EXTENDING GENERALIZED MULTI PROTOCOL LABEL SWITCHING TO CONFIGURABLE ALL-OPTICAL NETWORKS

Jane M. Simmons, Adel A. M. Saleh, Lotfi Benmohamed
Corvis Corporation
7065 Gateway Drive
Columbia, MD 21046
Phone: (301) 310-3410; Fax: (301) 310-3017
jsimmons@corvis.com

1. Introduction

Dense Wavelength Division Multiplexing (DWDM) has emerged as the favored transport technology in backbone networks to address the explosive growth in data traffic. However, while DWDM systems provide a large amount of bandwidth, that capacity may not always be optimally allocated within the network. Furthermore, with the large amount of traffic that is carried per fiber, coarse-granularity restoration is highly desirable. This has led to new network service requirements such as bandwidth-on-demand and restoration at the optical layer. Provisioning and restoring services rapidly, across multiple layers of a network, has become the focus of much effort in the industry.

An important factor for rapid provisioning and traffic management is that the process be scalable in terms of nodal processing requirements. One of the methodologies used to address scalability issues in the IP layer is Multiprotocol Label Switching (MPLS).[1] The goal of MPLS is to reduce the processing requirement of IP routers by assigning labels to traffic streams, and performing more efficient label-based switching as opposed to IP-header-based routing. Generalized MPLS (GMPLS), recently proposed in the Internet Engineering Task Force (IETF), extends the paradigm of label assignment and label-based switching to multiple network layers.[2] Provisioning in the optical layer is explicitly addressed in GMPLS. The Optical Internetworking Forum (OIF) is also focussed on developing a specification for rapid provisioning of the optical layer. The OIF efforts are embodied in the Optical User to Network Interface (O-UNI) specification.[3]

In the conventional IP-over-optical reference model, shown in Figure 1, the 'optical layer' consists of O-E-O crossconnects connected by WDM point-to-point links. In this model, every wavelength is converted to the electronic domain at each node; i.e., there is no optical bypass at a node. A reference model based on a configurable all-optical layer, which is the subject of this paper, is shown in Figure 2. In the configurable all-optical layer, traffic can be dynamically added/dropped at a node, or it can remain in the optical domain and optically bypass the node. In this model, an O-E-O crossconnect layer can still exist at the edge of the all-optical layer. However, as will be discussed below, the size of the O-E-O crossconnect in this scenario is smaller than that for the model of Figure 1.

While the efforts of the standards bodies and forums adequately address provisioning in the reference model of Figure 1, where O-E-O crossconnects are interconnected by point-to-point WDM links, they are not sufficient for the reference model of Figure 2, where a configurable all-optical layer is present. Configurable all-optical networks have specific properties (e.g., routing schemes intimately tied to the underlying technology) that are not adequately captured in current models. In this paper, we address the shortcomings in the current models, and propose extensions to them for provisioning in a configurable all-optical network. A precursor to this work was presented in the OIF.[4]

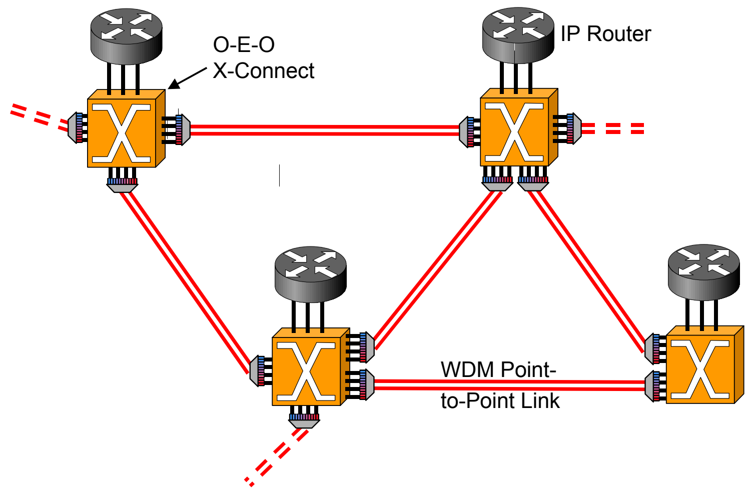

Figure 1 *The conventional IP-over-optical reference model, based on O-E-O crossconnects and WDM point-to-point links. The IP router layer is connected to the crossconnect layer.*

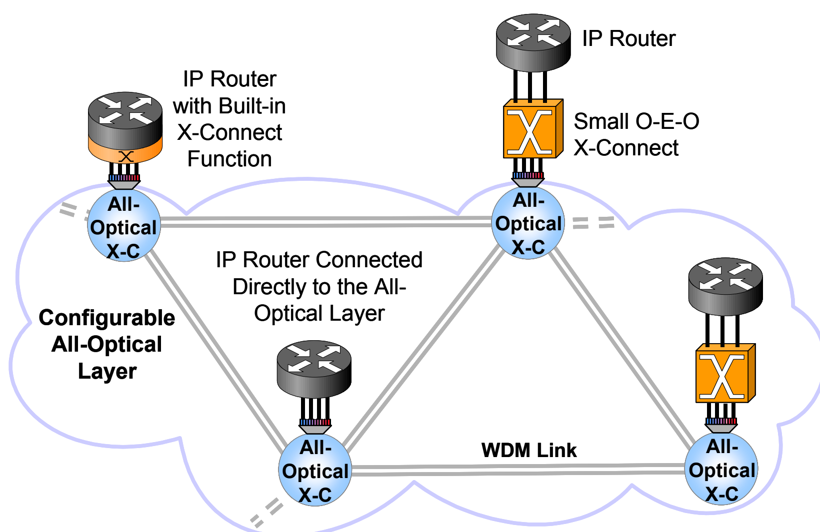

Figure 2 *The IP-over-optical reference model based on a configurable all-optical layer. A layer of small O-E-O crossconnects can exist over the all-optical layer. The IP router can be connected to an O-E-O crossconnect or directly to the all-optical layer. Some IP routers have a built-in crossconnect function.*

2. Benefits of Configurable All-Optical Networks

Configurable all-optical networks have just begun deployment for commercial applications. However, based on the many economic and scalability benefits of this technology and the growing acceptance in the industry, it is likely that all-optical networks will become a driving force in the industry. Thus, it is important that emerging standards address the configurable all-optical model.

In conventional point-to-point WDM systems, as exemplified by the O-E-O crossconnect layer, each wavelength is received electronically at every network node (see Figure 3). A WDM Transmit/Receive (Tx/Rx) line card is required for each one of the wavelengths. However, the majority of the wavelengths entering each node are not carrying traffic that is destined for that node; i.e., this traffic is 'express', as opposed to 'local', relative to that node.

With the advent of long reach optics, where signals can travel long distances without regeneration, it is not necessary to electronically terminate every wavelength entering a node. Rather, the express wavelengths can remain in the optical domain, and optically bypass the node, as shown in Figure 4. Tx/Rx line cards are not needed for the wavelengths optically bypassing a node, which dramatically reduces the cost of the overall network and significantly reduces the space and power requirements. [5, 6, 7] Furthermore, network provisioning is greatly simplified and accelerated because each additional circuit typically requires a Tx/Rx line card only at the circuit endpoints, as opposed to at each node or regeneration point along the circuit path.

Due to the dynamic nature of traffic, it is necessary that the optical layer be configurable. A static network, where all wavelength paths are fixed, leads to inefficient routing and stranded capacity. In a point-to-point WDM system, configurability can be accomplished with an O-E-O crossconnect that directs any input line to any output line. However, in order to have total reconfigurability, it is necessary that the O-E-O crossconnect be large enough to accommodate all wavelengths entering a node, as shown in Figure 3. As the number of wavelengths entering a node continues to soar from tens to hundreds to thousands, the size of the required O-E-O crossconnect will become prohibitive in terms of technology, cost, power, and space.

Reconfigurability in an all-optical network is accomplished through the use of all-optical switches and optical add/drop multiplexers. These network elements allow wavelengths to optically bypass a node or be dynamically dropped/added at a node. Furthermore, all-optical switches allow express wavelengths to be redirected onto any network port while remaining in the optical domain, as illustrated in Figure 4. Thus, reconfigurability is accomplished without the drawbacks of conversion to the electrical domain, leading to a much more scalable network. For those wavelengths that must be accessed locally, a small adjunct O-E-O crossconnect can be used for grooming and regeneration purposes, as shown in the central node of Figure 4. Since the bulk of the traffic does not pass through the O-E-O crossconnect, the required crossconnect size remains tractable even as the network grows.

In summary, the benefits of a configurable all-optical layer, in terms of cost, space, power, provisioning speed, and perhaps most importantly, scalability, will likely lead to widespread deployment of this technology as carriers grow their networks. It is very worthwhile to ensure that the interoperability paradigms and standards that are currently being developed adequately address this technology.

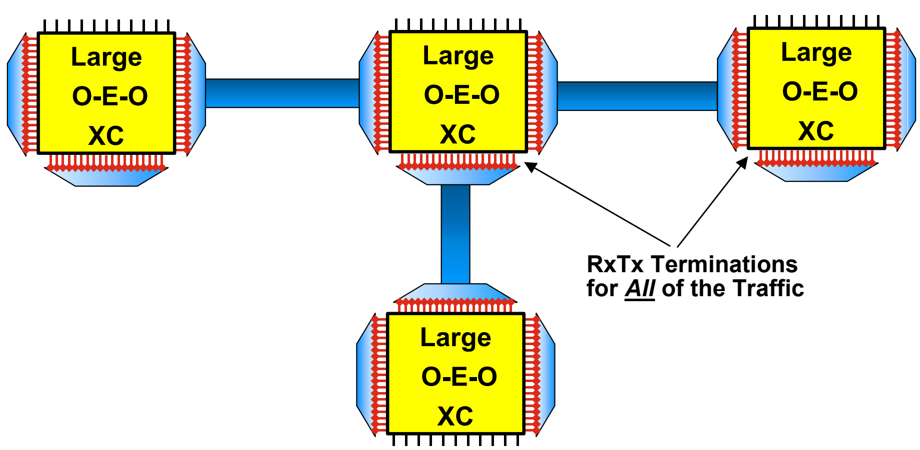

Figure 3 *In the conventional architecture, O-E-O crossconnects are interconnected by point-to-point WDM links. At each node, Tx/Rx line cards are needed for all of the traffic, regardless of whether it is express or local.*

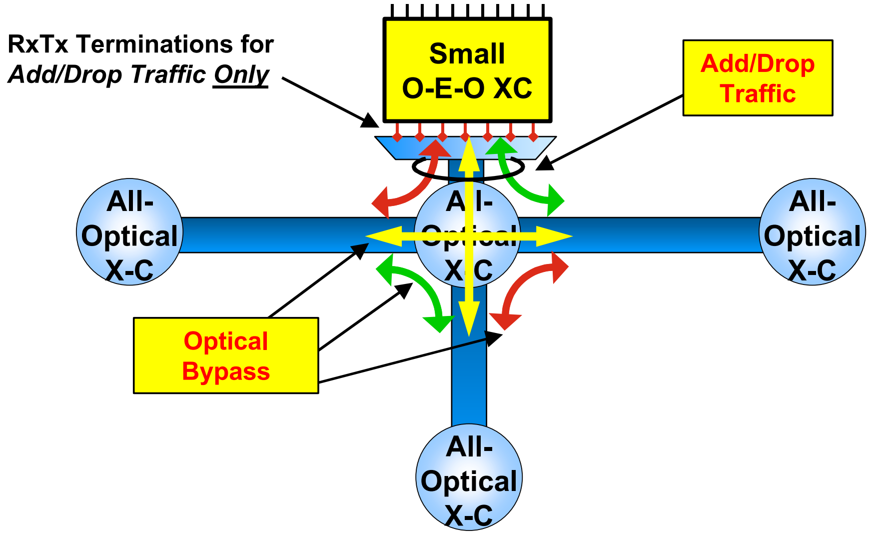

Figure 4 *In the all-optical architecture, ultra-long reach is combined with all-optical switches such that Tx/Rx line cards are required only for local add/drop traffic. The remaining wavelengths optically bypass the node.*

3. Peer-to-Peer vs. Client-Server Interoperability Models

Two reference models have emerged for interoperability with the optical layer: the peer-to-peer model and the client-server model (also known as the overlay model).[8] Either interoperability model is suitable for the architecture where O-E-O crossconnects are interconnected by point-to-point WDM links (although the OIF O-UNI addresses only the client-server model). However, the configurable all-optical layer has important properties such that, in order to reap the full benefits of this technology, we contend it is necessary for it to operate using the client-server model.

We will focus our discussion on the interoperability of an IP router layer with the optical layer, as shown in Figure 2. One option is for the IP router to interface with an O-E-O crossconnect, which in turn interfaces with a configurable all-optical layer. Another option is for the IP router, either with or without a built-in crossconnect functionality, to interface directly with the configurable all-optical layer.

In the peer-to-peer model, the IP routers are routing peers of the network elements in the optical layer (e.g., the O-E-O crossconnects or optical switches). This model relies on being able to capture the cost of routing within the optical layer and export this information to all IP routers, for example, through extensions to Open Shortest Path First (OSPF). This enables the IP routers to perform end-to-end routing, including selecting the optimal path to follow through the optical layer.

One of the important network costs that needs to be captured when establishing a connection in the optical layer is the amount of required regeneration. Regeneration incurs cost in terms of requiring a WDM Tx/Rx line card, and, often, in occupying ports on an O-E-O crossconnect. Minimizing regeneration should certainly be one of the important routing criteria. In the O-E-O layer, determining the amount of required regeneration along a path is relatively simple. The signal is regenerated at each node, and, depending on the distance of the link, possibly between nodes at fixed regeneration sites.

In the configurable all-optical layer, exporting the information required for optimal routing would be extremely difficult. There are a number of physical effects that affect where regeneration is required. For example, the optical reach in a given network is dependent on impairments such as accumulated noise, chromatic dispersion, polarization mode dispersion, nonlinear effects, and cross-talk. The effect of these properties on the optical signal depends on the transport and switching technology that is used, and thus can vary widely among vendors. Furthermore, some of these factors are dependent on the wavelength that is being routed; thus, routing is intricately tied to the choice of wavelength. This also implies the cost of a particular path may depend on which wavelengths are currently available along that path. Codifying a routing protocol that can take into account all of these effects is impractical.

In addition, many of the properties listed above depend on the type of fiber deployed in the network and the specific amplifier spacing. Since carrier networks typically include a wide range of fiber types and amplifier spacings, the IP layer would be forced to have knowledge of the deployed physical layer and would need to maintain a multitude of rules for each fiber type. This would create a large processing burden at the IP layer, and would mitigate the benefits of network layering.

A preliminary attempt to capture some of the routing impairments was proposed in [9], however, it did not fully capture the routing properties of a configurable all-optical network. An alternative proposal is to assume worst case regeneration. For example, assume there is a simple linear chain of nodes A, B, C, D and E. Assume that it is desired to set up a connection between A and E. Assume also that in the worst case scenario, the signal from A will need to be regenerated at C, while in the best case, it can travel all-optically to E. One simplifying assumption would be to always regenerate the signal from A at C, such that the effects of different wavelengths, etc. can be ignored. However, such simplifying rules will not allow the full benefits of the all-optical network to be realized.

Clearly, functions such as routing, selection of regeneration sites, and selection of wavelengths are necessarily intricately tied to the properties of the all-optical network. It is best that these functions be handled in the all-optical layer, implying that the client-server model should be used to interface to the all-optical layer. In such a model, the higher layer (e.g., the IP router, possibly via an OEO crossconnect) requests that the all-optical layer provide a connection from A to Z, and leaves it up to the all-optical layer as to how that connection should be provided. In the sections below, we will discuss the implementation of such a model in the context of GMPLS.

One extension to the client-server model would be to advertise to the O-E-O layer which O-E-O nodes can be connected by paths that remain entirely in the all-optical layer. However, due to wavelength continuity requirements, there may be a great deal of interdependence in this information. For example, the selection of a particular wavelength for one connection may eliminate all-optical reachability between some other nodes. This model would require a great deal of calculation and a large amount of information exchange between the all-optical and O-E-O layers; thus, similar to the peer-to-peer model, it may be too burdensome to implement.

4. GMPLS Hierarchy

In MPLS, labels are dynamically assigned at 'label switching routers' (LSRs) to represent a mapping from the ingress to the egress of the LSR. Labels have only local significance at a particular LSR, so that the total number of labels assigned at a LSR is relatively small. The label contained in an incoming packet is used to specify the next hop as well as the new label for that packet. (Essentially, labels serve the role of virtual circuit identifiers.) The traffic path can be represented by a series of labels on LSR interfaces, hence, it is referred to as a Label Switched Path (LSP). Since labels have only local significance, the label-lookup process can be very fast. This is especially beneficial in an IP network, where label-lookup can replace the more burdensome process of looking up the destination IP address in a global routing table.

One benefit of path setup through the labeling mechanism is that it provides control over the path of a traffic flow. For example, two flows may have the same source and destination points, but different requirements in terms of latency or reliability. At the ingress point, the packets belonging to the two flows could be assigned labels that will send them on different paths through the network. Labeling also provides a simple way of indicating a secondary path; changing the LSP mapping of a flow at the ingress can be used to direct the flow along a disjoint path.

Another powerful feature is that labels can be 'stacked'. For example, assume three paths are routed through three contiguous subnets. Assume the paths through the first and third subnets are different for all three flows, but that all three flows follow the same path through the second subnet. Each of the flows will be assigned labels such that they will follow different paths through the first subnet. However, at the ingress to the second subnet, another label can be 'pushed' onto the label stack, where this label is the same for all three flows (i.e., there is a nesting of labels). This 'outermost' label will be used to direct the path through the second subnet, until it is 'popped' at the egress of this subnet. All three flows are effectively aggregated into a single flow in the middle subnet, allowing for efficient traffic bundling and simplified traffic management.

The Multi Protocol Lambda Switching (MPλS) proposal extends the label switching paradigm of MPLS to optical networks.[10] This work was further extended in GMPLS to cover a multi-layered network. The GMPLS proposal covers a packet-switching layer (e.g., IP), a TDM layer (e.g., SONET), a wavelength-switching layer, a waveband-switching layer, and a fiber-switching layer. In general, in the optical layer, the optical crossconnect serves as the LSR, and the label is a wavelength, a waveband, or a fiber/port, depending on the granularity of the crossconnect. As opposed to MPLS where an explicit label is appended to the data, the optical layer labels are implicit.

GMPLS provides a unified paradigm for provisioning across multiple layers of a network. A flow at one layer of the network can 'tunnel' through a lower layer by adding a label appropriate for the lower layer. Multiple flows that have disparate labels at one layer can be assigned the same label at a lower layer. Consider the example of multiple SONET time slots multiplexed onto the same wavelength. The TDM-layer label will correspond to the individual time slot; however, the wavelength label assigned in the wavelength-switching layer will be the same for all of the multiplexed time slots.

GMPLS labels can be used to direct the flow of traffic across many layers. It is compatible with either the peer-to-peer or client-server model, depending on whether a lower layer advertises enough information to allow an upper layer to create the label assignments in the lower layer. The GMPLS hierarchy represents progressively coarser granularity. Thus, the hierarchy of GMPLS labels serves to aggregate traffic into coarser bundles, making GMPLS a good provisioning mechanism.

5. Extending GMPLS to a Configurable All-Optical Layer

The implementation of GMPLS as specified in reference [2] is appropriate for an optical layer with crossconnects interconnected by point-to-point WDM links. At each node, the crossconnect (the LSR) routes the traffic based on the incoming wavelength and port. However, as we will discuss, this is not the most natural labeling mechanism for an optical layer that is comprised of an O-E-O layer over a configurable all-optical layer. At the O-E-O layer, we propose that the label be a 'subconnection', as defined below (with the O-E-O crossconnect still serving as the LSR). We discuss how subconnections can be effectively used in the provisioning process, and why this is a good mechanism for interfacing to an all-optical layer.

5.1. Subconnections

A subconnection represents a flow that remains in the configurable all-optical layer except for the subconnection endpoints, where it is processed by the O-E-O layer. Figure 5 illustrates three distinct subconnections. Subconnection 1 extends over just one link; it enters the O-E-O crossconnect layer at both endpoints of the link. Subconnection 2 extends over three links. It enters the O-E-O layer at Nodes B and E, and optically bypasses C and D. Subconnection 3 includes regeneration at Node H. Even though it returns to the electrical domain at Node H, it does not enter an O-E-O crossconnect (the regeneration is handled in a static way, for example, in a patch-panel).

The O-E-O layer concatenates subconnections together to form an end-to-end connection. Thus, it is natural to let subconnections serve as the label in the O-E-O layer. The wavelength-labeling paradigm, while workable, is awkward when an all-optical layer is present. The output label of one O-E-O crossconnect may not be the same as the input label to the next O-E-O crossconnect, and thus may require a local mapping. For example, note that in Subconnection 3, the wavelength does not necessarily remain the same from the crossconnect at Node E to the crossconnect at Node I. A subconnection, however, has a consistent meaning at both of its endpoints. Thus, in the O-E-O layer, it is more natural to think of the end-to-end path as being comprised of a sequence of subconnections.

In the all-optical layer, the subconnections are comprised of a single 'all-optical fragment', or multiple 'all-optical fragments' if static regeneration is present (e.g., in Figure 5, the subconnection from E to I is comprised of the fragments E to H and H to I). (If all-optical wavelength conversion is present, the fragment could ride on multiple wavelengths.) Moreover, the fragments can be grouped and switched as part of a waveband-switching or fiber-switching layer. Thus, introducing the notion of subconnections does not violate the provisioning paradigm of GMPLS, where flows are grouped into coarser bundles to simplify routing and traffic management.

There is typically interplay between the O-E-O layer and the configurable all-optical layer in the provisioning process. On one hand, the all-optical layer relies on the O-E-O crossconnect to 'stitch together' the subconnections to form an end-to-end connection. On the other hand, the O-E-O layer must rely on the all-optical layer to select wavelengths and regeneration sites (for the reasons described in Section 3). Assume the O-E-O layer wishes to establish a connection between two nodes through the all-optical layer. It may not know a priori whether this connection can be established entirely in the all-optical layer. It may be necessary for the all-optical layer to establish two subconnections, with the O-E-O layer performing crossconnection at an intermediate point. Thus, when the O-E-O wishes to tunnel through the all-optical layer, it is not always certain how far the tunnel can extend.

The interplay between the O-E-O and all-optical layers is not unique to this combination of layers. For example, in the scenario where a wavelength-crossconnect interfaces to a waveband-crossconnect layer, the waveband layer may need to rely on the wavelength layer for regrooming of the wavebands.

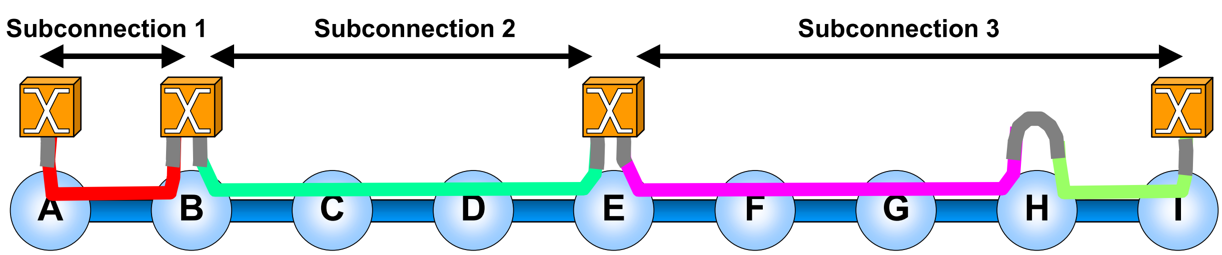

Figure 5 *Subconnections remain in the all-optical layer except for the endpoints, where they enter the O-E-O layer. Three subconnections are illustrated in the figure. Subconnection 1 extends over a single hop. Subconnection 2 optically bypasses Nodes C and D. Subconnection 3 optically bypasses Nodes F and G; the regeneration at Node H is static; i.e., it does not utilize an O-E-O crossconnect.*

5.2. Provisioning Over a Configurable All-Optical Layer

At any given time, there is likely to be a set of established subconnections that are not carrying traffic. For example, a carrier may choose to 'pre-position capacity' in the network by establishing subconnections based on expected traffic patterns. In this scenario, the Tx/Rx line cards are put in place and the appropriate wavelengths are reserved. In addition to pre-positioning subconnections, a carrier may choose to keep some of the constituent subconnections established when taking down an end-to-end connection; i.e., for one or more of the constituent subconnections of the deleted connection, the Tx/Rx line cards remain in place and the corresponding wavelengths remain reserved.

This process of establishing subconnections for future use is comparable to building up the pool of labels for the higher layer. The all-optical layer advertises its established subconnections to the higher layer. It is also important for it to advertise a cost associated with each subconnection. The cost can capture many different aspects, for example, the number of hops in the subconnection or the number of internal regenerations (if any). The cost may also depend on the state of the network. For example, the cost could capture the utilization of the links comprising the subconnection path so that lightly loaded paths can be preferentially selected. However, any state-based cost would need to be kept on a coarse basis, otherwise it would require too much overhead to keep updated. The all-optical layer must also provide Shared Risk Link Group information associated with the subconnections to enable diverse-path routing in the higher layers.[11]

Assume a request arrives at the O-E-O layer to establish a connection between two nodes. The O-E-O layer has knowledge of all subconnections that are available for use. These subconnections and their related costs can be included as adjacencies in a shortest path routing algorithm. One possible outcome of the routing algorithm is that a complete end-to-end path is found using established subconnections. However, this path may not necessarily be the best possible path, as exemplified by Figure 6. In this example, the desired end-to-end path is A to C, however, the only established subconnections are A to E and E to C, producing a somewhat circuitous path.

The path that is found has an associated cost that is based on the constituent subconnection costs as well as the cost of switching in the O-E-O layer (i.e., a critical resource, an O-E-O crossconnect port is utilized at each subconnection concatenation). The O-E-O layer must then decide whether to accept this path composed of established subconnections, or whether it should request that the all-optical layer create new subconnections. It is important that, up front, the O-E-O layer have knowledge of an 'acceptable' cost for each source-destination pair. If the path cost is greater than this threshold, then it requests that the all-optical layer provide a better path. There is, of course, no guarantee that a better path can be found. The decision of whether to request that a better path be found in the all-optical layer may also be influenced by the time requirements for establishing the connection. If the time requirements are very stringent, e.g., the new path is needed for restoration, then the O-E-O layer may be more willing to accept a sub-optimal path.

Another possible outcome of the routing process is that an end-to-end path between the desired endpoints does not exist using established subconnections. In this scenario, the O-E-O layer must request that the all-optical layer provide additional subconnections to complete the path.

The provisioning process we described above is a top-down approach, where the O-E-O layer is provided with routing options through the all-optical layer. An alternative scheme would be for some, or all, end-to-end connection requests to be passed down to the all-optical layer. The all-optical layer calculates the end-to-end path, which will typically be composed of a series of subconnections. It then requests that the O-E-O layer concatenate these subconnections to form the path. In either case, neither layer alone is capable of providing end-to-end provisioning.

It is important to note that requiring a costing mechanism such as the one that has been described is not unique to configurable all-optical networks. Such a mechanism is necessary for any multi-domain environment where some, or all, of the domains operate using a client-server model.[12] It is necessary for determining which domains the end-to-end path should traverse, and when to request that a domain establish a new internal path as opposed to utilizing the ones that are already established.

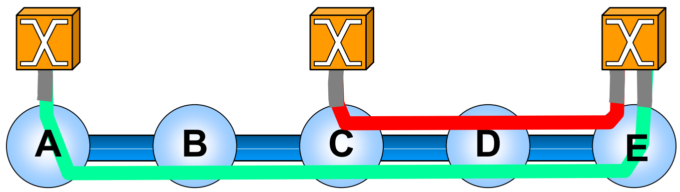

Figure 6 *The desired connection is from A to C, however, the only two existing subconnections are A to E and E to C. The O-E-O layer determines whether the cost of this path is within an acceptable threshold. If not, it can request that the all-optical layer establish more subconnections.*

5.3. Configurable All-Optical Layer Interface: UNI vs. NNI

A common point of debate is whether the interface between the O-E-O crossconnect and the configurable all-optical layer should be treated as a User to Network Interface (UNI) or a Network to Network Interface (NNI). Given that we assume the two layers are operating in a client-server mode, and that full information is not advertised to the O-E-O layer, it seems more appropriate to classify the interface as a UNI. Furthermore, the interface we described, where subconnections are advertised to a higher layer, is also suitable for the scenario where an IP router, either with or without built-in crossconnect functionality, is put directly on the configurable all-optical layer. This interface is more clearly a UNI since it involves the IP layer and the all-optical layer. Thus, referring to the all-optical interface in general as a UNI would provide consistency across all platforms.

6. Conclusions

The combination of ultra-long reach, optical bypass and dynamic configurability within the all-optical layer results in a more economical and scalable network. The many benefits of configurable all-optical networking is making this an important technology going forward. Due to the unique transmission rules in this layer and its dependency on vendor implementation in a rapidly changing technological environment, it is impractical to devise simple rules that can be exported to the O-E-O layer to enable end-to-end path selection in a peer-to-peer model. Thus, we argue that the all-optical layer should be operated in a client-server mode. We have extended the GMPLS paradigm to include provisioning over a configurable all-optical layer by using a judicious costing mechanism, combined with the notion of subconnections. This will allow the all-optical layer to be efficiently utilized and the full benefits of all-optical technology to be realized.

7. References

[1] E. C. Rosen, et al., "Multiprotocol Label Switching architecture," IETF, Internet Draft, July, 2000.

[2] P. Ashwood-Smith, et al., "Generalized MPLS - Signaling Functional Description," IETF draft-ietf-mpls-generalized-signaling-00.txt, work in progress, October, 2000.

[3] B. Rajagopalan (Editor) et al., "User Network Interface (UNI) 1.0 Signaling Specification," OIF2000.125.3, December, 2000.

[4] A. A. M. Saleh, et al., "Proposed Extensions to the UNI for Interfacing to a Configurable All-Optical Network," OIF2000.278, November, 2000.

[5] J. M. Simmons and A. A. M. Saleh, "The Value of Optical Bypass in Reducing Router Size in Gigabit Networks," *ICC'99*, Vancouver, British Columbia, June, 1999.

[6] A. A. M. Saleh, "Transparent Optical Networking in Backbone Networks," Invited Talk, *OFC '00*, Baltimore, MD, March 5-10, 2000, ThD7, pp. 62-64.

[7] R. E. Wagner, et al., "The potential of optical layer networks," *OFC'01*, Anaheim, CA, March 17-22, 2001, TuT3.

[8] OIF Carrier Group, "Carrier Optical Services Framework and Associated Requirements," OIF2000.155.1, September 2000.

[9] A. Chiu, et al., "Wavelength Routing in All-Optical Networks," OIF2000.251, November 2000.

[10] D. Awduche, et al., "Multi-protocol Lambda Switching: Combining MPLS Traffic Engineering Control With Optical Cross-connects," OIF00.014 (also IETF draft-awduche-mpls-te-optical-01.txt), January, 2000.

[11] S. Chaudhuri, et al., "Control of Lightpaths in an Optical Network," OIF2000.04, January, 2000.

[12] J. Strand and Y. Xue, "Routing For Optical Networks With Multiple Routing Domains," OIF2001.046.2, January, 2001.