Reprint of Proceedings of the 19th Annual National Fiber Optics Engineers Conference (NFOEC 2003), Orlando, FL, USA, 7-11 September 2003.

# QUANTITATIVE COMPARISON OF END-TO-END AVAILABILITY OF SERVICE PATHS IN RING AND MESH-RESTORABLE NETWORKS

**Matthieu Clouqueur, Wayne D. Grover**
TRLabs and Department of Electrical and Computer Engineering, University of Alberta
7th Floor, 9107 116 St NW Edmonton, Alberta, Canada T6G 2V4
Ph: +1 780 441 3800, Fax: +1 780 441 3600
Contact: {clouqueur, grover}@trlabs.ca

*Abstract* – A comparison of ring and mesh architectures for restorable transport networking is presented from the point of view of service path availability. The comparison is based on detailed simulations of the network's response to random sequences of failures and repairs. Each type of network serves the same demands and are tested with exact implementations of the two restoration mechanisms in the face of the identical sequences of physical layer failure and repair. Results show significantly higher average service path availability in the mesh architecture. The study also shows the potential of the mesh architecture to provide very high availability to a small fraction of selected high-priority service paths when prioritization in the restoration is introduced, while keeping the availability of lower-priority service paths almost unchanged.

## 1. INTRODUCTION

Industry interest in mesh-based optical transport has increased greatly in recent years, from an era where ring-based transport dominated. Many network operators are now considering mesh as the way to go in the future [1]. The reason for this change is a growing appreciation of the advantages of the mesh architecture. The most widely recognized advantage is that mesh-based transport networks are considerably more capacity efficient than their ring counterparts [2]. This is attributable to the network-wide sharing of protection capacity over non-simultaneous failure scenarios and to the fact that service paths are not ring constrained – they can follow true shortest path routes over the graph. Rings not only require a minimum of 100 % redundancy but also suffer from "stranded capacity" effects in which working capacity cannot be as highly utilized as when capacity is placed along pure mesh growth principles. Mesh-based survivable networking also supports simpler provisioning, especially of multiple service classes, and easier accommodation of growth where it was not foreseen [3][4][5][6]. Mesh networks also scale more easily in the sense that growing a mesh network only requires the addition of capacity wherever capacity has been exhausted whereas rings require the addition of entire rings.

Recent work has also shown how gracefully and efficiently a span-restorable mesh can support multiple classes of protection and provide extremely high restorability to priority service paths [7][8]. In [8] it was shown that mesh-restorable networks exhibit very high average restorability against dual failures even when strictly designed only for single-failure restorability. This is attributable to the ability of mesh networks to use spare capacity in a very general way and, if also adaptive, to find restoration paths to the greatest extent possible under any circumstances. In contrast, rings (and 1+1 APS) provide only one predefined protection option and can never guarantee full dual-failure restorability to any paths because there are always dual failures that are guaranteed to bring down both the service path and its backup path.

This brings us to an interesting remaining question that industry colleagues have recently posed: "Which provides higher service-path availability – a ring or a mesh-based network?" So far the question seems to have been debated only qualitatively. One view is that rings are more redundant, spatially localized, and self-contained, so they should automatically be higher in availability. The opposite view is that although mesh is less redundant, it is also more general in its rerouting characteristics and can employ a failure-adaptive backup (even if more time is taken to do so), if initial pre-plans are overwhelmed. So in this viewpoint, mesh may do as well or even better. Another *a priori* view is that because rings switch in "50 ms", and mesh may take longer, rings should exhibit less outage time. This, however, is an easily corrected misunderstanding about what dominates the unavailability of either scheme. While the duration of a restoration "hit" may be of concern in its own right, it has virtually no influence on service availability, which depends almost wholly on the likelihood of an unrestorable dual failure, not the speed of response to restorable single failures (more on this point follows.)

A few past studies can be found addressing availability analysis of ring or mesh [7][9][10][11][12][13] but these studies are either devoted to one or the other architecture (not a comparison between them) or, where they do offer comparison, there are approximating simplifications that one or the other side of the comparison would find objectionable. Typical simplifying approximations are for instance that all paths in a ring are down if two failures hit the same ring (which is not the case), or, to model shared mesh restoration as approximately equivalent to 1+1 APS, or even to "smooth out" all details of capacity, topology, and mechanism to employ average-case Markov state transition modelling. Our aim has been to deliberately avoid any such simplifying approximations to provide an extremely precise and fair comparison using *exact* implementations of both survivability mechanisms. The study is thus essentially experimental in nature and is scrupulously detailed as to the exact re-routing mechanisms involved and how they interact with the topology, the spare capacity present, and the network state due to any prior failures.

### 1.1. Unimportance of the Restoration Speed to Availability

One thing we would like to do right away is establish that the reconfiguration time to restore single failures is essentially irrelevant in the debate about availability. Historically, great importance has been given to the issue of restoration speed with ring advocates saying "50 ms is essential" and mesh advocates saying "anything under a second is absolutely suitable." As a consequence it has often been assumed that fast restoration is required to achieve high availability. But we would like to demonstrate the fallacy of this assumption. Consider, as an example, a service path that undergoes 12 successful single-failure restoration events in five years (a relatively large number of such events) but then suffers an outage of 6 hours (half of a typical physical repair time) from a dual failure. We will calculate the unavailability of the service assuming (a) 50 ms, (b) 1 second [1] single-failure reconfiguration times. In (a) we have:

$$U_{\text{service}} = \frac{12 \cdot 50 \times 10^{-3} + 6 \cdot 3600}{5 \cdot 8766 \cdot 3600} = 1.36896 \times 10^{-4} \text{ or } A_{\text{service}} = 0.99986310$$

In case (b) we have:

$$U_{\text{service}} = \frac{12 \cdot 1 + 6 \cdot 3600}{5 \cdot 8766 \cdot 3600} = 1.36969 \times 10^{-4} \text{ or } A_{\text{service}} = 0.99986303$$

The point is that despite the intuition that fast restoration times goes hand in hand with high availability, there is virtually no practical connection between these measures. What absolutely dominates the availability of a restorable network is whether or not the service is exposed to an unrestorable dual (or higher-order) failure. Almost any imaginable number of successful reconfigurations to single failures can occur without the availability moving numerically from virtually unity. This is true whether those reconfigurations take 50 ms or 1 s. But if a single outage is experienced that relates to the time required to complete one of the physical repairs required under a dual failure scenario, then the availability is dramatically impacted. This was also demonstrated in [7].

We are not arguing that restoration time is not an important consideration in its own right. Obviously, if all else is equal, faster is better. But the point is that if one's customer wants to talk service availability, it has essentially nothing to do with restoration speed if any reasonable automated scheme is employed within a corresponding spare capacity environment that is designed to ensure full single failure restorability. What such clients absolutely need to avoid are multi-hour outages that could cause them losses of millions of dollars. As a performance measure, availability *does* reflect the latter concern but is utterly unresponsive to any practical differences in restoration switching times. The contribution of single-failure restoration times themselves are accordingly ignored in this study and only unavailability due to unrestorable multiple failures is taken into account.

We also do not mean to imply that mesh schemes are necessarily inferior to rings in restoration speed. Modern cross-connects have been designed with mesh restoration in mind as a key application and can reduce restoration times to little more than the network propagation times involved. In fact large BLSR rings are similarly limited by signalling propagation times. In practice, 250 ms is reportedly more typical of BLSRs [14], not the "50 ms" of folklore automatically associated with any kind of ring (only 1+1 APS and UPSR rings routinely achieve 50 ms.) Moreover, any remaining doubts about mesh restoration speeds are addressed by the Distributed Pre-Planning (DPP) concept [15]. Under DPP mesh restoration trials are exercized constantly in the background so that efficient and up to date pre-planned paths are already known in advance of a failure. With fast cross-connects, the real-time activation phase is then limited only by physical alarm propagation speeds.

---

[1] In practice many vendors will now state that pre-planned shared mesh restoration can, however, actually be implemented to operate within 250 ms or so, but we use 1 second just to be especially conservative for the sake of this numerical example.

### 1.2. Outline

From here, Section 2 describes the approach taken in simulating ring and mesh availability, detailing ring and mesh restoration behavioural models, the simulation method and the test networks used. In Section 3 we present and analyze the experimental results. In Section 4 we introduce a slight variation of the mesh restoration model that uses prioritization and we discuss additional experimental results obtained with that revised model showing how priority services can obtain extremely high availability. Section 5 concludes and discusses future work directions.

## 2. RING/MESH AVAILABILITY COMPARISON METHOD

### 2.1. The Two Survivability Mechanisms

The simulation of each survivability mechanism relies on the precise knowledge of the routing of all working paths and the reaction of all structures to single, dual, and higher order failures. In the ring case we know precisely which segment of each end-to-end path traverses each ring. When failures occur on a particular span each affected service path is restored following a BLSR-protection mechanism provided that another failure is not already employing the rings’s protection. Note that this *does* permit certain paths through a ring to survive dual failures, depending on how the two failures hit the ring. More formally the precise logic governing the outage of a path transiting a ring is as follows:

Let the set of $W$ spans and $(W+1)$ nodes in the normal route of a signal path segment X within the ring be called the forward path of X, For{X}. The set of other nodes and spans in the ring is then defined as Rev{X}, i.e., Rev{X} = {{R} – For{X}} where set {R} is the complete ring. Rev{X} contains $S$-$W$ spans, where $S$ is the number of spans on the ring. For outage of path X in the ring, it is then necessary and sufficient that one failed element belong to For{X} and the other to Rev{X}. This follows the detailed considerations given in [11]. Note, however, that unlike [11], this work is not concerned with special measures such as matched node interconnection between rings to avoid node failure. In the language of [11] this study considers only single-fed paths through ring-based networks, subject only to span failures.

The mesh restoration mechanism is implemented as follows. Upon a single failure, restoration is effected by rerouting affected paths around the failed span in a k-shortest paths like manner through the exact spare capacity present. This means that first all feasible paths on the shortest route are taken. Next, all paths feasible through the second shortest route around the failure span are taken, not using any spare channels of the first set of paths, and so on. This functional routing model goes on until either all failed paths are restored or no more restoration paths are feasible. This is known to be extremely close a single commodity maximum-flow solution for span restoration [16]. When a second failure happens, the reaction to that failure is functionally identical but takes into account any spare capacity usage from the first failure. Moreover, if a subsequent failure directly hits restoration paths of a first failure, the affected working paths and restoration paths are viewed as a single revised number of failed working channels for which restoration is sought in the reaction to the second event. Overall this results in the so-called “partly adaptive” span-restoration behaviour in [7]. All channels bearing working service are initially considered equally important for assignment of the available restoration paths. Later, tests with priority assignment give available restoration paths to priority service paths first. In all cases the environment of spare capacity in use and spare capacity available for restoration of any failure is always updated to reflect current usage arising from any previous still outstanding span failure(s). When any failure state terminates, all related working paths are returned to normal routing, the associated restoration paths are collapsed, and their spare capacity is returned to an available spare state. Figure 1 illustrates some of the key concepts of the two survivability mechanisms under comparison.

Part (a) and (b) of Figure 1 show both architectures under normal operations with two identical service paths for illustration of effects. In the ring architecture, path A traverses one span of the ring before travelling on and path B is completely contained between nodes of the ring. Part (c) and (d) of Figure 1 show the reaction of both mechanisms to the same first failure. In the ring architecture, For{A} of path A is the single-hop segment between the upper two ring nodes. When failure hits that span, it is restored using path Rev{A} through the backup capacity around the rest of the ring. In the mesh, path A is restored by assigning it to one of the set of dynamically found “ksp” restoration paths described above. For clarity in Figure 1 we only show one path on the affected span but there is no implication intended that the mechanism restore all affected service paths over a single restoration route (this is sometimes a misconception about span restoration.) At the same time as the one rerouted path shown takes its route, many other paths (not shown but affected by the same failure) would also be restored over other routes of the ksp path-set.

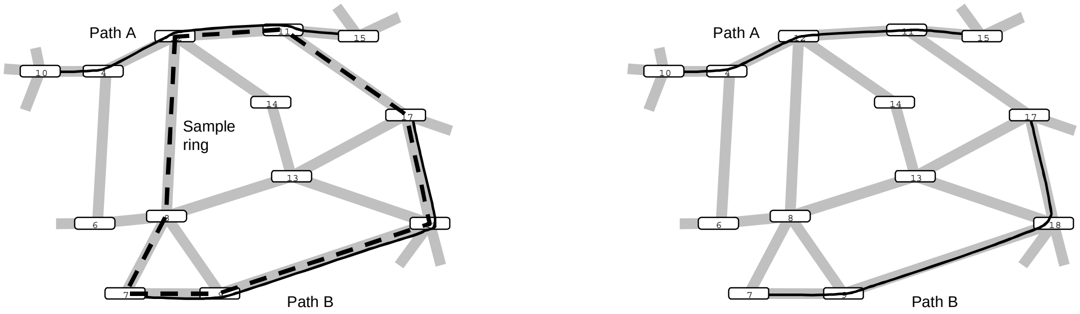

(a) Ring - Normal Operation

(b) Mesh - Normal Operation

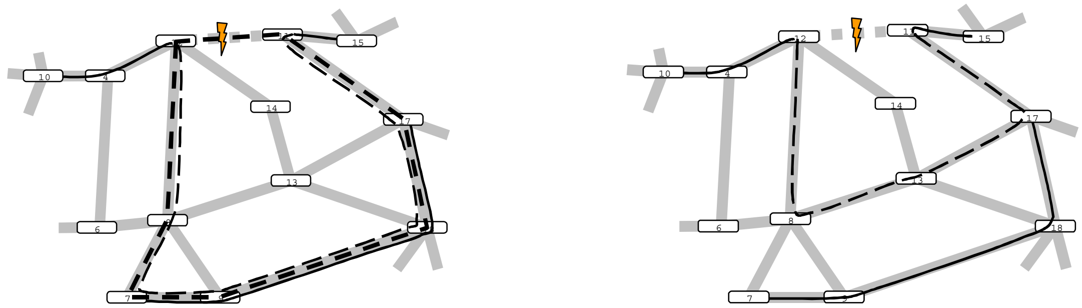

(c) Ring - First Failure

(d) Mesh - First Failure

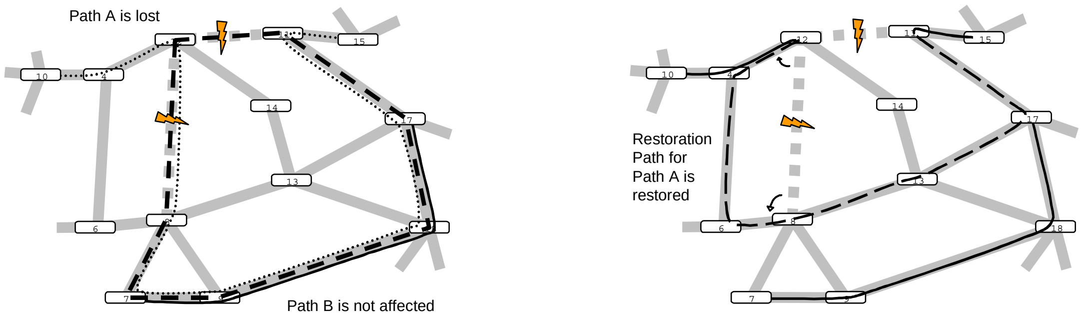

(e) Ring - Second Failure

(f) Mesh - Second Failure

Figure 1. Illustrating some details of the ring and mesh mechanisms under comparison

Part (e) and (f) of Figure 1 show the reaction of both mechanisms to a second failure. This second failure hits a span in Rev{A} and thus brings down service path A. That path is now experiencing hard outage, although path B remains unaffected in the same ring and is still routed over its normal working channels. The second failure also hits the restoration path of path A for the mesh architecture. In this case the failed restoration path is unified with the rest of the failed working paths on the second failed span and another maximum feasible ksp-type restoration path-set is developed in response to the second failure in the presence of the first span failure and the already used spare capacity removals from the network. Now in the mesh architecture, path A may or may not be restored following the second failure — it all depends on what is feasible within the remaining spare capacity following the first failure. If the number of restoration paths for the second failure is less than fully required, it also depends on whether path A has either luck or priority in the assignment by the failure end nodes of such restoration paths as are feasible for the second failure. As drawn we illustrate the possibility that the restoration path of path A is itself re-restored on the second failed span, showing how service paths can survive dual failures in the mesh. The example shows that the two architectures have very different responses to failures and that in neither architecture does a dual or triple etc., failure necessarily mean outage for service paths. Not shown are the cases of dual failures that hit different rings or are spatially separated enough in the mesh that they do not even interact. Neither of these cases is outage-causing to any paths.

### 2.2. Simulation Method and Test Networks

For the simulation, importance is placed on making a true "apples-to-apples" comparison between the two architectures. To achieve this the following conditions were set:

1) Ring and mesh are compared on identical facilities graphs serving identical end-to-end demands: The two topologies displayed on Figure 2 were used. The net32 topology, which has 32 nodes and 45 spans, was taken from [17] and was tested with two different demand matrices. The first demand matrix and the corresponding ring design were also taken from [17]. This first demand matrix corresponds to a hub-type of demand where all demands originate or terminate at one of three different hub nodes. The corresponding test case is denoted net32-A. The second demand matrix was created for the present study by generating random demands between nodes in proportion to the product of their nodal degrees (gravity-based approach) and serving them within the same ring set as net32-A following ring-constrained shortest path routing until the first demand request was blocked. This resulted in a demand pattern that achieves about 78 % loading of working capacity in the corresponding ring set. The corresponding test case is denoted net32-B. The same end to end demand pattern is used for the corresponding mesh network design on the topology of net32. The second topology is an arbitrary manually designed topology that has 25 nodes and 50 spans. The corresponding demand matrix was generated using a purely deterministic gravity-based model.

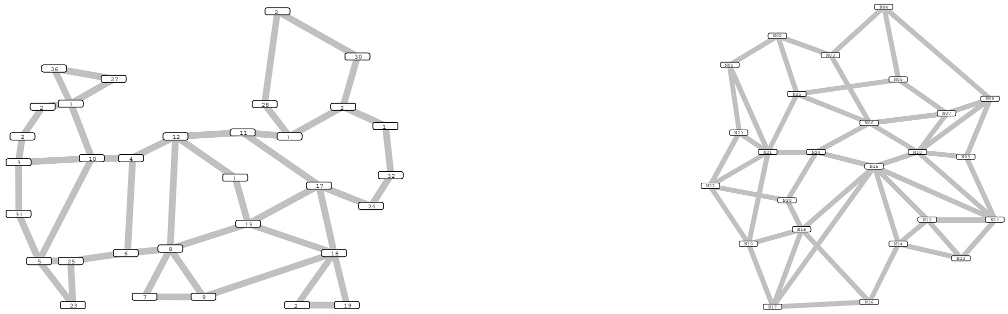

(a) net32

(b) 25n50s1

Figure 2. Test networks topologies

2) Efficient fully restorable designs are used for both architectures: The ring networks are based on efficient sets of BLSR ring placements and routing of demands through the rings. The ring design used for net32 (identical for net32-A and net32-B) is taken from the work of [17] and is based on the *fcrip* design method. The ring design for 25n50s1 was produced for this study using the Tabu Search method [17][18]. The mesh designs were produced using the Modular Joint Capacity Placement (MJCP) linear programming formulation in [19] using ten eligible working routes for each demand pair and ten eligible restoration routes for each span-failure scenario and the same capacity modularity and end to end demands used by the ring designs (a modularity of 48 in all designs).

3) Exact survivability mechanisms are emulated: As explained in Section 2.1, exact emulation of both survivability mechanism is performed, taking into account the capacity usage of all capacity units and producing a detailed report of the status of all service paths after each failure or repair event.

4) Both architectures experience identical span-failure sequences: A sequence of random physical layer failures and repairs is generated and both architectures are tested with it. The sequence is obtained by first generating a unique timeline of failures and subsequent repairs for each span as an independent entity and then combining them all into one composite history of failure and repair for the whole network. Note that under these circumstances, nothing prevents the network from experiencing dual, triple or higher order failures. In addition, this type of simulation is effectively at steady state the instant it starts because the up/down state probability of each span already reflects the long-term average when its individual failure/repair history was generated. Thus, there is no need for a period of transient simulation to reach steady state in this approach to simulating the failure environment. For simplicity, it was assumed that each span in the network had the same MTBF of 1 year with times to failure being negative-exponentially distributed. Times-to-repair were also negative exponentially distributed with a mean

of 12 hours. The simulation tool behind the work can however be easily employed to model length-dependent failure rates on each span and/or any particular distribution of repair times.

Table I summarizes some details of the test networks used. The working capacity totals correspond to the sum of working channels along all service paths in the network. The total capacity values correspond to the total number of capacity modules in each design multiplied by the module capacity. Confirmation appears right away of the great difference in capacity efficiency between ring and mesh designs. The ring designs require much more capacity than the corresponding mesh designs. It is particularly noticeable with the 25n50s1 test case, in which the topology is favourable to the mesh architecture. With the less connected net32 topology, mesh is also less efficient (especially with the hubbed demand) but still requires significantly less capacity than ring. Also, path lengths are slightly higher for mesh in the net32 topology. This sparse topology tends to force paths to go more out of their way to enhance spare capacity sharing under MJCP design, which is a joint optimization model. In contrast in the highly connected 25n50s1 topology paths are longer for the ring design. It will be seen later that path length has a considerable importance in terms of availability. With such total capacity differences, one could wonder whether mesh can provide comparable availability to that of the ring architecture. The following section answers the question.

Table I – Details of Test Network Designs

| Network Design | Demand | Number of Rings | Working Capacity | Total Capacity | Average Path Length | Network Redundancy |
|---|---|---|---|---|---|---|
| net32-A Ring | Hubbed | 12 | 1824 | 8448 | 5.15 | 363 % |
| net32-A Mesh | Hubbed | - | 1943 | 4656 | 5.48 | 140 % |
| net32-B Ring | Rand./grav-bsd | 12 | 3304 | 8448 | 2.01 | 156 % |
| net32-B Mesh | Rand./grav-bsd | - | 3500 | 7152 | 2.14 | 104 % |
| 25n50s1 Ring | Gravity-based | 19 | 4966 | 17520 | 3.07 | 253 % |
| 25n50s1 Mesh | Gravity-based | - | 4574 | 6960 | 2.83 | 52 % |

## 3. EXPERIMENTAL RESULTS

As the results are essentially experimental in nature, let us first address issues of statistical confidence. All results are based on a thousand simulations of one year of the network’s life. As a check on confidence intervals, we inspected the variance of the point estimation of one of the lowest overall unavailability predictions resulting from the study. This is for the 25n50s1 test network in the mesh case for the average path unavailability independent of path length. The result was that for the ensemble of 1000 one-year simulations $U_{ave} = 3.297\times10^{-5}$ with std dev = $8.91\times10^{-7}$ which is 2.7 % of $U_{ave}$. For higher unavailability values the accuracy of estimation is only increased because such cases are based on more outage-contributing events. In addition, the total number of dual or higher-order failures (over the 1000 one-year trials) was 2619 for the net32 topology and 3180 for the 25n50s1 topology. Each of the individual span failure/repair sequences was produced using the negative exponential random number generator (RNG) in [20]. With suitable scaling, the generator is called alternately to generate the time-to-next-failure and the time-to-repair for each span until the total simulation period has been reached. The number of calls that this produces to the RNG is much lower than its period (in the order of $10^8$) so we perceive no risk of correlation between sequences for the different spans.

For each design, the following measures of the resulting availability were computed:

- Average path unavailability: This is fraction of time that a given service path is experiencing outage. Because it is reasonable to expect that individual path availability’s depend on path length, results are presented as a function of path length in terms of number of hops in the path. Each data point presented is the average over all paths of a given length and over all 1000 test periods.
- Average number of outages per year (network total) and per path: These values indicate the probability of each service path of experiencing an outage during a given one-year period and characterize the total number of outage-causing events the operator could expect over the network as a whole per year. This is also presented versus path length.
- Statistical Frequency of total path outage times per year: This data summarizes the proportions of service paths expected to undergo no outage over a one-year period and the fraction of all paths that experience various higher levels of annual outage. This also allows visibility of the worst case scenarios experienced by any path at all.

In Figure 3 we can see that in all cases the unavailability of paths in ring networks is significantly higher than the paths of same lengths in the mesh networks. The difference in unavailability also increases as path length increases. The difference is about a factor of two at the longest paths. For the net32 topology, this difference is reduced slightly if we take into account the fact that service paths are on average slightly longer in the mesh design. For the 25n50s1 topology however, which already shows the biggest advantage for mesh, the difference between the two architectures is even higher if we consider the fact that service paths are on average longer in the ring architecture. The curves on the right of Figure 3 are almost identical in shape but they show the expected number of outages per year and indicate an expected number of outages per path per year reaching a maximum of about 0.17 for mesh and 0.26 for ring. If, for the sake of the argument, we neglect the probability of two outages happening to the same service path in a year, this is close to saying that there is about a 17 % chance of the longest mesh paths, and about a 26 % chance in the ring designs, to experience an outage during any given year.

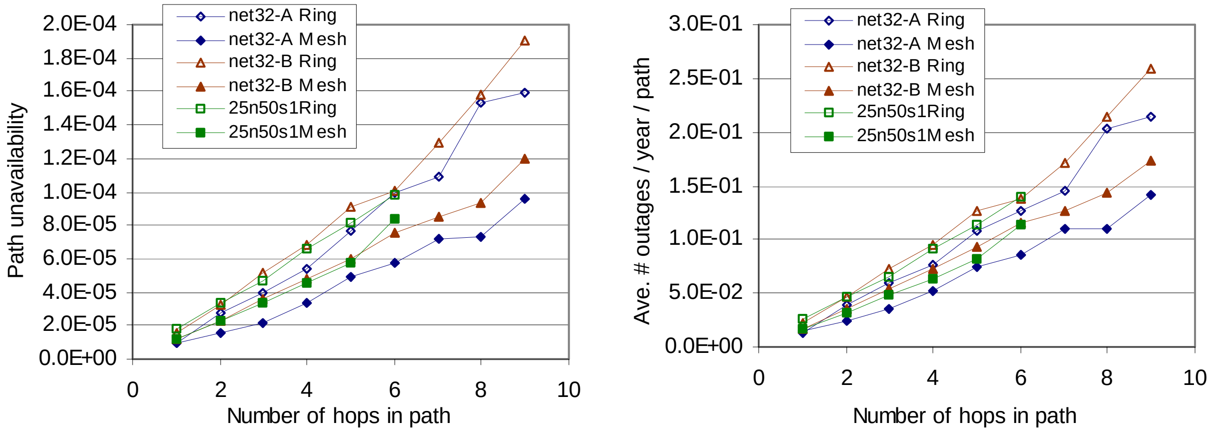

Figure 3. Simulation results for the three test networks

Figure 4 shows the distribution of actual outage times experienced by paths in both architectures for the 25n50s1 test case with annotation of the expected fraction of paths that experience no outage in any one year. Results show a slightly higher proportion of paths expected to experience no outage in a year with mesh (95.4 % for mesh vs. 93.7 % for ring.) The worst-case experience of any one path is also 72 hours of total outage in one year compared to 48 in total for the worst-case mesh path. Similar results are observed for the other test cases.

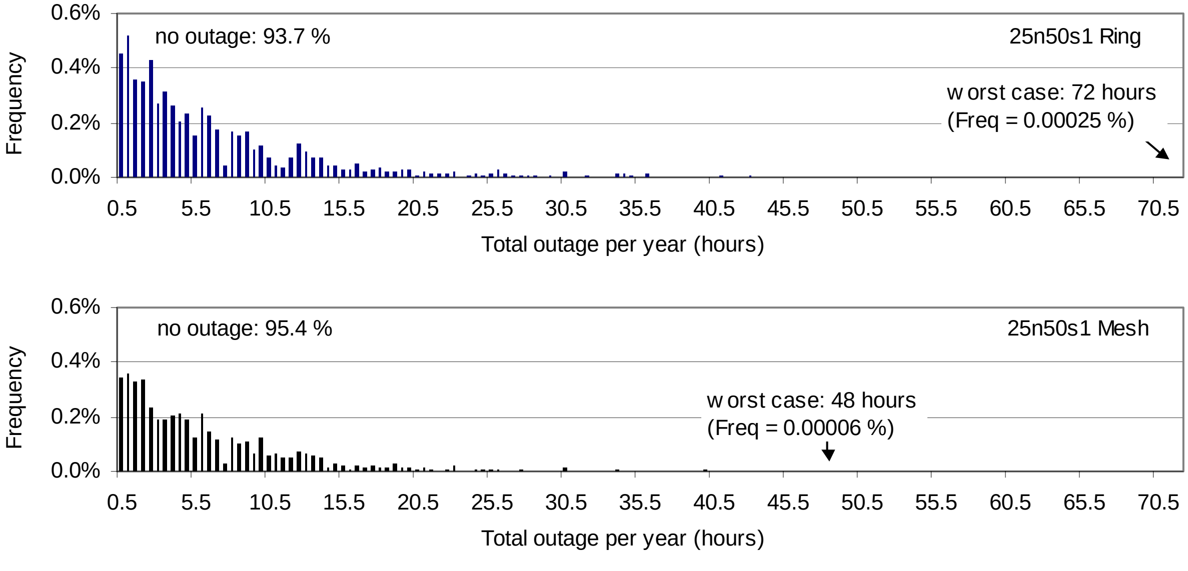

Figure 4. Distribution of outage times for test case 25n50s1

Overall, these results show an advantage for mesh on all availability measures. Note, however, that even with mesh restoration long paths experience relatively high probabilities of an outage during any given year. For critical services it would clearly be desirable to have a probability of outage lower than 17 % per year. This is the motivation for introducing *prioritization* in mesh restoration in the next section. This is an option that only has meaning in the context of the mesh restorable architecture. With ring protection, when failure occurs on one span of a given ring, all affected paths are either restorable or not because protection switching inherently occurs at the optical line rate, not at the individual channel level where priorities can be implemented in cross-connect based mesh restoration.

## 4. EFFECT OF MESH RESTORATION WITH PRIORITIES

### 4.1. Concept of Priorities

Unlike in previous studies where different classes of service were treated differently from the point of view of the capacity *design* itself [8][21], the present distinction between the high-priority and the low-priority paths is purely from an operational point of view. Both classes are still guaranteed full restorability to any single failure by design. The difference is only that the high priority service paths are always considered first for restoration in any circumstances where full restoration may not be possible for all affected paths. High priority service paths are still not guaranteed dual-failure restorability by design (as in [8]) but they have a higher chance of being restored. To test the effect of applying priority distinctions, we tag a certain fraction of the service paths on each demand pair as "high priority." We tested three service mixes: "10/90", "30/70", "50/50" where the first number is the proportion of high-priority service paths.

### 4.2. Experimental Results for Mesh with Priority Service Classes

Figures 5-7 show the effects of priority for the three test cases. In all cases we see a clear improvement in availability for the high-priority service paths. As expected, the improvement is greater for the priority class when fewer paths have high-priority status. A more surprising outcome is that giving high priority status to some service paths does not affect the availability of the remaining paths very much. In fact, when the proportion of high-priority services is only 10 %, the priority group benefits very noticeably, but the availability of the low priority paths is almost unchanged. In other words, the Mesh LP 90 curves in Figures 5-7 are almost identical to the basic unprioritized results for mesh in Figure 3. This is an interesting result from a revenue point of view because it indicates the prospect of selling a fraction of services at a higher price without having to reduce quality to other services. One explanation for this effect is that the small high priority group benefits mainly from the coherence of applying available restoration path resources in a consistent, as opposed to random way. In other words to provide the best service for a few, opportunities to avoid outage must be consistently applied to them, but this makes little difference to the remainder who receive their opportunities (in the event of any shortage) on an essentially random basis of allocating the hardship. In this thinking the randomness has more to do with the experience of the non-priority group than does the preference given to the smaller priority group when it is possible and necessary to do so for them.

Another interesting result is that the effects of priorities appear to depend on the type of topology. The two test cases using the net32 topology show an improvement for the priority services that is less pronounced than in the 25n50s1 test-case. The difference in availability improvement between the demand mix scenarios is also not as high as in the case of 25n50s1. The most dramatic effect is seen in the highly connected 25n50s1 network. In particular, the 10/90 demand mix shows an almost six-fold reduction in unavailability in the best case and the

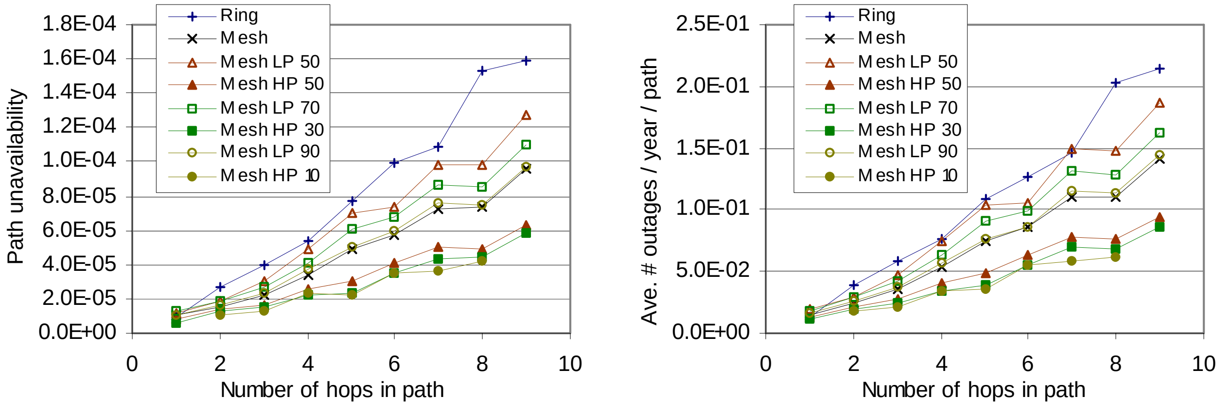

Figure 5. Results for test network net32-A

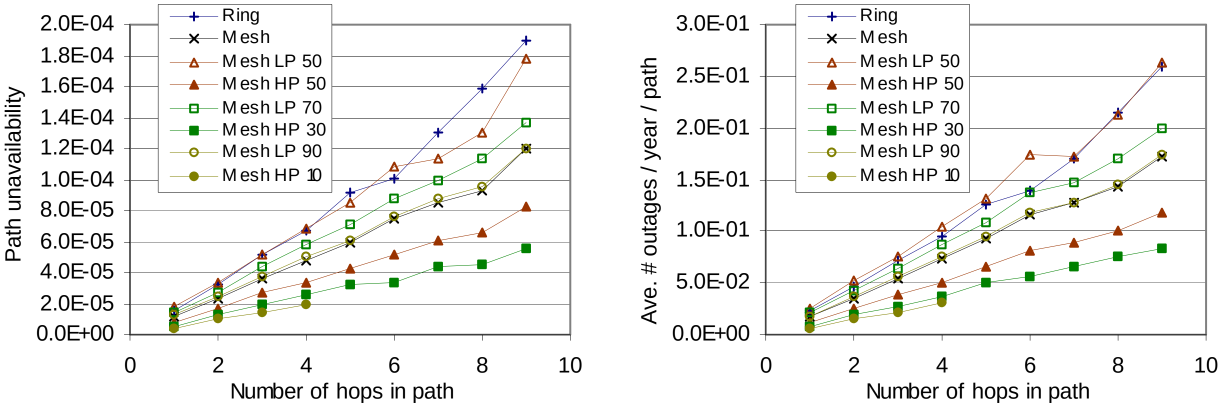

Figure 6. Results for test network net32-B

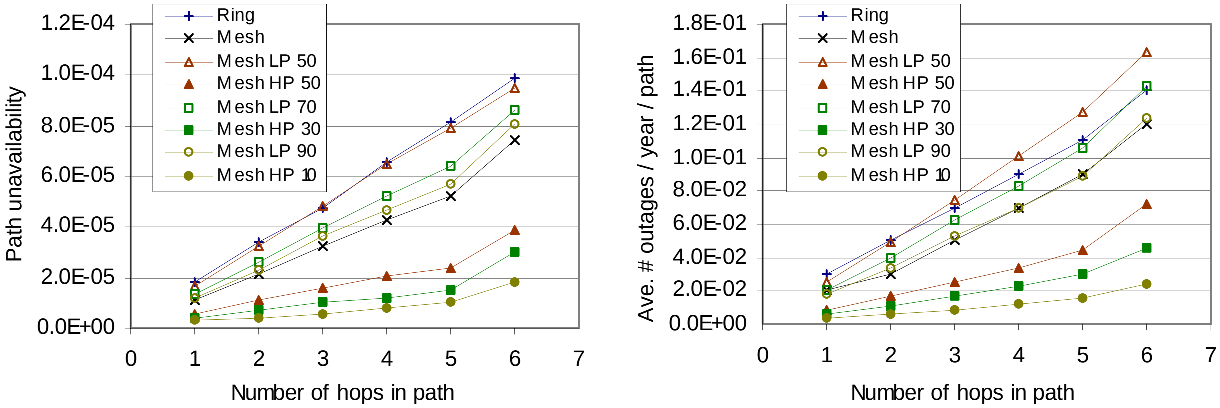

Figure 7. Results for test network 25n50s1

availability of the 90 % of low-priority paths is virtually unchanged. The probability of experiencing an outage in a given year is in that case reduced to about 2 % in the worst case (compared to about 12 % without priorities).

Finally Figure 8 shows the distribution of outage times for the high-priority class in the 25n50s1 test case with the 10/90 demand mix. The improvement is very clear: the expected proportion of paths in that class experiencing no outage in a given year is now 99.2 % (compared to 95.4 % without priorities) and the worse case scenario of total annual outage is reduced from 48 to 40 hours (happening with a probability of 0.0017 %).

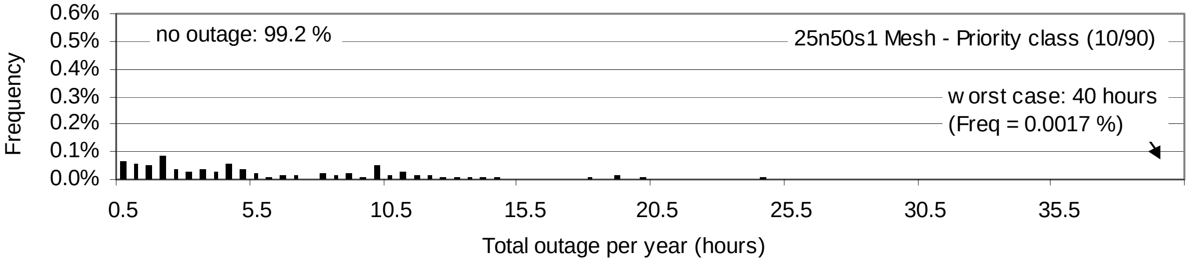

Figure 8. Sample results of total outage distribution for 10 % priority class in mesh

## 5. CONCLUDING DISCUSSION

This work has sought to address an open question about ring and mesh network availability through a carefully controlled study in which the network and failure environments are absolutely identical for each alternative and where we mechanize the exact reaction of each architecture in the face of failures. Results are sufficient to put aside the argument that because of its lower capacity requirement mesh-restorable networks cannot be as high-availability as rings. In fact, what we actually see is that mesh flexibility wins out over ring redundancy largely because of the “locked up” nature of the ring investment in protection capacity. Results show that service paths enjoy significantly higher availability in the mesh architecture despite the much lower capacity requirements of the mesh designs. The advantage is even more pronounced in the more connected test topology. The study also investigated the availability of mesh service paths with two levels of priority. This can only be implemented with the mesh architecture and has the advantage of enhancing the availability of high-priority paths even further. If the proportion of high-priority paths is about 10 % the availability of low-priority paths is almost unaffected while the availability of high-priority paths is greatly increased.

It must be acknowledged that node failures where omitted from this study. This partly is justified by the high internal redundancy already provided in optical cross-connects and ADMs and the security measures taken to enforce the very high reliability of network nodes. Taking node failures into account would therefore not have significantly affected the results and the main finding that mesh does at least as well as rings in terms of availability.

We believe that this study provides yet more confirmation of the great potential of mesh-restorable networks and clearly shows that there is no simple link between network redundancy and availability. High redundancy is not sufficient to provide high availability. Based on this work an increasingly clear view is emerging that what is ideal for very high availability is to be a priority service path in a highly connected mesh-based restorable network that employs an adaptive response to any circumstances where a pre-planned single failure response is overwhelmed. When these conditions are met then more redundancy could reasonably be indeed expected to translate into higher availability. This is a subject of interest in further studies in this area.

## REFERENCES

[1] B. Mukherjee, “WDM optical communication networks: progress and challenges,” *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, October 2000, pp. 1810-1824.

[2] W. D. Grover, “Case studies of survivable ring, mesh, and mesh-arc hybrid networks,” in *Proceedings of IEEE GLOBECOM 1992*, Orlando, FL, USA, December 1992, pp. 633-638.

[3] R. D. Doverspike and B. Wilson, “Comparison of capacity efficiency of DCS network restoration routing techniques,” *Journal of Networks and System Management*, vol. 2, no. 2, June 1994, pp. 95-123.

[4] O. Gerstel and R. Ramaswami, “Optical layer survivability – an implementation perspective,” *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, October 2000, pp. 1885–1899.

[5] N. Mandal, “Ringbusters: With optical mesh carriers can stop going in circles for «new-world bandwidth»,” Fiber Optics News, June 2001, http://www.tellium.com/documents/RINGBUSTERS_FINAL.pdf

[6] W. D. Grover, D. Li, “The forcer concept and applications to express route planning in mesh survivable networks,” *Journal of Network and Systems Management*, vol. 7, no. 2, June 1999, pp. 199-223.

[7] M. Clouqueur, W. D. Grover, “Availability analysis or span-restorable mesh networks,” *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 4, May 2002, pp.810-821.

[8] M. Clouqueur, W. D. Grover, “Mesh-restorable networks with complete dual failure restorability and with selectively enhanced dual-failure restorability properties,” *Proc. OptiComm 2002*, Boston, July 2002, pp. 1-12.

[9] P. Arijs, W. Van Parys, P. Demeester, “Cost and availability comparison of WDM mesh and ring network architectures,” in *Proceedings of DRCN 2000*, Munich, Germany, April 2000, pp. 33-38.

[10] M. Bolduc, S. Baddipudi, J. Babbitt, A. Wright, “Optical backbone architecture comparison,” in *Proceedings of NFOEC 2000*, Denver, CO, USA, August 2000.

[11] W. D. Grover, “High availability path design in ring-based optical networks,” *IEEE/ACM Transactions on Networking*, vol. 7, no. 4, August 1999, pp. 558-574.

[12] D. A. Schupke, “Reliability models of WDM self-healing rings,” *DRCN 2000*, Munich, Germany, April 2000, pp. 174-180.

[13] H. C. Cankaya, V. S. S. Nair, “Reliability and availability evaluation of self-healing SONET mesh networks,” in *Proceedings of IEEE GLOBECOM 1997*, Phoenix, AZ, USA, November 1997, pp. 252-256.

[14] K. Raychaudhuri (Lucent Technologies), Panel Discussion at *IEEE Bandwidth Management Workshop IX*, Montebello Quebec, 25-28 June 2001.

[15] W. D. Grover, “Distributed restoration of the transport network,” *Network Management into the 21st Century*, T. Plevyak and S. Aidarous, Eds. Piscataway, NJ, IEEE/IEE Press, February 1994, ch. 11, pp. 337-417.

[16] D. A. Dunn, W. D. Grover, M. H. MacGregor, “A comparison of k-shortest paths and maximum flow methods for network facility restoration,” *IEEE J. on Selected Areas in Comm.*, vol. 12, no. 1, January 1994, pp. 88-99.

[17] G. D. Morley, Analysis and Design of Ring-Based Networks, Ph.D. thesis, University of Alberta, Spring 2001.

[18] G. D. Morley, W. D. Grover, “Tabu Search Optimization of Optical Ring Transport Networks,” in *Proceedings IEEE GLOBECOM 2001*, San Antonio, TX, USA, November 2001, vol. 4, pp. 2160-2164.

[19] J. Doucette, W. D. Grover, “Influence of modularity and economy-of-scale effects on design of mesh-restorable DWDM networks,” *IEEE JSAC*, vol. 18, no. 10, October 2000, pp. 1912-1923.

[20] W. H. Press et al., Numerical Recipes in C - The Art of Scientific Computing, Cambridge University Press.

[21] W. D. Grover, M. Clouqueur, “Span-restorable mesh networks with multiple quality of protection (QoP) service classes,” to appear in *SPIE Optical Networks Magazine*, Fall 2003.