing.)

```
       C                              C
       |                            /   \
       |                          /       \
       |                        /           \
A------X------B               A---------------B
```

```
(a) Fiber Cable Topology         (b) Fiber Topology
```

Figure 2-3. Fiber Cable vs Fiber Topologies

The imminent deployment of ultra-long (>1000 km) Optical Transport Systems introduces a further complexity: Two OTS's could interact a number of times. To make up a hypothetical example: A New York - Atlanta OTS and a Philadelphia - Orlando OTS might ride on the same right of way for x miles in Maryland and then again for y miles in Georgia. They might also cross at Raleigh or some other intermediate node without sharing right of way.

Diversity is often equated to routing two lightpaths between a single pair of points, or different pairs of points so that no single route failure will disrupt them both. This is too simplistic, for a number of reasons:

- A sophisticated client of an optical network will want to derive diversity needs from his/her end customers' availability requirements. These often lead to more complex diversity requirements than simply providing diversity between two lightpaths. For example, a common requirement is that no single failure should isolate a node or nodes. If a node A has single lightpaths to nodes B and C, this requires A-B and A-C to be diverse. In real applications, a large data network with N lightpaths between its routers might describe their needs in an NxN matrix, where (i,j) defines whether lightpaths i and j must be diverse.

- Two circuits that might be considered diverse for one application might not be considered diverse for in another situation. Diversity is usually thought of as a reaction to interoffice route failures. High reliability applications may require other types of failures to be taken into account. Some examples:
  - o Office Outages: Although less frequent than route failures, fires, power outages, and floods do occur. Many network managers require that diverse routes have no (intermediate) nodes in common. In other cases an intermediate node might be acceptable as long as there is power diversity within the office.
  - o Shared Rings: Many applications are willing to allow "diverse" circuits to share a SONET ring-protected link; presumably they would allow the same for optical layer rings.
  - o Disasters: Earthquakes and floods can cause failures over an extended area. Defense Department circuits might need to be routed with nuclear damage radii taken into account.

- Conversely, some networks may be willing to take somewhat larger risks. Taking route failures as an example: Such a network might be willing to consider two fiber cables in heavy duty concrete conduit as having a low enough chance of simultaneous failure to be considered "diverse". They might also be willing to view two fiber cables buried on opposite sides of a railroad track as being diverse because there is minimal danger of a single backhoe disrupting them both even though a bad train wreck might jeopardize them both.

These considerations strongly suggest that the routing algorithm should be sensitive to the types of threat considered unacceptable by the requester. Note that the impairment constraints described in the previous section may eliminate some of the long circuitous routes sometimes needed to provide diversity. This would make it harder to find many diverse paths through an all-optical network than an opaque one.

[Chaudhuri00] introduced the term "Shared Risk Link Group" (SRLG) to describe the relationship between two non-diverse links. The above discussion suggests that an SRLG should be characterized by 2 parameters:

- Type of Compromise: Examples would be shared fiber cable, shared conduit, shared ROW, shared optical ring, shared office without power sharing, etc.)
- Extent of Compromise: For compromised outside plant, this would be the length of the sharing.

Two links could be related by many SRLG's (AT&T's experience indicates that a link may belong to over 100 SRLG's, each corresponding to a separate fiber group. Each SRLG might relate a single link to many other links. For the optical layer, similar situations can be expected where a link is an ultra-long (3000 km) OTS). The mapping between links and different types of SRLGÆs is in general defined by network operators based on the definition of each SRLG type. Since SRLG information is not yet ready to be discoverable by a network element and does not change dynamically, it need not be advertised with other resource availability information by network elements. It could be configured in some central database and be distributed to or retrieved by the nodes, or advertised by network elements at the topology discovery stage. On the other hand, in order to be able to perform distribute path selection at each node that satisfies certain diverse routing criterion, each network element may need to propagate the information of number of channels available for each channel type (e.g., OC48, OC192) on each channel group, where channel group is defined as a set of channels that are routed identically and should be given unique identification. Each channel group can be mapped into a sequence of fiber cables while each fiber cable can belong to multiple SRLGÆs based on their definitions.

2.4 Other Unique Features of Optical Networks

There are other major differences between optical networks and IP networks that have significant impacts on the design of the Optical Layer control plane. They include the following two areas.

- Bi-directionality: In an IP network, Label Switched Paths (LSPs) are inherently unidirectional. However, current transport networks are bi-directional oriented, mostly due to the evolution of two-way transmission in Public Switched Telephone Network and by SONET/SDH line protection schemes [Doverspike00]. This often requires the bi-directional connections provided by the optical layer to use the same numbered channel in each direction. As a result, a channel contention problem may occur between two bi-directional request traveling in opposite directions. Signaling mechanisms have been proposed to resolve this type of contention [Ashwood00].

- Protection and restoration: In an IP network, when a backup LSP is pre-established to protect against failure(s) on a working LSP, the backup LSP does not occupy any physical resources before a failure occurs. However, in an optical network, a pre-established optical connection for backup does occupy the ports and channels on the path of the connection. This can be used for the 1+1 protection, but not for shared mesh protection. Instead with shared mesh protection, the backup path can be pre-selected with or without the associated channels being chosen prior to any failure, then cross-connect ports/channels physically after a failure on the working path has been detected. See [Doverspike00] for more detailed discussions on various protection/restoration schemes.

2.4 Discussion and Summary

Dealing with diversity seems to be an unavoidable requirement on optical layer routing. It requires dealing with additional constraints in the routing process but most importantly requires additional state information to be available to the routing process.

The physical constraints of optical technology apply inside an all-optical ``domains of transparencyÆÆ. TodayÆs OTS is a simple ``domain of transparencyÆÆ consisting of WDM Mux/Demuxers and Optical Amplifiers. Because an OTS is not easily reconfigurable these constraints are dealt with at the time of installation and donÆt complicate routing and the control plane.

As domains of transparency become both larger and software reconfigurable as discussed earlier, these physical constraints on connectivity and transmission quality become increasingly of concern to the control plane. It is important to note that at present this evolution is largely technology driven: vendors pushing the technology envelope are competing fiercely to provide solutions which have higher capacity, can go further all-optically, are more reconfigurable, and are more cost-effective. Routing constraints, which are essentially a by-product of this competitive dynamic, may well become more complex. As vendors pursue their diverse visions it is quite plausible that the optical layer of the future will be made up of heterogeneous technologies which differ significantly in their routing implications.

What are the control plane architecture choices in such an eventuality? Alternative approaches that deserve consideration are:

- Per-Domain Routing: In this approach each domain could have its own tuned approach to routing. Inter-domain routing would be handled by a multi-domain or hierarchical protocol that allowed the hiding of local complexity. Single vendor domains might have proprietary intra-domain routing strategies. This approach has the advantage of providing carriers a flexible way to digest technologies and subnetworks that for whatever reason are not using a standard intra-domain routing protocol.

- Enforced Homogeneity: The capabilities of the control plane would impose constraints on system design and network engineering. As examples: If control plane protocols did not deal with non-linear impairments carriers would require their vendors to provide transport systems where these constraints were never binding. Transmission engineers could be required to only deploy domains where every possible route met all constraints not handled explicitly by the control plane even if the cost penalties were severe.

- Additional Regeneration: At (selected) OLXCÆs within a domain of transparency, the control plane could insert O/E/O regeneration into routes with transmission problems. This might make all routes feasible again, but at the cost of additional cost and complexity and with some loss of rate and format transparency.

- Standardized Intra-Domain Routing Protocol: The examples discussed in Section 2 suggest that a single standardized protocol which tries to deal with the full range of possible topological and transmission constraints will be extremely complex and will require a lot of state information. However when combined with limited application of the two previous approaches it might be more plausible.

Given the complexity of physical and connectivity impairments and diversity requirements, a valid question to ask is whether a centralized routing model, where routing is done centrally using a centralized database with a global network view would be better than the distributed model favored in the Internet. Here, we provide some pros and cons on each model.

To the extent that the per-domain routing approach just discussed is used, the choice of model might be different depending on the characteristics of the domain. For example, in a domain like Fig. 2-1 it seems likely that a centralized model is more appropriate because network elements like tunable lasers and reconfigurable OADM's seem on the surface to be unlikely peers to much more complex devices like OXC's or routers. On the other hand, a purely "opaque" domain where impairment constraints play no role in routing would appear to be an excellent candidate for the distributed model.

In the context of the complexities discussed in this paper, a centralized model has some advantages:

- Information such as SRLGÆs and performance parameters which change infrequently and are unlikely to be amenable to self-discovery could reside in a central database and would not need to be advertised.

- Routing dependencies among circuits (to ensure diversity, for example) is more easily handled centrally when the circuits do not share terminals since the necessary state information should be more easily accessible in a centralized model.

- Pre-computation of restoration paths and other computations that can benefit from the use of global state information may also benefit from centralization.

There are, of course, significant disadvantages to the centralized model when compared to a distributed model:

- If rapid restoration is required, it is not possible to rely on a centralized routing system to compute a recovery path for each failed lightpath on demand after a failure has been detected. The distributed model arguably will not have this problem.

- The centralized approach is not consistent with the distributed routing philosophy prevalent in the Internet. The reasons which drove the InternetÆs architecture-scalability, the inherent problems with hard state information, etc.-are largely relevant to optical networking. In addition there is the major disadvantage that a centralized approach would seem to preclude integrated routing across the IP and optical boundary.

A related issue is whether routes should be pre-computed. It has been suggested, for example, that all routes (or at least a large number) be pre-computed and stored in a central database. This potentially might allow more sophisticated algorithms to be used to filter out the routes violating transmission constraints. There are however serious disadvantages (in addition to the disadvantages of the centralized model given above):

- In a large national network there are just too many routes that might be needed, by orders of magnitude. This is particularly true when diversity constraints and restoration routing may force weird routings.
- Every time any parameter changes anywhere in the network all routes using the impacted resource will need to be reexamined.

3. Business and Operational Realities

The Internet technologies being applied to define the new Optical Layer control plane evolved in a very different business and operational environment than that of today's transport network provider. The differences need to be clearly understood and dealt with if the new control plane is going to be a success. The Optical Interworking Forum, one of the principal standards groups in this area, has recently formed a Carrier Subgroup to provide guidance from this perspective for their standards activities.

In this section we touch on two aspects of this problem: Business Models and the management of the introduction of new technology.

3.1 Business Models

The cost of providing gigabit connections is expected to drop rapidly, but will still require dedicated use of expensive and periodically scarce capacity and equipments. Therefore the ability to control network access, and to measure and bill for usage, will be critical. Also, lightpath connections are expected to have quite long holding times (weeks-months) compared to LSPs in an IP network. Therefore the collection of usage data and the nature of the connection establishment process have very different characteristics in the Optical Network than in an IP network.

In addition, industry revenues from legacy services (voice and private line) are expected to dwarf those from IP transport for the next few years. ?? John: could you find a reference from RHK? Meeting the needs of these services and migrating them to the operatorÆs newer service platforms will also be a critical need for operators with extensive embedded revenues. Thus the needs of services based on SONET/SDH, Ethernet, ATM, etc. will need to be given attention. In addition most operators hope that they will have many different ISP's and Intranets as customers. Thus the customer base for most operators will be quite diverse.

Another area of prime concern is Operations Systems (OSÆs). The opportunity to create a thinner and more nimble network management plane by off-loading many provisioning and data-basing functions onto a vendor-provided control plane and/or Element Management System (EMS) holds the promise of large and immediate benefits to operators in the form of reduced software development and more rapid deployment of new functionality. This is a critical area to achieve scalability.

In the short term the principal benefits of the proposed control plane are two: rapid provisioning and a reduction in the cost and complexity of OSÆs and operations. Both of these benefits require that circuits be controlled end-to-end by the new control plane, for otherwise the provisioning times will be determined by those of the older, much slower segments and OS costs and OS and operations complexity may actually go up because of the need to interwork the old and the new worlds. To avoid this the capabilities of the new control plane need to be available end-to-end as soon as possible. This will put a premium on the rapid development of standards for interworking across trust boundaries, for example between Local Exchange Carrier's and national networks.

3.2 Managing The Introduction Of New Technology

We expect optical layer hardware technology to continue to evolve very rapidly, with a very real possibility of additional "disruptive" advances. The analog nature of optical technology compounds this problem for the control planes because these advances are likely to be accompanied by complex technology-specific constraints on routing and functionality. (Sections 2.1 and 2.2 above provide examples of this.) An architecture which allows the gradual and seamless introduction of new technologies into the network without time-consuming and costly changes to embedded technologies and especially control planes is highly desirable.

When compared to the IP experience several distinctions stand out:

- The optical layer control plane seems more likely to be buffeted by hardware changes than is the IP control plane.
- Optical layer innovations are currently being driven by start-up companies, with product innovation well ahead of the standards process. Efforts at control plane standardization are much less mature than comparable IP efforts. This is a matter of considerable concern because neither rapid provisioning nor the operational improvements desired are likely if each vendor has a proprietary control plane, with interworking between vendors (and hence between networks, in most cases) left as a problem for operators' OS's to solve.

3.3 Service Framework Suggestions

For the reasons given above and others, we expect that the best model for an optical layer control plane within a trust domain is one that pays heavy attention to the management of heterogeneous technologies and associated service capabilities. This might be done by hiding complexities in subnetworks. These subnetworks would then advertise only a standardized abstraction of their connectivity, capacity, and functionality capabilities. Hopefully this would allow even disruptive technologies such as all-optical subnetworks to be introduced with a minimum of impact on preexisting parts of the trust domain.

Each network operator will have a need to define "branded" services - bundles of service functionality and SLA's with a specific price structure. In a heterogeneous network it will be necessary to map a customer request for such a "branded" service onto the specific capabilities of each subnetwork. This suggests a hierarchical model, decisions about these mappings, and also about policies for peering with other networks and overall management of the service offerings available to specific customers managed centrally but application of these policies handled at the local or subnetwork level.

4. Security Considerations

The solution developed to address the requirements defined in this document must address security aspects.

5. Acknowledgments

This document has benefited from discussions with Michael Eiselt and Mark Shtaif.

References:

[Ashwood00] Ashwood-Smith, P. et al., "MPLS Optical/Switching Signaling Functional Description", Work in Progress, draft-ashwood-generalized-mpls-signaling-00.txt.

[Awduche99] Awduche, D. O., Rekhter, Y., Drake, J., and Coltun, R., "Multi-Protocol Lambda Switching: Combining MPLS Traffic Engineering Control With Optical Crossconnects", Work in Progress, draft-awduche-mpls-te-optical-01.txt.

[Chaudhuri00] Chaudhuri, S., Hjalmtysson, G., and Yates, J., "Control of Lightpaths in an Optical Network", Work in Progress, draft-chaudhuri-ip-olxc-control-00.txt.

[Doverspike00] Doverspike, R. and Yates, J., "Challenges For MPLS in Optical Network Restoration", IEEE Communication Magazine, February, 2001.

[ITU] ITU-T Doc. G.663, Optical Fibers and Amplifiers, Section II.4.1.2.

[Kaminow97] Kaminow, I. P. and Koch, T. L., editors, Optical Fiber Telecommunications IIIA, Academic Press, 1997.

[Moy98] Moy, John T., OSPF: Anatomy of an Internet Routing Protocol, Addison-Wesley, 1998.

[Ramaswami98] Ramaswami, R. and Sivarajan, K. N., Optical Networks: A Practical Perspective, Morgan Kaufmann Publishers, 1998.

[Tkach98] Tkach, R., Goldstein, E., Nagel, J., and Strand, J., "Fundamental Limits of Optical Transparency", Optical Fiber Communication Conf., Feb. 1998, pp. 161-162.

[Yates99] Yates, J. M., Rumsewicz, M. P. and Lacey, J. P. R., "Wavelength Converters in Dynamically-Reconfigurable WDM Networks", IEEE Communications Surveys, 2Q1999 (online at www.comsoc.org/pubs/surveys/2q99issue/yates.html).

Authors' Addresses:

Angela Chiu
AT&T Labs
200 Laurel Ave., Rm A5-1F06
Middletown, NJ 07748
Phone:(732) 420-9057
Email: alchiu@att.com

John Strand
AT&T Labs
200 Laurel Ave., Rm A5-1D06
Middletown, NJ 07748
Phone:(732) 420-9036
Email: jls@att.com

Robert Tkach
Celion Networks
1 Shiela Dr., Suite 2
Tinton Falls, NJ 07724
Phone:(732) 747-9909
Email: bob.tkach@celion.com

James Luciani
PO Box 1010
Concord, MA 01742
Phone:(978) 266-2854
james_luciani@mindspring.com