T_1 onto these lightpaths. The switching granularities of the nodes 1, 2, 3, and 4 are STS-1, OC-3, STS-1, and STS-1, respectively. The capacity of a wavelength channel is OC-12 for this illustration. Figure 1(a) shows the switching state of the OXCs. Since the switching granularity of node 2 (OC-3) is coarser than the bandwidth granularity of the traffic demand (OC-1), there is a free STS-1 timeslot (timeslot 3 in L_1) switched onto L_2 (timeslot 9 in L_2) by the OXC at node 2. Although this timeslot goes through node 2, it cannot be accessed by node 2. Any traffic carried by this timeslot will bypass node 2 and directly reach node 3, where it can be switched to any free outgoing grooming port. This is equivalent to having an STS-1 circuit directly connecting node 1 and node 3. Figure 1(b) shows the network state (virtual connectivity) after routing T_1. These circuits form another topology above the virtual topology, and traffic demands should be routed on this topology instead of on the virtual topology.

4. Network Design Framework

To accommodate characteristics of multi-granularity networks, we extend the graph model proposed in [2]. The graph model can route a traffic demand according to the current network state, and update the network state after carrying the traffic. The extended graph model can also intelligently choose the appropriate type of OXC to carry the current traffic demand, given there are several types of OXCs at a node. Due to space limitation, the extended graph model is not shown here.

Based on the extended graph model, we propose a design procedure as follows.

Step 1	Place one OXC of each type at each node.
Step 2	Compute a route for each traffic demand, and choose the most suitable OXC at each node along the route using the extended graph model, until all the traffic has been carried.
Step 3	For each type of OXC at each node, move the traffic going through the other types of OXC to this type of OXC, estimate the port cost of this type of OXCs, and choose the type of OXC with the least cost at each node.
Step 4	After determining the type of OXC at each node, reroute all the traffic demands and calculate the network cost.

5. Numerical Results

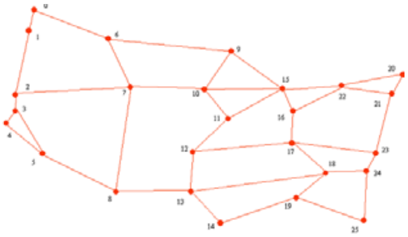

Figure 2. A 26-node nation-wide backbone network.

We conducted experiments on a typical nation-wide backbone network. The topology is shown in Fig. 2. It has 26 nodes and 40 bi-directional links. The capacity of a wavelength channel is OC-192. The bandwidth granularity of a traffic demand can be STS-1, OC-3, OC-12, OC-48, and OC-192, and the total traffic bandwidth requirement distribution of these 5 granularities is α_1: α_2: α_3: α_4: α_5, respectively. The traffic is uniformly distributed between all the nodes. There are 3 types of OXCs, shown in Table 1.

The per-port cost ratio of Type I, Type II, and Type III OXCs is β_1: β_2: β_3.

We compare the port cost in four scenarios. In Scenario 1, 2, and 3, there is only a Type I, II, and III OXC at each node, respectively; in Scenario 4, we use the above network design framework to determine the type of OXC at each node, and all three types of OXCs can coexist in the network. In the experiment reported here, the ratio of α_1: α_2: α_3: α_4: α_5, is 5:1:1:3:3, which is based on the projected traffic distribution of a typical nation-wide WDM backbone network, and the per-port cost ratio β_1: β_2: β_3 is 1:3:4. Note that these ratios are just inputs to our network design procedure, and more-accurate data, when available, can be plugged into our model.

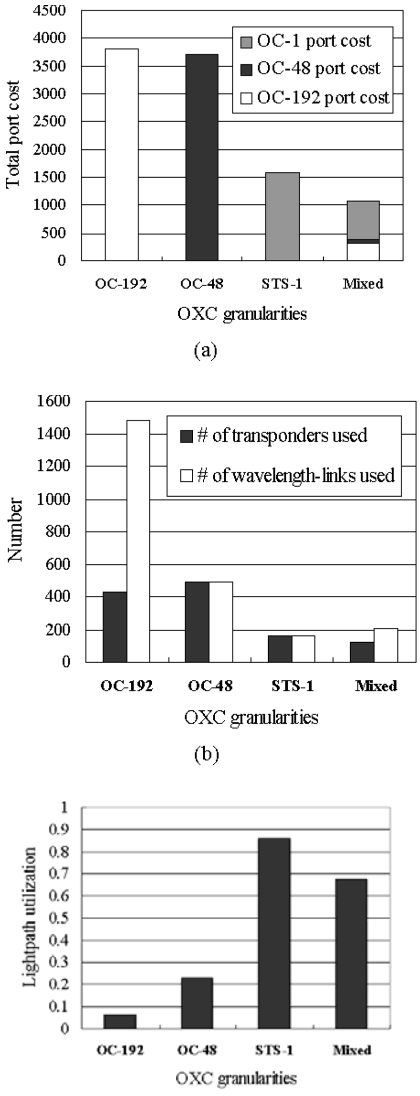

(a)

(b)

(c)

Figure 3 (a) Comparison of total port cost in the four scenarios. (b) Comparison of number of transponders and wavelength-links used in the four scenarios. (c) Comparison of the lightpath utilization in the four scenarios.

Figure 3(a) shows the total port cost, which is normalized by the per-port cost of all-optical OXCs, in the four scenarios, Fig. 3(b) shows the number of transponders used and wavelength-links used in the network, and Fig. 3(c) shows the lightpath utilization in the four scenarios. For the given traffic distribution and port cost ratio, the total port cost of the network in Scenario 1 is the highest, followed by the cost in Scenarios 2 and 3, and Scenario 4 achieves the lowest port cost. In Scenario 1, since the OXCs do not have grooming capability, the lightpath utilization is very low (6.5%) and 3822 OXC ports are used, resulting in highest total port cost despite the lowest per-port cost. In addition, this scenario uses the largest amount of wavelength-links to carry all the traffic. In Scenario 3, although the per-port cost of the type of OXCs used is the highest, the total port cost is less than that in Scenarios 1 and 2. This is because Type III OXCs can efficiently pack low-speed connections onto high-speed wavelength channels, making the lightpath utilization relatively high (86%). Hence, the total number of OXC ports (394), WDM transponders used (160), and wavelength-links used (160) are lower than those in Scenarios 1 and 2.

However, there is still room for improvement. For instance, not all of the nodes need such high flexibility in grooming fabric; some nodes may achieve similar performance with coarser grooming granularity or even no grooming capability, with the coordination of other nodes, thus further reducing the cost. This can be observed in Scenario 4. In this scenario, we choose an appropriate type of OXC for each node. Compared with Scenario 3, although more OXC ports may be used, the total port cost and the number of transponders used in the network are reduced about 33% and 23%, respectively, at the price of using more wavelength-links and lower lightpath utilization. This is because some Type III OXCs at some nodes are replaced with Type I and Type II OXCs, which have lower per-port cost than Type III, and Type I OXCs do not need transponders for bypassing traffic.

6. Conclusion

We proposed a framework for WDM backbone network design to better utilize the benefit of different type of OXCs, which have different bandwidth granularities. Our results demonstrate that using different type of OXCs will yield better network performance, and a design using our framework can reduce the network-wide OXC port cost.

References

[1] R. Ranganathan, L. Blair, and J. Berthold, “Architectural Implication of Core Grooming in a 46-node USA optical Network,” *OFC 2002*, ThP2, pp. 498-499.

[2] H. Zhu, H. Zang, K. Zhu, and B. Mukherjee, “A Novel, Generic Graph Model for Traffic Grooming in Heterogeneous WDM Mesh Networks,” *IEEE/ACM Trans. Networking*, to appear, 2003.

FR3 **11:00 AM**

Pre-Emptive Reprovisioning in Mesh Optical Networks

R. Ramamurthy, A. Akyamac, J-F. Labourdette,
S. Chaudhuri, *Tellium Inc., Oceanport, NJ, Email: jlabourdette@tellium.com.*

Pre-emptive reprovisioning is a method to perform reprovisioning of a backup path in advance of a second failure, to reduce the time to recover service from seconds (reprovisioning) to milliseconds (restoration). We evaluate the tradeoff between benefits and operational complexity.

1. Introduction

In shared-mesh restoration [1,2,5], each working path has a diverse backup path. In one restoration architecture [1,2], backup routes are pre-computed, and shared protection channels on the backup path are pre-assigned at the time of path provisioning. Channels on the backup path may be shared between primary paths whose working paths are diverse. Upon a single failure event, the lightpaths whose primary paths are affected by the failure are restored on their backup paths. If restoration fails (because the shared-protection channels are either in a failed state or are already being used by another lightpath - which can happen in the case of a double failure), then re-provisioning of the backup path is attempted. If reprovisioning of the backup path is successful, then the newly reprovisioned backup path is used to carry traffic. If reprovisioning fails (due to lack of capacity), then the lightpath is not restorable. Reprovisioning is a time-consuming process since it is performed at a centralized management system, and lightpaths are sequentially reprovisioned to avoid contentions for capacity. Pre-emptive reprovisioning is a method to perform reprovisioning of a backup path in advance of a failure. The motivation for pre-emptive reprovisioning is to reduce the time it takes to restore service from several seconds or tens of seconds with reprovisioning to order of 10s to 100s of milliseconds with restoration. However, there is operational complexity in pre-emptive re-provisioning. In this paper, we evaluate the trade-off between the benefits of pre-emptive reprovisioning and the additional operational complexity.

2. Pre-emptive Reprovisioning

Pre-emptive Reprovisioning automatically “re-provisions” a new backup path for a demand whose backup path has become unavailable. As a result, upon a failure, pre-emptively reprovisioned lightpaths can be restored in the order of milliseconds as opposed to being reprovisioned in the order of seconds to minutes. Backup paths can become unavailable due to several events that include:

Backup Channel Failure: A backup shared-channel fails (for example due to failure of electronics at one of the end-switches). In Fig. 1(b), a shared channel on the backup paths of P1 and P2 fails rendering the backups B1 and B2 unavailable.

Primary Channel Failure resulting in Lightpath Switch: A shared-mesh lightpath switches to its backup (rendering the backup channels unavailable for other lightpaths that are sharing them). In fig. 1(a), P1 switches to its backup B1, rendering the backup B2 of P2 unavailable.

Fiber failure: A fiber fails resulting in several lightpaths whose primaries use the failed fiber to switch to their backups.

The value of pre-emptive reprovisioning is to allow to restore rather than reprovision upon a double failure. We evaluate the performance of pre-emptive reprovisioning on 3 representative real networks

-Network 1 (45 nodes, 75 links, 72 shared-mesh demands, 97 shared channels)
-Network 2 (17 nodes, 26 links, 102 shared-mesh demands, 203 shared channels)
-Network 3 (50 nodes, 88 links, 300 shared-mesh demands, 476 shared channels)

The goal of the study is to evaluate the benefits of performing pre-emptive reprovisioning. The three failure events listed earlier are considered. Table 1 illustrates the “average” number of lightpaths whose backup paths become unavailable for each of the three events. The averages represent the average over all possible events.

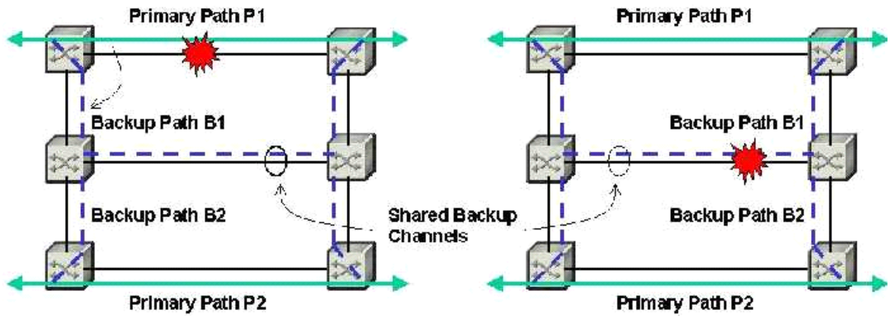

Figure 1. (a) illustrates shared-channels becoming unavailable due to a lightpath switching to protection resulting from a channel failure on the primary path, (b) illustrates a backup path becoming unavailable due to channel failure.

	Network 1	Network 2	Network 3
Backup channel failure	4	2.2	4.2
Primary channel failure	13.6	7.6	25.1
Fiber failure	17	55	192

Table 1. Number of lightpaths that would benefit from pre-emptive reprovisioning for different types of failures (average over all possible failures)

For backup channel failure, the number of lightpaths that need to be pre-emptively reprovisioned is equal to the number of lightpaths that share a protection channel [3,4,6], which is about 4 lightpaths for the above networks. For a lightpath switch to back-up, the number of lightpaths that need to be pre-emptively reprovisioned is equal to the number of backup hops multiplied by the number of lightpaths that share a protection channel, which is about 20 depending on the size of the network, but is independent of the network loading. For fiber failure, the number of lightpaths that need to be pre-emptively reprovisioned is equal to the average number of lightpaths on a fiber multiplied by the average backup hops multiplied by the average lightpaths that share a shared back-up channel, which is dependent on the network loading. Table 2 illustrates the “average” percentage of lightpaths that are successfully “pre-emptively reprovisioned” assuming that no extra capacity has been added for pre-emptive reprovisioning. Also included in the table is the percentage of lightpaths for which there is a “diversity violation” after pre-emptive reprovisioning. Diversity violation occurs when fully diverse primary and backup routes are not available and primary and backup routes have to share links. For backup channel failure, ~90% of lightpaths can be successfully pre-emptively reprovisioned (since there are only a few lightpaths that need pre-emptive reprovisioning), for lightpath switch, ~75% can be successfully pre-emptively reprovisioned, and for a fiber failure about ~50% can be successfully pre-emptively reprovisioned.

As shown in Tables 1 and 2, pre-emptive reprovisioning is not effective against fiber failures because a) a large number of lightpaths require pre-emptive reprovisioning (large operational impact) and b) only a limited percentage of lightpaths can indeed be successfully reprovisioned. In the remainder of this paper, we only consider pre-emptive reprovisioning in the case of an initial single channel failure. We now estimate the frequency of pre-emptive reprovisioning requests due to channel failures resulting from switch port failures. Given the FIT rates of switch ports, we can estimate the rate at which a given port fails within a certain duration (e.g., 1 week). We then estimate the number of channel failures in a given duration by multiplying the number of channels by the probability of a channel. We assume typical FIT rates for 2.5G, 10G ports and other equipment. We assume that 4 lightpaths need to be reprovisioned for a protection channel failure and 20 lightpaths need to be reprovisioned for primary channel failures. We find that the frequency of pre-emptive reprovisioning is dominated by primary channel failures, and the number of lightpaths pre-emptively reprovisioned depends on the size of the network. For example, we find that for Network 3, an average of 3 lightpaths are pre-emptively reprovisioned per week due to shared channel failures, whereas an average of 50 lightpaths are pre-emptively reprovisioned per week due to working channel failures.

The likelihood that pre-emptive reprovisioning is beneficial is the likelihood of a second failure impacting a lightpath during the Mean Time to Repair (MTTR) of the first failure (that caused the lightpath to become unprotected). We consider the first failure to be a primary channel failure causing a lightpath to switch to its backup, thereby resulting in 20 lightpaths to become unprotected, and pre-emptively reprovisioned. We then consider the likelihood of a second failure (being either another channel failure or a fiber failure) impacting one of 20 pre-emptively reprovisioned lightpaths. The following figure plots the likelihood that pre-emptive reprovisioning is beneficial against the MTTR of the first failure.

	Network 1		Network 2		Network 3	
	% successful	**% with diversity violation**	**% successful**	**% with diversity violation**	**% successful**	**% with diversity violation**
Backup channel failure	96	28	90	23	93	11
Primary channel failure	76	19	85	38	90	11
Fiber ailure	71	20	58	28	58	15

Table 2. Percentage of lightpaths that are successfully pre-emptively reprovisioned within existing capacity

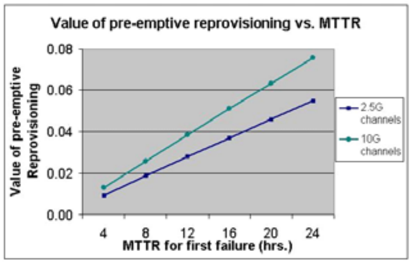

Figure 2. Probability that pre-emptive reprovisioning is successful

We find that when the MTTR is about 4 hours, the probability that pre-emptive reprovisioning will be beneficial is about 1 in 100 times. As the MTTR increases, the probability that pre-emptive reprovisioning is beneficial increases proportionally.

3. Conclusion

Pre-emptive reprovisioning of backup paths improves the restoration time upon multiple failures. We evaluated pre-emptive reprovisioning on several representative real networks. We find that the success rate of pre-emptive reprovisioning is about 90% for protection channel failures, 75% for working channel failures, and about 60% for fiber failures. We find that the likelihood that pre-emptive reprovisioning is beneficial, i.e., the chance of a second-failure impacting the network within the MTTR of the first-failure, is proportional to the MTTR and is about 1 in 100 when the MTTR is 4 hours. The decision to support pre-emptive re-provisioning is thus a trade-off between the operational complexity of re-provisioning after single failures and the chances that re-provisioning would be beneficial in those (rare) cases of double failures before the first failure is repaired. The benefit is that the lightpath would then be restored (10s to 100s of msec) rather than re-provisioned (sec.) after a double failure. This work provides some preliminary results on the benefits of preemptive re-provisioning.

References

[1] J-F. Labourdette, et. al., *"Routing Strategies for Capacity-Efficient and Fast-Restorable Mesh Optical Networks*," Photonic Network Communications, 4:3/4, pp. 219-235, July-December 2002.

[2] G. Ellinas, et. al., "*Routing and Restoration Architectures in Mesh Optical Networks*," Optical Networks Magazine, Issue 4:1, January/February 2003.

[3] E. Bouillet, J-F. Labourdette, R. Ramamurthy, S. Chaudhuri, "Enhanced Algorithm Cost Model to Control tradeoffs in Provisioning Shared Mesh Restored Lightpaths," in Proc. OFC 2002, Anaheim, CA, March 2002.

[4] J. Doucette, W. Grover, T. Bach, "Bi-Criteria Studies of Mesh Network Restoration: Path Length vs. Capacity Tradeoffs," in Proc. of OFC 2001, Anaheim, CA, March 2001.

[5] B. Doshi, et al, "Optical Network Design and Restoration," Bell-Labs Technical Journal, Vol. 4, No. 1, pp. 58-84, January-March 1999.

[6] C. Qiao, Y. Xiong and D. Xu, "Novel Models for Shared Path Protection," in Proc. OFC 2002, Anaheim, CA, March 2002.

FR4 **11:15 AM**

Implementation of Statistical Lambda Multiplexing Network (SLAMNet)

H. Nakamura, H. Yokoyama, S. Nomoto, *KDDI R&D Laboratories Inc., Saitama, Japan, Email: nakamura@kddilabs.jp*.

This paper discusses implementation of a novel signaling-free DWDM network architecture that provides multi-wavelength service. In the service, the number of assigned wavelengths is dynamically modified according to the volume of traffic on a best-effort basis.

I. Introduction

In a future large-scale Dense Wavelength Division Multiplexing (DWDM) network, multiple wavelengths will be used by a pair of routers to transmit high-volume traffic. A basic concept of a novel architecture called Statistical Lambda Multiplexing Network (SLAMNet) was proposed for providing multi-wavelength service in the DWDM network [1]. In comparison with the other architectures proposed for the DWDM network [2-5], the most striking feature of the architecture is signaling-free wavelength path switching to dynamically adjust the number of wavelengths to the varying volume of traffic. Taking the advantage of the signaling-free wavelength path switching, the SLAMNet can afford to be implemented to a part of a network with keeping its scalability, to be operated and maintained in a simple and economical way, and also to shorten the switching cycle regardless of a round trip time necessary for signaling. Furthermore, the signaling-free wavelength path switching can harmonize with signaling-based wavelength path management mechanisms such as GMPLS. The architecture was realized on a testbed. This paper discusses implementation of the SLAMNet and clarifies the target and effectiveness of the architecture.

II. Basic Concept of SLAMNet

A basic concept of the SLAMNet is illustrated in Figure 1. A pair of routers uses two types of optical channel (OCh) paths, referred to as initial paths and additional paths. The initial paths are fixedly established between the router pair and are used for transmission of traffic at all time. On the other hand, the additional paths are temporally utilized for transmission of burst traffic that overflows the capacity of the initial paths. The additional paths are dynamically setup and released by a pair of OCh path switches (OPSWs) and time-share wavelength resources assigned via optical

A basic configuration of an OPSW is illustrated in Figure 2. A monitor measures the volume of traffic transmitted by established initial and additional paths and observes its variation on a realtime basis. Additional paths are set up and released by making use of the variation of traffic volume as a trigger in source and destination OPSWs independently. Although the OPSWs

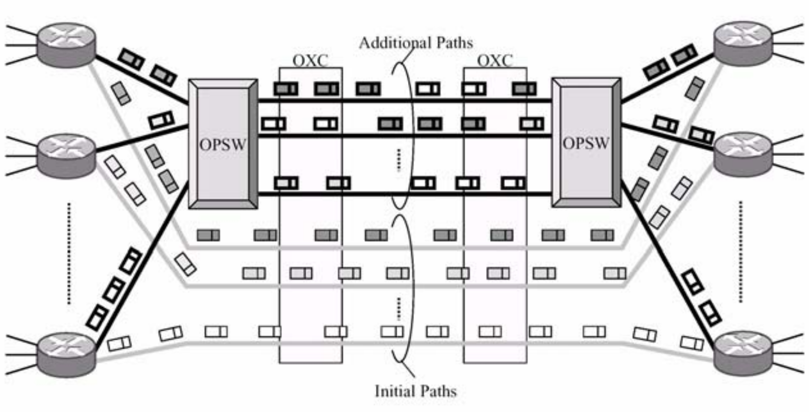

Figure 1 : Basic Concept of SLAMNet

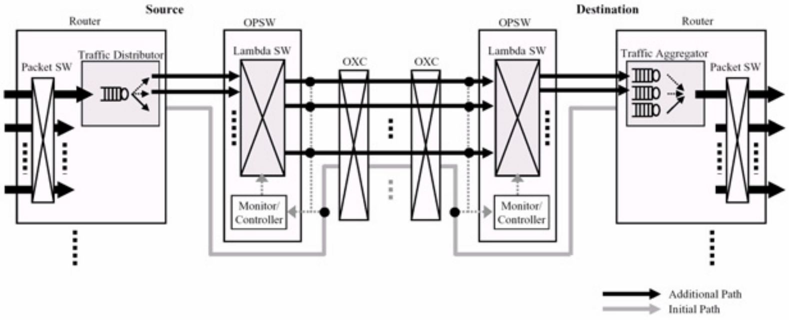

Figure 2 : Configuration of OPSW

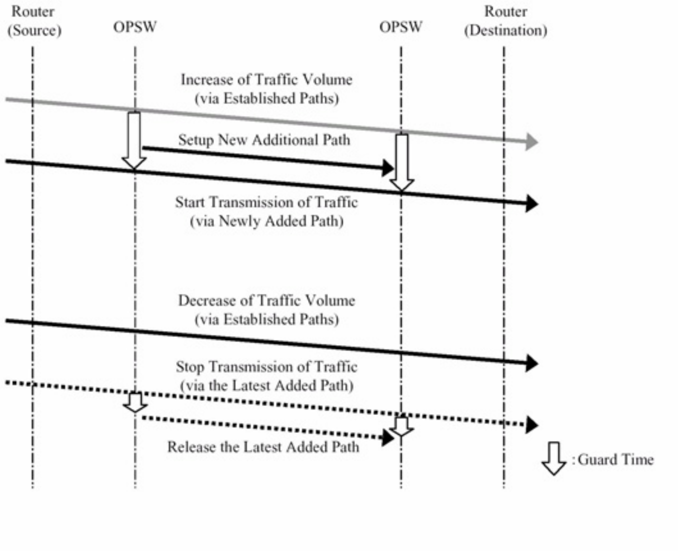

Figure 3 : Basic Flow of Signaling-Free Path Management