rated a 500 m long highly nonlinear fiber (HNLF) with a nonlinearity coefficient $\gamma = 10.9\ W^{-1}km^{-1}$ and a zero dispersion wavelength $\lambda_0 = 1552$ nm. The output from a continuous wave (CW) external cavity feedback laser at an operating wavelength of 1558 nm was used as a probe signal for the wavelength converter. By appropriately setting the polarization of the light within the NOLM and filtering out the 1549 nm control pulses at the output port, we were able to generate a high-quality 10 GHz train of 3.3 ps pulses at 1558 nm. These wavelength converted pulses were modulated to provide a 2^{31}-1 pseudorandom data sequence at 10 Gbit/s using a high-speed $LiNbO_3$ modulator. The 10 Gbit/s pseudorandom data stream was multiplexed upto an aggregate bit rate of 80 Gbit/s using a two-stage passive multiplexer (MUX) and subsequently sent into the data port of the 80 Gbit/s demultiplexer.

The high-speed data demultiplexer we used in this experiment was also based on a NOLM using a 1km length of nominally the same HNLF. The 1558 nm 80 Gbit/s data stream incident to the demultiplexer was passed through a tunable optical delay line to allow us to adjust the arrival time of the data pulses relative to the 1549 nm rectangular control pulses.

We performed switching window measurements by launching a CW probe signal into the data port of the demultiplexer and measuring the temporal width of autocorrelation traces. A good switching window characteristic with a ~10 ps FWHM triangular trace was obtained using the rectangular control pulses in contrast to a value of ~4 ps when driving the demultiplexer directly with 1.7 ps pulses from the EFRL. These results show that we can expect at most a 10 ps timing jitter tolerance by using the rectangular control pulses.

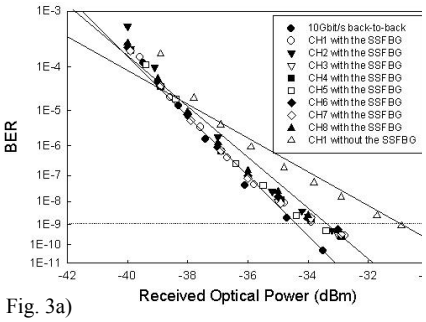

Fig. 3a)

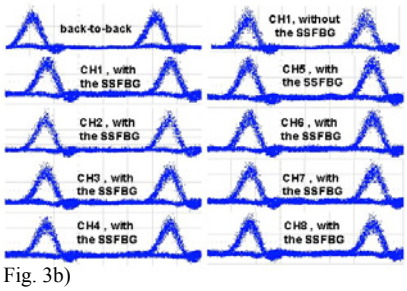

Fig. 3b)

Fig. 3. (a) BER and (b) Eye diagrams at –32 dBm, on the 80 Gbit/s to 10 Gbit/s demultiplexed signals as a function of received optical power under optimal time synchronization between the control and data pulses.

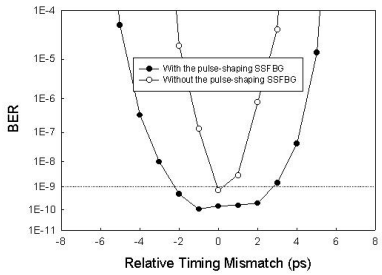

Fig. 4)

Fig. 4. The BER measured at a fixed received optical power of -31 dBm as a function of relative timing mismatch between control and data pulses.

In order to confirm the benefits of using rectangular control pulses in practical high-speed OTDM demultiplexing systems, we performed bit-error-rate (BER) measurements on the 80 Gbit/s demultiplexer. At first we measured the BER under optimal time synchronization between control and data pulses with the pulse shaping SSFBG (10 ps rectangular control pulses) and without it (soliton control pulses). The results are summarized in Fig. 3(a). Error-free demultiplexing operation for all 8-channels was achieved with 1 dB power penalty relative to 10 Gbit/s base rate back-to-back using a 10 ps wide rectangular shape control pulses. A 2 dB power-penalty improvement was also achieved with respect to data demultiplexing without the SSFBG. We attribute this penalty improvement to the enhanced tolerance to the ~sub-picosecond timing drift inherent to our experimental setup caused by environmental temperature/air flow variations during this particular experiment. The corresponding eye diagrams are shown in Fig. 3(b). Both intensity noise and timing jitter are clearly evident in the eyes of channel no.1 operated using soliton control pulses compared to those of all 8-channels operated using rectangular control pulses. Next, we measured BER at a fixed received optical power of -31 dBm as a function of relative timing mismatch between control and data pulses and the results are shown in Fig. 4. Error-free performance was readily achieved over a +/-3 ps timing mismatch range for the rectangular pulse driven demultiplexer versus a +/-0.5 ps range for the demultiplexer driven directly with the soliton pulses.

3. Conclusion

We have experimentally demonstrated that SSFBG based rectangular pulse switching technology can be used to provide improved timing jitter tolerance in a high-speed 80 Gbit/s OTDM demultiplexing system. Error-free demultiplexing operation with ~5 ps timing jitter tolerance was achieved using ~10 ps rectangular control pulses generated using a pulse shaping SSFBG. The 2 dB power-penalty improvement compared to demultiplexing without the SSFBG highlights that we can obtain significant OTDM system \performance enhancement simply adding by a SSBFG to the system. This approach could be applied to 160 Gbit/s OTDM systems since state-of-the art SSFBG technology can produce high-quality gratings capable of generating rectangular pulses with rise/fall time less than 1 ps.

*The authors would like to thank OFS for providing the HNLF.

4. References

[1] M. Nakazawa, "Tb/s OTDM Technology", in *Proc. 27th European Conference on Optical Communications (ECOC'2001)*, Paper Tu.L.2.3, pp.184-187, (2001).

[2] M. C. Gross, M. Hanna, K. M. Patel, and S. E. Ralph, "Reduction of power fluctuations in ultrafast optically time-division-multiplexed pulse trains by use of a nonlinear amplifying loop mirror", *IEEE Photon. Technol. Lett.*, **14**, pp.690-692 (2002).

[3] K. Uchiyama, and T. Morioka, "All-optical signal processing for 160Gbit/s channel OTDM/WDM systems", in *Proc. Optical Fiber Communications Conference (OFC'2001)*, Paper ThH2 (2001).

[4] M. Jinno, "Effects of crosstalk and timing jitter on all-optical time-division demultiplexing using a nonlinear fiber Sagnac interferometer switch", *IEEE J. Quantum Electron.*, **30**, pp.194-202 (1994).

[5] J. H. Lee, P. C. Teh, P. Petropoulos, M. Ibsen, and D. J. Richardson, "Timing jitter tolerant all-optical modulator and demultiplexing systems incorporating pulse-shaping fiber Bragg gratings", in *Proc. Optical Fiber Communications Conference(OFC'2001)*, Postdeadline Paper PD30 (2001).

[6] P. Petropoulos, M. Ibsen, A. D. Ellis, and D. J. Richardson, "Rectangular pulse generation based on pulse reshaping using a superstructured fiber Bragg grating", *J. Lightwave Technol.*, **19**, pp.746-752 (2001).

TuI 2:00 PM - 4:00 PM

Murphy1

Network Protection and Restoration 1

Georgios Ellinas, *Tellium, Inc., USA, Presider*

TuI1 **(Invited)** **2:00 PM**

p-Cycles, Ring-Mesh Hybrids, and "Ring Mining:" Options for New and Evolving Optical Transport Networks

W. Grover, *TRLabs / University of Alberta, Edmonton, AB, Canada, Email: grover@trlabs.ca*.

This paper gives an overview of three new concepts in which ring and mesh-like aspects are both present in the network architecture or strategy. *p*-Cycles yield mesh-like efficiency while retaining ring-like speed. Forcer clipping ring-mesh hybrids enhance overall capacity efficiency. Ring-mining gets valuable re-use of ring capacity while evolving to a mesh. These ideas provide operators with new options for planning or evolving optical transport networks.

1. Introduction

For many years networks operators have faced a black-or-white choice between ring or mesh-based architectures. Rings offer speed and the simplicity of a contained system with supporting standards and many vendors. In contrast mesh networks are a more general, flexible, and capacity-efficient, but are not standardized nor as supported by vendors. Interest in mesh is increasing, however. One reasin is the planning complexity, inefficiency and inflexibility experienced with multi-ring networks. Dynamic lightpath demand and multiple levels of protection are also more easily implemented over a mesh. But even these re-considerations of mesh versus ring are still in an "either-or" framework. Are there ways to get the best of both ring and mesh together? This paper explains some approaches which do involve both technologies in a single transport network architecture or strategy. In different senses these schemes capture the best of both ring and mesh principles.

2. *p*-Cycles

p-Cycles are like BLSR rings with one technical difference: the ring protection capacity can be accessed by *straddling* span failures, as well as the usual failures that occur and are protected on the ring itself. In terms of real-time switching for either type of failure, the action is completely BLSR-like loop-back reaction except that signals heading into a failed straddling span may switched to protection as well as the usual ring line signal. In Fig. 1(a), the ring-like reaction to failure of a span on the cycle of the ring is shown. But if the same protection capacity is operated as a *p*-cycle, it can also protect against failure of a "straddling" span such as in Fig. 1(b)

.

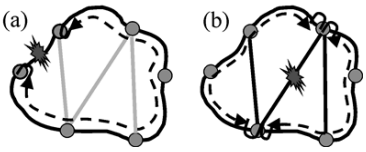

Figure 1(a): BLSR ring, (b) *p*-Cycles

This is not a major change to the equipment for a BLSR or optical shared protection rings. In fact

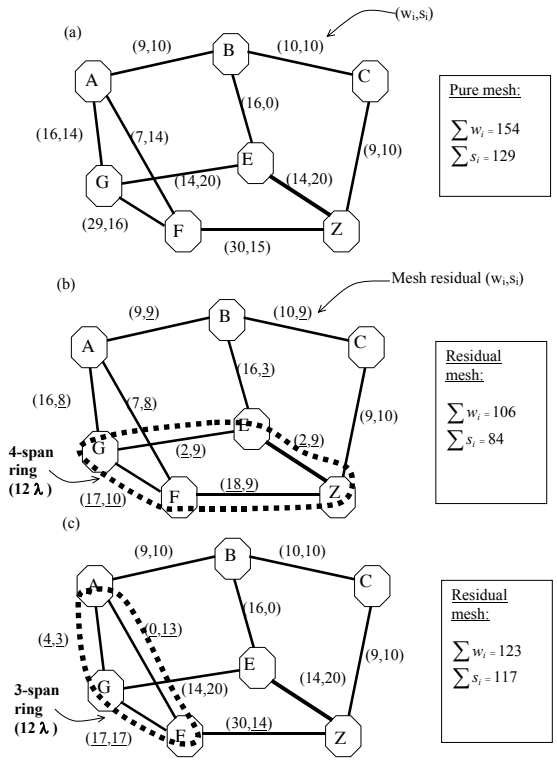

Figure 2(a) Optimized pure-mesh design, (b) ring with CRR = 0.97, (c) CRR=0.60

existing ADMs can be upgraded to operate as *p*-cycle nodes through addition of a straddling-span interface unit (described in the talk). But the difference it makes in efficiency is dramatic. Once straddling spans are introduced, there is no need for working paths to follow ring-constrained routes. Working paths can go shortest-path on any spans. This can save as much as 30% in working capacity, not to mention the greater flexibility and simplicity in routing. Secondly, protection redundancy becomes essentially that of a span restorable mesh; typically well under 100%, instead of well over 100% with rings. The reason is that straddling span relationships require no new spare capacity and use the existing protection with twice the leverage : when a straddling span fails, the cycle itself remains intact and can offer *two* protection paths for each of unit of protection ring capacity. A *p*-cycle of *N* hops and *M* straddling spans will have a logical spare to working ratio of *N/(N+2M)*. For example, above, the redundancy would be 50%.

p-Cycles can either be cross-connect based or based on an ADM-like "capacity slice" nodal device structure [1]. On OXCs *p*-cycles can be formed from individual spare wavelength channels offering the greatest flexibility to adapt and evolve the *p*-cycle configuration to protect dynamic demand. On the other hand, the ADM-like *p*-cycle nodal element offers the "pay as you grow" advantage of conventional rings. In either case, *p*-cycles avoid the structural association between the routing of working demands and the configuration of protection capacity.

Studies to date have considered self-organization of the *p*-cycle sets [2, 3], application of *p*-cycles to MPLS [4], and DWDM [5,6] and joint optimization of working paths and spare capacity [7]. In [5] it was found that as little as 39% redundancy was needed and [6] reports *p*-cycle designs using fewer OEO conversions and wavelength kms than other alternatives without significant availability penalty. Thus, *p*-cycles offer the "best of both" in the sense of attaining ring-like speed, with mesh-like efficiency.

3. Forcer-clipping Ring-Mesh Hybrids

Ring-mesh hybrids offer another way to get the best of both. The most common hybrid is in the form of rings accessing a mesh core. This is a loose hybrid, however, as ring and mesh domains are essentially separate subnetworks. We recently investigated the prospects for a more integrated hybrid involving strategic ring placements embedded within a span-restorable mesh network. Figure 2 gives an illustration of the basic concept and shows that certain ring placements can have a more than proportional benefit on the surrounding mesh in which they are embedded.

Figure 2(a) shows a 7 node, 10 span, span-restorable mesh of unit capacity modularity. (w,s) gives the working capacity on the span and the spare capacity assigned for 100% restorability (hop limit of five) from an optimal solution for the pure mesh design. The pure mesh has a total capacity of 283 channels. In (b) a 4-span 12 channel ring is overlaid. This represents an investment of 4 (2*12) = 96 (working plus spare) capacity units. The ring takes up all the working capacity it can from the mesh and the re-optimized residual mesh has 190 channels. Thus, the particular ring reduces the mesh by 93 in return for an equivalent capacity of 96 invested in the ring. The ratio, called the capacity return ratio (CRR) is = 93/96 = 0.969. The high CRR is an indication that this ring placement is very likely to be economic because, most rings cost less that 97% of the cost for the equivalent channel capacity and termination in a mesh environment. Thus, CRR numerically represents the discount factor for ring capacity (relative to mesh) below which a given ring placement would yield a net savings. Figure 2(c) is a counterexample where the CRR is 0.60, which is far less likely to be an economic ring placement.

So why is the ring in (b) cost-effective in the hybrid, but (c) is not? A general theory that explains this is called Forcer Analysis [8]. Forcer analysis reveals how the working span quantities drive the dimensioning of spare quantities needed for restoration. The best ring placements turns out to be good "Forcer clipping" rings [9]. The central idea of optimized ring-mesh hybrids based on this is that selectively placed rings can "clip off" the worst of these forcers, yielding network-wide relief in sparing of the residual mesh component, that more than pays for the ring placement. Especially where one ring can group together several strong forcers, there can be a more than proportional cost reduction in the remaining mesh network.

Results in [9] show that rings chosen by a forcer-clipping algorithm can yield up to 25% savings relative to pure-mesh design (depending on the network, demand, and relative ring-mesh costs), but that an all-ring network is also not optimal. The feasibility of such forcer-hybrids is increased by the advent of cross-connect systems with integrated ADM functionality, eliminating costs associated with a signal-transition from ring to mesh environments en-route through the hybrid transport network.

4. Ring Mining

Ring mining is a strategy for reaping maximum benefit from already-deployed rings while also obtaining the benefits of evolution to a mesh target architecture. Ring mining is based on logical "disassembly" and reuse of ring transmission spans to form a new mesh (or hybrid) target architecture [11]. A baseline strategy for ring-to-mesh evolution is to "cap the rings" and serve growth with a new mesh overlay. But in ring mining we view the installed capacity of rings as a sunk investment to be "mined." Ring mining was recently studied [11] on 17 fully-loaded ring network designs as test cases. One of the simplest experiments was to test the extent to which one could support additional demand with *no* added capacity, simply by breaking its spans up for reconfiguration as a mesh. In three of the test networks the result was up to 290%. In over 35% of test cases a complete doubling or more in demand could be supported with no additional capacity investment. In other test cases where the growth potential with zero capacity additions was not so high, the allowance of a small budget for new capacity additions serves as a catalyst to growth reusing the ring capacity. In the 17 test cases we saw complete deferral of new expense for factors of 1.5 to 3 times in service growth. Figure 3 shows selected results illustrating how capacity to serve growth is obtained from existing rings first significantly deferring additions of new capacity while making the evolutionary transition to a mesh.

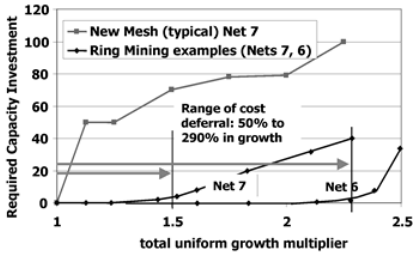

Figure 3: Typical Results of Ring Mining

Ring mining supports such potentially high growth multiples without new capacity additions by (1) freeing working paths from ring-constrained routes, (2) reclaiming the 100% ring protection capacity for more general use as mesh capacity, and (3) unlocking ring stranded capacity. In the most general case, a full transitional planning model based on the ring mining concept can minimize the total cost of serving growth through a ring to mesh or ring to hybrid target evolution, including costs of ring "disassembly".

In summary, ring mining, forcer-clipping ring-mesh hybrids, and *p*-cycles are all new networking ideas that in various ways obtain the best of both ring and mesh technologies which were previously almost always two exclusively separate options.

References

[1] W. Grover, D Stamatelakis, "Bridging the ring-mesh dichotomy with p-cycles", DRCN 2000, Munich.

[2] -------, "Cycle-oriented distributed preconfiguration: Ring-like speed with mesh-like capacity&...," ICC'98, 1998, pp. 537-543.

[3] ------, "OPNET Simulation of Self-organizing Restorable SONET Mesh Transport Networks",

OPNETWORKS '98 (CD-ROM), Wash. D.C., April 1998, paper 04.
[4] -----, "IP layer restoration &...," IEEE JSAC, Oct. 2000, pp. 1938 - 1949.
[5] D. A. Schupke, et.al., "Optimal configuration of p-cycles in WDM networks," ICC'02, NYC, 2002.
[6] L.Lipes, "Understanding the trade-offs associated with sharing protection," OFC 2002, ThGG121.
[7] W. D. Grover, J. E. Doucette, "Advances in optical network design with *p*-cycles:...," LEOS Topical Meetings, Quebec, July 15-17, 2002.
[8] W.D. Grover, D.Y. Li, "The forcer concept and its applications to express route planning in mesh survivable networks," JNSM, June 1999, pp. 199-223.
[9] W. D. Grover, R. G. Martens, "Optimized Design of Ring-Mesh Hybrid Networks," DRCN 2000, April 2000.
[10] M. Clouqueur, et. al. "Mining the Rings: Strategies for Ring-to-Mesh Evolution," DRCN 2001, Budapest, October 2001.
[11] M. Clouqueur, W.D. Grover, D. Leung, O. Shai, "Mining the Rings: Strategies for Ring-to-Mesh Evolution," Proc. 3rd DRCN 2001, Budapest, Hungary, October 2001, pp. 113-120.

TuI2 **2:30 PM**

Benefits of Restoration Signaling Message Aggregation

M. Goyal, *The Ohio State University, Columbus, OH;* J. Yates, *AT&T Labs - Research, Florham Park, NY;* G. Li, *AT&T Labs - Research, Florham Park, NJ;* W. Feng, *Oregon Graduate Institute, Beaverton, OR, Email: mukul@cis.ohio-state.edu*.

Signaling aggregation provides a practical approach to fast failure recovery in path-based restoration schemes. This paper proposes several signaling aggregation schemes and demonstrates their potential for improving restoration performance compared with schemes without aggregation.

1. Introduction

Path based restoration schemes [1] have emerged as a cost effective approach to achieve fast recovery from failures in optical networks. These schemes use signaling along restoration (backup) paths to re-establish connectivity after a failure disrupts the primary paths of the connections. Current restoration signaling proposals [1,2] can be characterized as "per-connection" in nature since each failed connection is restored using a separate set of signaling messages. As the number of connections affected by a failure increases, there is a corresponding linear increase in the number of signaling messages generated to restore these connections and hence in the queuing delays suffered by these messages at the optical cross-connects (OXCs) [2]. These queuing delays impact failure recovery times which become unacceptably high for moderately large number of connections in the network. Clearly, if a single signaling message could restore multiple connections, the number of signaling messages can be reduced, presumably with a corresponding reduction in the queuing delays and recovery times. In this paper, we propose several restoration signaling aggregation schemes that combine individual signaling messages into an aggregate message. We demonstrate significant improvements in recovery times achieved with these schemes in comparison to per-connection signaling.

2. Signaling Aggregation Schemes

While we experimented with several proposals for the path-based restoration signaling [1,2], in this paper we focus on the procedure described in [1]. In this scheme, the OXC detecting a failure sends a *failure indication* (or *alarm*) message to the source OXC of each of the failed connections. Upon receiving the *alarm*, the source OXC initiates failure recovery by sending a *switchover request* message (or simply a *request*) towards the destination OXC along the restoration path. As the *request* travels through the intermediate OXCs, they select channels for the failed connection. Upon receiving the *request*, the destination OXC generates a *switchover response* message (or a *response*), which travels back towards the source OXC along the restoration path. As the intermediate OXCs receive the *response*, they initiate channel cross-connection for the failed connection and forward the *response* further upstream without waiting for the cross-connection to complete. The connection recovery is completed when all OXCs along the restoration path have finished channel cross-connections for the connection.

The first step in signaling aggregation is aggregating the *alarms* generated by the OXC detecting a failure. Rather than sending individual *alarm*s for each failed connection, the *alarm*s going to the same OXC can be combined into a single *aggregated alarm*. We refer to this as *alarm aggregation*, which, besides reducing the number of *alarms*, allows a source OXC to simultaneously learn about multiple failed connections.

Among the failed connections identified by an *aggregated alarm,* the connections with the same restoration path can be restored using common *aggregated request* and *response* messages. We refer to this scheme as *aggregation over common path* signaling. Further reduction in the number of signaling messages can be obtained if, rather than aggregating signaling messages going over same end-to-end path, the OXCs aggregate signaling messages going to the same neighboring (next-hop) OXC. We refer to such a scheme as *aggregation over next hop* signaling. In this scheme, the source OXC groups the failed connections according to the next OXC along their restoration paths and sends a single *aggregated request* for each group. If the OXC receiving an *aggregated request* is the destination for a sub-set of the connections identified therein, it generates an *aggregated response* for these connections and sends it back towards the source OXC along their common restoration path. The remaining connections are re-grouped according to the next OXC on their restoration paths and the OXC transmits a single *aggregated request* message for each group. This procedure is repeated at all OXCs receiving the *aggregated request* message. Thus, an *aggregated request* splits as the restoration paths of the failed connections diverge. Note that *aggregation over common path* is actually a special case of *aggregation over next hop*.

Using the above procedures, an OXC may receive multiple *aggregated request* and *response* messages during a short period of time. If we allow a signaling message to wait for a short while at an OXC after it is processed, then it is feasible for the OXC to re-aggregate multiple aggregated messages based on common next hops. We refer to this approach as *aggregation over next hop with delay*. Note that *aggregation over next hop* is a special case of *aggregation over next hop with delay*, when the delay is set to 0.

3. Benefits of Restoration Signaling Aggregation

To evaluate the performance of restoration signaling aggregation, we implemented the signaling schemes described above in the NS2 simulator [3]. The signaling message processing times were obtained from measurements in the AT&T prototype testbed [4] and are listed in Table 1. We simulated several topologies that demonstrated similar conclusions; we present results here for the 21-node, 26-link ARPA2 network [5] with connections established between randomly selected source and destination nodes. Each simulation involved sequentially failing every link in the network and observing the recovery times for the affected connections; the reported results are the 90th percentile of the observed recovery times. It is assumed that the OXCs are capable of executing cross-connections in parallel - this was shown in [2] to be a prime requirement for fast restoration.

Simulation Parameters	
Processing delay for Request and ALARM messages.	0.418ms
Processing delay for Response messages.	0.326ms
Forwarding delay for a message in transit.	0.1ms
Channel cross-connection delay	2-3ms

Table 1 Simulation Parameters

We begin the performance evaluation by comparing the recovery times obtained with per-connection signaling with those obtained using the signaling aggregation schemes for different numbers of connections, as reported in Figure 1. As discussed earlier, the recovery times with per-connection signaling increase linearly with the number of connections. The *aggregation over common path* scheme dramatically improves recovery times by using common signaling messages to restore all connections sharing a common restoration path. As the number of connections increase from 2000 to 8000 in the ARPA2 topology, the average number of connections sharing a common restoration path increases from 4.8 to 19. Assuming that message processing times are unaffected by the increase in message processing load, the increase in number of connections has virtually no impact on the recovery times. The more sophisticated *aggregation over next hop* scheme provides further improvements in recovery times, as illustrated in Figure 2 which also shows marginally improved recovery times achieved with the *aggregation over next hop* when *response* messages are delayed by 1ms at each OXC to allow for greater aggregation. In general, the benefits of delayed signaling are marginal and it is difficult to determine the correct delay duration. Hence *aggregation over next hop* signaling appears to be the most appropriate practical choice.

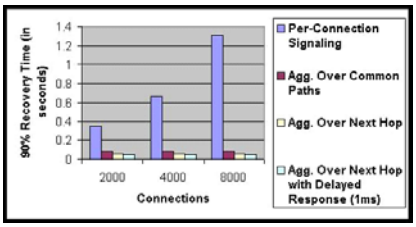

Figure 1 Comparing 90% Recovery Times obtained with Per-Connection Signaling and different Signaling Aggregation schemes.

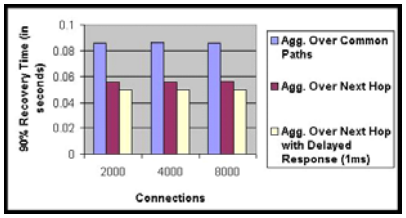

Figure 2 Performance of different Signaling Aggregation schemes.

The performance of the signaling aggregation schemes reported in Figures 1 and 2 is based on the assumptions that 1) the processing times for aggregated messages is the same as for non-aggregated messages; 2) there is no limit on the number of connections that can be signaled within a single message; and 3) the restoration path is calculated using shortest path routing. The first assumption is based on the observation made in other control plane protocols that the per-packet processing overhead dominates the total message processing [6]. The validity of the second assumption depends on how much connection-related information is carried in the signaling message. If it is only the connection ID, which would be adequate for shared mesh restoration