

Nodal Architectures for Shared Mesh Restoration of IP and Wavelength Services

Jane M. Simmons, *Senior Member, IEEE*

Abstract— Future networks are expected to support heterogeneous traffic, including both IP and wavelength services. IP services are typically restored using IP-layer rerouting mechanisms, whereas wavelength services are restored in the optical layer. With architectures where IP services are mapped into Layer-2 or TDM protocols at the network edge, a common restoration mechanism can be utilized for both IP and wavelength services. This paper examines various nodal architectures that can be utilized to support such a unified restoration paradigm.

Index Terms—Edge cross-connect, IP, nodal architecture, optical networks, shared mesh restoration, wavelength services

I. INTRODUCTION

As networks evolve, various service layers are introduced into the network, typically with an accompanying restoration mechanism. Two services that are expected to be prevalent in future networks are Internet Protocol (IP) services and wavelength services. IP services are already well established in carrier networks, with restoration typically accomplished using IP-layer protocols that re-route packets around the failure. Wavelength services, where the end-to-end service request occupies a full wavelength, currently make up a small fraction of carrier traffic, but are expected to burgeon as applications such as e-science emerge. Wavelength services are likely to be restored in the optical layer, where the full wavelength is rerouted as a unit using core optical switches.

One architecture for supporting combined IP and wavelength services is shown in Figure 1(a). In this IP-over-Wavelength-Division-Multiplexing (WDM) architecture, traffic passes directly from the IP routers to the core switch. Any grooming of IP traffic occurs in the router; grooming refers to the bundling of traffic into wavelengths at various points along a data path. Such an architecture requires a large number of IP router ports, which has implications for network cost and power consumption. Supporting IP restoration strictly in the IP layer imposes further demands on the required size of the router.

To alleviate the burden on the IP router, an alternative architecture is to pass some or all of the IP traffic through a Layer-2 switch or a Time-Division-Multiplexing (TDM) switch (e.g., an Optical Transport Network (OTN) switch), as shown in

Fig. 1(b), where the grooming and the bulk of the restoration of the IP traffic is shifted to this layer [1-3]. For simplicity, the switch will be referred to here as a Layer-2 switch.

The chief rationale for this architecture is that Layer-2 switches are significantly lower in cost and power per Gb/s as compared to IP routers [1, 2, 4]. For example, a study using the national network and traffic set of Section 8.1 of [5] yields approximately: 1,700 IP ports for the Fig. 1(a) architecture; 750 IP ports and 1,650 Layer-2 ports for the Fig. 1(b) architecture. With IP ports on the order of three times more costly than Layer-2 ports [1] and requiring on the order of four times more power [4], the architecture of Fig. 1(b) provides a 25% savings in overall port cost and a 30% savings in port power.

With restoration occurring at a lower layer, the architecture is more amenable to using a shared restoration platform that is common with that used for wavelength services. In Section II, we summarize shared protection based on pre-lit subconnections, which is well suited for protecting these services. Sections III and IV present nodal architectures that allow the pre-lit subconnections to be used to protect either type of service. The discussion applies generally to sharing protection among sub-rate and wavelength services, as long as the Layer-2 switch bundles the work and protect sub-rate traffic into separate wavelengths.




Fig. 1. (a) IP-over-WDM. (b) IP traffic groomed and restored in Layer-2. Full-wavelength IP traffic can go directly to the core switch and be treated as a wavelength service. The core switch could, for example, be a ROADM.

II. PRE-LIT SUBCONNECTIONS

Consistent with the architecture of many carriers, we assume that the core switches enable optical bypass. With this technology, traffic that is transiting the node can remain in the optical domain, thereby eliminating much of the required electronic terminating equipment. For example, the core switch could be a Reconfigurable Optical Add/Drop Multiplexer (ROADM) or an all-optical switch. Because signals that remain in the optical domain over several links are susceptible to optical amplifier transients, it is desirable to maintain constant power in the fiber as much as possible. Thus, restoration schemes that rapidly turn on or off lasers, or rapidly switch wavelengths in the optical domain, can be undesirable.

One shared restoration scheme that avoids problems with optical amplifier transients is based on pre-lit protection subconnections. Pre-lit subconnections, where a wavelength is lit

Manuscript received May 12, 2009. J. M. Simmons is with Monarch Network Architects, Holmdel, NJ, USA.

Copyright (c) 2009 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

over one or more links prior to it actually being utilized, were first introduced in [6]. Details of the implementation of shared protection based on pre-lit subconnections can be found in [7]. Essentially, at the time of failure recovery, the appropriate protection subconnections are concatenated together to form an end-to-end protection path. The protection subconnections are pre-lit; thus, there is no need to rapidly turn on lasers at the time of failure.

An example of this scheme is shown in Fig. 2. The work paths are shown by the solid lines, and the protection subconnections are shown by the dotted lines. In this figure, there are seven subconnections to protect the four service demands. It is assumed that no more than one link/node failure occurs concurrently. If, for example, the work path between Nodes *A* and *C* fails, the subconnections along *A-D*, *D-F*, and *F-C* are concatenated together to form a protection path.

There are many nodal architectures that allow the protection subconnections to be shared by both the sub-rate and wavelength services. One distinguishing characteristic is whether the concatenation of subconnections is accomplished with an edge cross-connect or with the core switch. These architectures are considered in the following two sections.




Fig. 2. Seven pre-lit protection subconnections and four service demands. H-J and H-K are sub-rate demands that can be carried in a single wavelength. The remaining demands can be either sub-rate or a wavelength service.

III. SHARED RESTORATION WITH AN EDGE CROSS-CONNECT

Assume that an edge cross-connect is employed, as shown in Fig. 3, which depicts Node *F* from Fig. 2 (with some extra services). The edge cross-connect operates on the granularity of a full wavelength; it provides flexibility on the client side of the core switch. The IP router is not shown in the figure. The four protection subconnections are terminated on the edge cross-connect. There are also connections between the Layer-2 switch and the cross-connect to allow the sub-rate traffic to ‘grab’ the protection capacity. Under no failures, the work traffic of the wavelength services passes through the cross-connect. However, it is not necessary to pass the sub-rate work traffic through the cross-connect.

The restoration switching occurs in the edge cross-connect. First, assume that Link *A-B* fails, such that the *A-C* service requires protection; this could be either a sub-rate or wavelength service. At Node *F*, the *D-F* and *F-C* subconnections are concatenated in the edge cross-connect as shown by the dashed line labeled ‘1’ in Fig. 3.

In a second scenario, assume that Link *D-E* fails, bringing down the work path that extends between Nodes *D* and *F*. If this path carries a wavelength service, then restoration is ac-

complished at Node *F* by configuring the edge cross-connect to connect the wavelength-service client to the *F-J* subconnection as shown by the dashed line labeled ‘2’ in Fig. 3. If this path carries a sub-rate service, then the edge cross-connect is configured as shown by the dashed line labeled ‘3’, thereby connecting the Layer-2 switch to the *F-J* subconnection.

In a third scenario, assume that the paths between *H* and *J* and between *H* and *K* both carry sub-rate services, where the sum of the bandwidths is less than that of a single wavelength. Assume Link *H-I* fails. The two services are bundled together such that they use the same protection subconnections along *H-D* and *D-F*. The protection paths diverge at Node *F*, which requires that the Layer-2 switch be accessed. Thus, at the time of failure recovery, it is necessary for the protection subconnections to extend to the Layer-2 switch. This is accomplished by configuring the edge cross-connect as shown by the dashed lines in Fig. 4 (Fig. 4 is identical to Fig. 3; it is re-drawn to show this particular concatenation process).

In many carrier networks, core traffic grooming occurs in only a subset of the nodes. If Node *F* were not equipped with a switch capable of grooming, then the operation illustrated in Fig. 4 could not be supported. Assuming Node *D* also did not have a grooming switch, then additional protection subconnections would be required along *H-D* and *D-F*, such that the *H-J* and *H-K* sub-rate services could be protected separately.




Fig. 3. The restoration switching occurs in the edge cross-connect. The traffic is assumed to be bi-directional, such that, for example, subconnection ‘DF’ is equivalent to subconnection ‘FD’.




Fig. 4. The Layer-2 switch is also involved in the restoration process.

Note that in addition to providing protection against link/node failures, the edge cross-connect architecture also provides protection against failures of the transponders used for wavelength services; i.e., the service would be switched to a spare transponder. Transponders are transmitters/receivers plugged into the core switch that convert a short-reach signal to a long-reach signal. Furthermore, because the edge cross-connect is switching short-reach signals, it can be either opti-

cal, e.g., based on micro-electrical-mechanical-system (MEMS) technology, or electronic; an optical cross-connect is depicted in the figures. In either case, the switch speed should be fast enough to enable rapid restoration.

In the architecture of Figs. 3 and 4, restoration of the sub-rate traffic occurs at the sub-rate level. With sub-rate-level restoration, the work and protect paths of each sub-rate service are routed over diverse paths, and then the sub-rate paths are independently bundled into wavelengths. Another option is protecting sub-rate traffic at the optical layer, where the work paths are first bundled into wavelengths, and then the wavelengths are protected as a unit. Protecting sub-rate services at the optical layer is generally less capacity efficient (e.g., the operation illustrated in Fig. 4 is not supported with optical-layer protection), but may reduce the overall cost due to fewer required Layer-2 switch ports.

A nodal architecture that accommodates optical-layer protection of the sub-rate traffic is shown in Fig. 5 (this figure does not correspond to any of the nodes of Fig. 2). With this architecture, the restoration process for a wavelength carrying sub-rate services is identical to that for a wavelength service.




Fig. 5. Optical-layer protection of sub-rate services.

IV. SHARED RESTORATION WITH THE CORE SWITCH

Figure 6 depicts the architecture of Node *F* when the core switch is used for restoration (assuming restoration of the sub-rate traffic occurs at the sub-rate level). Both the sub-rate and wavelength services directly feed into the core switch (leaving the wavelength services vulnerable to transponder failures). The role of the Layer-2 switch in restoration is similar to that described in Section III. Some of the protection subconnections extend to the Layer-2 switch to allow grooming of the protect traffic. Note that the speed of the core switch must be fast enough to meet the restoration requirements, which may be challenging for core switches based on wavelength-selective switches (WSSs).

This architecture requires more advanced core switches than those of Section III. This is illustrated in Fig. 7, which addresses the concatenation of two protection subconnections. Figure 7(a) shows the configuration prior to failure, with the transponder details explicitly shown. The transponders are in the loop-back mode, with the receiver internally connected to the transmitter (if these subconnections were later needed to direct protect traffic to the Level-2 switch, the transponders could be reconfigured to the pass-through mode). Subconnection 1 is pre-lit on λ_1 and routed on Link 1, whereas Subconnection 2 is pre-lit on λ_2 and routed on Link 2. Figure 7(b) shows the concatenation of the two subconnections after failure recovery, where the receiver of Subconnection 2 now re-

ceives the pre-lit λ_1 signal and the receiver of Subconnection 1 receives the pre-lit λ_2 signal. The core switch must be both *colorless* and *directionless* to support this operation. Colorless refers to being able to receive (and transmit) any wavelength on any slot of the switch; directionless refers to any transponder being able to access any network port.

Another issue arises in core switches that have WDM add/drop ports, where a wavelength can only be used once on each port (this is typical of WSS-based core switches). If, in Fig. 7, the two subconnections are located on different add/drop ports, then, for example, λ_1 must be free on the drop port corresponding to Subconnection 2, otherwise wavelength contention would occur during failure recovery. This adds to the complexity of the restoration algorithms.




Fig. 6. Restoration using the core switch.




Fig. 7. (a) Protection subconnections prior to failure. (b) The core switch is reconfigured such that the subconnections are concatenated. The transmitters are unaffected, thus avoiding transients. (T=transmitter, R=receiver)

V. CONCLUSION

We have presented several nodal architectures that support unified shared mesh restoration for both wavelength and sub-rate services using pre-lit subconnections. Architectures based on an edge cross-connect tend to be simpler. Furthermore, the cost of a MEMS-based edge cross-connect has significantly decreased. For example, in the study noted in Section I, the edge cross-connects represent only on the order of 1% of the total network equipment cost.

REFERENCES

- [1] S. Sengupta, V. Kumar, and D. Saha, "Switched optical backbone for cost-effective scalable core IP networks," *IEEE Communications*, vol. 41, no. 6, pp. 60-70, Jun. 2003.
- [2] D. Stoll, "Constraints in multilayer optimization for IP traffic over optical WDM networks," *OFC/NFOEC'08*, San Diego, CA, Feb. 24, 2008.
- [3] S. Elby, "Bandwidth flexibility and high availability," *OFC/NFOEC'09*, San Diego, CA, Mar. 25, 2009.
- [4] R. Tucker, "A green Internet," *IEEE LEOS Annual Meeting*, Nov. 9 - Nov. 13, 2008, Newport Beach, CA.
- [5] J. M. Simmons, *Optical Network Design and Planning*, Springer, New York, 2008.
- [6] J. M. Simmons, A. A. M. Saleh, and L. Benmohamed, "Extending Generalized Multi-Protocol Label Switching to configurable all-optical networks," *NFOEC'01*, Baltimore, MD, July 8-12, 2001, pp. 14-23.
- [7] J. M. Simmons, "Cost vs. capacity tradeoff with shared mesh protection in optical-bypass-enabled backbone networks," *OFC/NFOEC'07*, Anaheim, CA, Mar. 25-29, 2007, Paper NThC2.