from Proceedings 17th National Fiber Optic Engineers Conference NFOEC 2001, Baltimore, July 8-12, 2001

QUANTIFYING AND MANAGING THE INFLUENCE OF MAINTENANCE ACTIONS ON THE SURVIVABILITY OF MESH-RESTORABLE NETWORKS

Wayne D. Grover, **Matthieu Clouqueur**
TR*Labs* and University of Alberta
{grover,mclouque}@edm.trlabs.ca

Thuthuy Bach
Nortel Networks
tbach@nortelnetworks.com

Summary – In an extensive quickly growing DWDM transport network the frequency of in-service upgrades or other required maintenance activity may be high. We are therefore interested in defining, quantifying and managing the theoretical risk that a "maintenance outage" may have for the rest of a mesh-restorable network. The point of view is of putting one span in maintenance and then asking, *if* a failure occurs to another span while in the maintenance state, what is the risk of incomplete restoration? In ring networks, the corresponding exposure to failure is contained within the specific ring involved. But the price for this containment is the 100% redundancy of rings and that the risk is of 100% restorability loss on other spans when in a maintenance state. We find that mesh networks exhibit a more distributed risk field, but much lower risk magnitudes. In over 2500 trials with a mesh equivalent of roll-to-protection as the maintenance model no other span sustained over 70% theoretical loss of restorability and 99.6% of the time the theoretical risk to restorability was < 50%. This gives an important operational advantage in support of priority service paths. We give a procedure for calculating the 'theoretical risk field' of a given maintenance action in a mesh-restorable network. We then study the effects of various mesh-network design differences on the extent and magnitude of the theoretical risk field. We produce results for a worst-case ("equivalent to failure") model for maintenance and the mesh equivalent of "roll to protection" in rings. The basic methods can be adapted for other maintenance models or models of the restoration process or adapted to handle multiple maintenance scenarios. The work is aimed at understanding and managing the effects of maintenance in a mesh-restorable network and ultimately providing new operational and design capabilities for future mesh-based network planning and operations systems.

I. Motivation and Background

One of the motivations for mesh-restorable networking is the greater capacity efficiency of mesh networks relative to ring-based transport networks. The capacity efficiency comes from the more direct routing of working paths, the need for less spare capacity for restoration, and the avoidance of "stranded capacity" effects in rings where one or more ring spans may exhaust while other spans of the ring have valuable but unusable remaining working capacity. Of course capacity efficiency is not the only consideration in a choice of network architecture, but in an increasingly competitive environment the combination of capacity efficiency, ease of growth, and service-provisioning flexibility that a mesh network can offer may provide certain carriers with a productivity edge over their competitors.

Why would a more capacity-efficient mesh imply corresponding dollar-cost savings? One factor is optical networking technology based on DWDM: The distance-dependent transmission costs for EDFAs, plus transmission-related nodal termination and routing / switching equipment quantities can place a significant emphasis on capacity efficiency at the lightwave system utilization level. If relatively low-cost Gb/s routers replace circuit switches and cross-connects in the nodes, but handle thousands of times more traffic over DWDM transmission, there is a shift the historical balance of transmission versus switching costs back to where transmission costs *do* matter. In addition, sustained rapid growth of demand creates an environment where there is a premium for capacity efficiency *even if* the cost of capacity was zero, driven purely by deployment speed as the rate-limiting step to earnings growth. If the optical transmission capacity can barely be provisioned quickly enough to keep up with demand then at any point in time the more efficient architecture will be serving more demand (and earning more revenue) for a given installed base of transmission systems. Despite the perennially fashionable exaggeration that "bandwidth is free", the fact is that yearly *incremental* capital expenditures of 300 to 400 million $US for transport equipment is not untypical. If, say, a 40% increase in capacity efficiency attributable to mesh architecture advantages leads to even a 20% savings on a budget like that, then "going mesh" would seem to be a well-qualified project for corporate productivity enhancement and profit.

Having thus motivated the benefits of mesh restoration, there are some costs of entry. Mainly, the company that operates a mesh-restorable network needs to be able to accept somewhat more sophisticated and automated operational processes. This can be as much a cultural barrier as technical. At least technically, however, there is now a considerable literature on mesh-restorable network design, distributed restoration, service provisioning, restorability-audit, reversion after failure, automatic integration of new service paths into the restorability plan, and so on. See for instance [1], [2], [3], [10] and references therein. In further support of the viability of mesh-based survivable networking, half a dozen optical networking start-ups have made the technical and business judgements that there is enough accumulated knowledge to now proceed with optical mesh-restorable networks and their corresponding operational support systems.

One operational aspect of mesh-restorable networking that has apparently *not,* however, received attention, is the way in which maintenance activity on mesh spans may create a theoretical exposure to reduced restorability if a failure occurs elsewhere during the maintenance state. In rings, the corresponding properties are very clear: The set of spans at risk from maintenance on one span is all other spans in the same ring and the magnitude of the risk is 100% loss of restorability[1]. For span-restorable mesh networks, the risk field will be more diffuse and generally less than 100% in magnitude on any span. A main aim of this work is to characterize the corresponding theoretical risk field for span maintenance in a mesh network. Of particular interest is the prospect that in a mesh-restorable network it may be feasible to *guarantee zero risk due to maintenance actions for high priority service paths*. This is an important question that this work confirms.

A. Possible goals for managing maintenance-related restorability risk

The overall theme of this project is to support the prospect of mesh-based restorable networking with means to analyze and then explicitly control and manage the impact of maintenance–related withdrawals of spare capacity on the restorability of a mesh network. In increasing steps of complexity we conceive of the following measures or goals in this regard: This paper is devoted primarily to the first goal. We touch briefly on considerations related to the second goal in closing the paper. The third goal is left to future research.

(i) A first requirement is a basic analysis capability to assuring that the network-wide extent of "exposure to a subsequent failure" of a proposed maintenance action is at least known in advance by some analysis. Such an initial tool or capability would simply tell a network operator what the extent of theoretical exposure is to a single-span failure anywhere else in the network, while doing a certain maintenance action. The "extent" of influence of a maintenance action can be quantified by a vector that indicates, for each other span, the loss of restorability of the other span under the condition of a single real failure arising while in the maintenance state. Given a proposed maintenance action, the tool could give a diagnosis of any particular other spans that would accrue some non-zero restorability risk or simply flag other spans sustaining risk above some threshold of significance due to the proposed maintenance.

(ii) A second operational capability could be to be able to assure that multiple simultaneously scheduled maintenance outages are in some sense (to be further defined) "mesh-orthogonal" in that the geographical / topological ranges of failure exposure do not overlap or compound each other. This type of future tool or capability would tell a network operator what combinations of simultaneous maintenance activities avoid compounding any theoretical risk or possibly *avoid maintenance-*

---

[1] In either a BLSR or UPSR type of ring, or their DWDM counterparts, maintenance that rolls working to protection, or in any other way uses the tributary (UPSR) or line-level (BLSR) protection capacity fully, makes all other spans (or tributaries) unrestorable (100% restoration risk) for the duration of the maintenance action. Note that this is not to say that no demands will survive in the ring, only that all demands crossing the failure span will be unrestorable.

related risk entirely for priority service paths. If X and Y are single span maintenance projects, the aim would be to know if the combined (parallel) maintenance plan (X||Y) runs a significantly greater overall risk than is inherent in serial maintenance plan (X+Y) [2] This is covered further in Section 3.

(iii) A third capability would be to provide support for planning operational *schedules* for simultaneous dispatch of lists of prioritized maintenance actions. In this concept, one would go beyond the *analysis* capabilities (i) and (ii) to *synthesis* of recommended schedules for simultaneous activities. Here we envision an "input hopper" of required maintenance activities. Each input to the hopper may have an associated urgency of time priority for its completion. There may also be a designation of which regional work-crew resources will be busied out by the associated action. The aim of the scheduler will be to dispatch the items in its input hopper as quickly as possible, subject to the priorities of each job, limits on the number of simultaneous maintenance actions due to regional resource-use considerations and subject to an ongoing assurance of mesh orthogonality of every multi-action maintenance state. Such a system might be imagined scheduling a large number of work crews on simultaneous maintenance actions continent-wide while always assuring that the collective single-failure risk exposure is no greater than if all the same activities were done one-at-a time in succession. We do not address this idea further in the present work.

B. Models for restorability-related effects of span maintenance

We can see at least three functional models for the effect of maintenance-related activities from a network restorability viewpoint. In order of increasing network influence impact they are:

Type 1: (Ring-like roll-to-protection): In this model all of the spare capacity of the maintenance-span is withdrawn from the remainder of the network with no other effects. This assumes that a working channel block is rolled intact as a modular unit onto an identical sized protection channel block, as in a 4-fiber BLSR, to facilitate maintenance on the working channels. This refer to as a ring-like roll-to-protection. The main feature is that the "roll" is completely contained within the maintenance span.

Type 2: (Mesh equivalent of roll-to-protection): More generally in a mesh-restorable network, the number of working and spare channels on each span can be separately assigned. In an optimized mesh capacity design at the per-channel capacity level, most spans will have $w_i \geq s_i \geq 0$ but $w_i \leq s_i$ and $s_i = 0$ outcomes are possible on some spans as well. In this case the "roll to protection" model becomes generalized:

(i) If $w_i \leq s_i$:

- The 'move to protection' is completely contained within span *i* itself.
- The network wide effect is a withdrawal of w_i units of spare capacity from span *i* but some spare capacity may remain. i.e., the effective s_i is $s_{i,eff} = s_i - w_i$.

(ii) If $w_i > s_i$:

- The 'move to protection' first uses up all the s_i capacity on span *i*, the maintenance span.
- The remaining working amount $w_{i,eff} = w_i - s_i$ is re-routed over replacement paths through spare capacity on other spans of the network, around the maintenance span.
- The network wide effect is withdrawal of all spare capacity from the maintenance span *plus* a network withdrawal of spare capacity on other spans as required to support the creation of maintenance replacement paths for $w_{i,eff}$

---

[2] The following notation is suggested for later use in discussing serial or parallel maintenance actions: (X||Y) –> X and Y occur simultaneously on the network, (X+Y) –> X and Y are scheduled in succession, (X ⊥ Y) –> X and Y are "mesh-orthogonal" maintenance actions, |(X||Y|)| = k –> there are k spans that are in the risk field of *both* X and Y. Given this, |(X||Y|)| = 0 could be thought of as the definition of this type of mesh maintenance orthogonality.

Type 3: (Equivalent to failure & restoration): The extremely worst-case model for the maintenance action is to assume the complete withdrawal of all spare capacity on the maintenance span *and* the complete re-routing of all working demands crossing the span via a restoration-like set of replacement paths using spare capacity on other spans. We do not know if this model is actually required in practice as we understand most maintenance or upgrading is done under roll-and-cut type of procedures. Presumably any type 3 maintenance action would be very carefully considered, scheduled late at night, and / or avoided if at all possible by other strategies because (in *either* a ring- or mesh-based network) such actions put the network in an initial sate that is equivalent to already having sustained and restored a single failure. The network is then completely reliant on any inherent dual-failure restorability should a true failure arise while in this state. For any given demand in a ring, the restorability risk in such a circumstance is already known: it is *all or nothing*: either 100% or 0% for each other span cut. For all demands that cross a failure span while there is a Type 3 maintenance action elsewhere on the ring the risk is 100%. We will characterize the corresponding risk for mesh networks under this worst case model as well as the more benign Type 2 model.

C. Optimizing mesh-networks for Type 2 maintenance

As a bit of an aside, this section shows a simple sense in which a mesh network capacity design can be optimized with Type 2 maintenance considerations in mind. The basic idea is to note that under type 2 maintenance activities, capacity designs that tend towards $w_i \cong s_i$ on every span (rather than either $w_i << s_i$ or $w_i >> s_i$) would be preferred over otherwise equivalent-cost solutions that have no such propensity towards numerical balance of working and spare on each span. In strictly optimal solutions to the mesh-network design problem [3][4] one fairly often sees spans with $s_i = 0$ and others with $s_i > w_i$ but it is evident that many cost-equivalent (or nearly so) other arrangements also exist with better balance in the per-span working / spare ratios. This, in and of itself, suggests that the recent bi-criterion approach [4] could again be applied to maximize the trade-off between the locality of maintenance-related single-failure exposure and the added amount of spare capacity. i.e., the objective function would be[3]:

$$\min\left\{\sum_{i\in S} c_i s_i + \alpha \cdot \sum_{i\in S} \max(0, w_i - s_i)\right\} \qquad (1)$$

When α is arbitrarily high this would result in mesh-restorable networks with complete containment of working capacity on same-span spare capacity for Type 2 maintenance albeit with an increase in spare capacity. At α = 0 we have a conventional mesh spare capacity design. Thus, at one extreme one could generate 100% "maintenance-localized" mesh networks. Varying α through more moderate values would produce a family of designs that embody all the feasible trade-off points between this form of maintenance locality and total spare capacity cost. As found in [4] for restoration path tightening with a similar bi-criterion method, a significant improvement in the average maintenance-locality may arise before any total capacity penalty is required by using a small but non-zero α which effectively instructs the solver to select amongst otherwise capacity-equivalent designs for enhanced spare to working balance on spans.

II. The Theoretical Risk Field of a Mesh Span-maintenance Action

A. Transforming Type 2 maintenance actions to equivalent Type 3 for simulation

Henceforth, we will consider Type 2 and Type 3 logical models for the capacity-related effects of mesh span maintenance. Type 3 maintenance is functionally equivalent to complete failure and restoration of the maintenance span in terms of the network environment seen for restoration of any actual failure that arises while in this state. The Type2 maintenance model can be thought of as a special case of Type 3 maintenance where the actual w_i of the span is first transformed to $w_{i,eff} = \max(0, w_i - s_i)$ and $s_{i,eff} = \max(0, s_i - w_i)$.

[3] As stated this is a nonlinear objective function (because of the max( ) operator). This can be overcome in practice by defining an additional set of w_i-s_i "difference variables" in the formulation and asserting positive-only values.

Type 2 and Type 3 maintenance effects can therefore be conveniently approached though re-use of basic capability from recent studies of mesh network availability involving dual-failure restorability considerations [5]. In this view maintenance on a span is equivalent to voluntarily putting the span in a failed mode and "restoring" all its working channels.[4] For Type 3 the true w_i and s_i are used. For Type 2 modelling $w_{i,eff}$ and $s_{i,eff}$ are used instead.

B. How maintenance creates a theoretical "risk field"

For both maintenance re-routing of working channels and for actual failures we assume a span-restoration mechanism that produces path-sets equivalent to the k-shortest paths [6][7] through available spare capacity between the end-nodes of the failure or maintenance span involved. The theoretical risk field from maintenance on one span *m*, can therefore be evaluated by means similar to computing the dual-failure restorability of the network over all dual-failure scenarios involving span *m* as one of the "failures."

There are two effects through which a maintenance state coupled with an actual span failure can lead to an outcome that is not fully restorable:

1) Contention for spare capacity: After the first span is under maintenance the required number of restoration paths for restoration of a second (failure) span may not be feasible. This could be due to the fact that not enough spare capacity is left after the deployment of maintenance replacement paths or because a given restoration mechanism is not adaptive enough to changes in the spare capacity layer to find feasible paths.

2) Failure on another span used for the maintenance replacement paths: In this situation a failure occurs on a span that is currently supporting one or more of the paths used for maintenance replacement. The outcome of such a situation depends on the amount and distribution of remaining spare capacity in the network and may result in un-restorable fractions on both the maintenance span and failure span.

Under the equivalence of a Type 3 maintenance model to a restored span failure, and the simple transformation of a Type 2 maintenance action to an equivalent Type 3 action of lesser impact, we can go ahead and speak about maintenance state and failure state combinations as if it was two actual failures being considered. In a dual span-failure case the previously defined [5] dual-failure restorability $R_2(a,b)$ is:

$$R_2(a,b) = \frac{(w_a - N_a) + (w_b - N_b)}{w_a + w_b} \tag{2}$$

where N_a and N_b are the number of channels that are *not* restorable on span *a* and span *b* respectively and w_a , w_b are the number of working channels on span *a* and *b* respectively. To apply this to maintenance, we can define the *theoretical loss of restorability on a span i from maintenance on a given span m*, in terms of the percentage of all working channels on span *i plus span m* that would be non-restorable if a failure occurred on *i* in the presence of the maintenance state on span *m*.

$$L_m(i) \triangleq \frac{N_m + N_i}{w_m + w_i} = 1 - R_2(m,i) \tag{3}$$

For each span taken as a maintenance span *m*, one can therefore define a vector of risk $\vec{L}_m$ whose components are the $L_m(i)$ for each other span *i*. The theoretical *risk field* caused by a maintenance action on

[4] Regarding the term "channel": The domain of this problem is really one of logical architecture and is not specifically tied to any one technology such as DWDM or Sonet. The term *channel* can thus represent a wavelength, a waveband, an OC-n, STS3, a whole fiber, and so on - whatever is the basic unit of capacity management at the given transport level. The methods can be extended to accommodate multiple demand sizes and channel sizes but here we assume a single unit-capacity channel entity and demand quantities that are in the same units.

span *m* is then the set of all other spans *i* for which the theoretical loss of restorability $L_m(i)$ is non-zero. (We call it theoretical to emphasize that there is no actual outage unless a failure really occurs on span *i* during the maintenance state). Within the field of risk of a maintenance action working channels on other spans are exposed to a *possible* outage due to single span-failures, whereas in a maintenance-free state every span of the mesh-restorable network is guaranteed 100% restorability to any single failure, by design. Prospective failures on spans outside of the risk field retain their guarantee of full single-failure restorability. Once the maintenance risk fields of each span are quantified an obvious use of them is to co-ordinate actions so that maintenance on a span *i* which is in the risk field of a current maintenance *m*, is deferred until *m* is complete.

C. Method for evaluating the theoretical risk fields due to span maintenance:

In this section we outline a computational procedure to define the theoretical field of risk of any given span maintenance. The procedure computes risk fields with the following properties:

(i) Type 2 or Type 3 maintenance actions are modelled. Span maintenance effectively deletes some or all spare channels on the span and may force some or all working channels onto restoration paths depending on the maintenance model and the working and spare capacities.

(ii) We assume that the restoration mechanism is intelligent and adaptive to the extent that restoration of a span *i*, while span *m* is under maintenance, recognizes and adapts to any preceding consumption of spare capacity on the maintenance span and / or the network by the maintenance state.

(iii) Effects of failure on both the failure and maintenance span are considered: If a failure of span *i* severs one or more maintenance replacement paths for span *m*, a restoration action is triggered for span *i* followed by a restoration effort for any lost working channels from span *m*. The restoration mechanism will try to find new restoration paths between the end nodes of the maintenance span over spare channels not used by the restoration action for the failure itself. This includes a release of surviving spare capacity on maintenance-replacement paths for the span *m* and use of the restoration process in the context of span *m* to find new maintenance replacement paths. If restoration of the failure span *i* was incomplete, any further restoration paths for span *i* that may be feasible after dissolution of the failed replacement paths from span *m*, are then sought, completing the restoration reaction to the failure.

Figure 1 details the overall algorithm for calculating the risk field. The procedure in (iii) is a functional approximation for the effects of a distributed restoration protocol that is self-updating, an example of which is in [2]. Such a protocol immediately gives “working” status to any spare channel it uses in a restoration path or as a maintenance replacement path. This inherently updates itself should it be triggered to act again to protect either new capacity or maintenance replacement paths in the event of a failure during maintenance. Note that in the event of failure on an already deployed restoration (or maintenance replacement) path it is preferable to fall back to the end nodes of the original failure (or maintenance) span, rather than secondarily restore these paths from the new failure span.

Figure 2 shows an example of a calculated risk field for maintenance of the indicated span “M”. The test network of 22 nodes and 41 spans was manually created as an arbitrary test case. The test case is based on a gravity demand model with an average of 3.2 units of demand per O-D pair. Demands are routed via shortest-path routing followed by a Herzberg-type [3] optimal spare capacity design with a hop limit of five. The capacity design is at the theoretical minimum of spare capacity to still be restorable. Real networks will rarely be this minimally capacitated and will have smaller risk fields as a consequence. In Fig.2(a) and (b) the identical unit-channel integer capacity design applies but the maintenance model differs. In Fig.2(c) the corresponding minimum-capacity-cost *modular* capacity design is used, obtained with the method of [8] with a hop limit of 5 and module sizes 12, 24, 48. The non-modular capacity design ((a) and (b)) has an average of 34.3 working and 18.4 spare channels per span and 53.5 % redundancy. Fig.2(c) has a redundancy of 73% due to modularity effects. The numbers on each non-maintenance span are the entries of the $\vec{L}_m$ risk vector portrayed as the percentage of the working channels on those spans plus the maintenance span that would not be restorable *if* the indicated span fails during the maintenance state for span M. The maintenance span chosen for the illustration is a typical span for this network in that it has 32 working channels and 15 spare channels. The computation of risk fields for all spans based on simulation of all combinations of maintenance span and failure span combinations for Figure 2 takes only a few seconds.

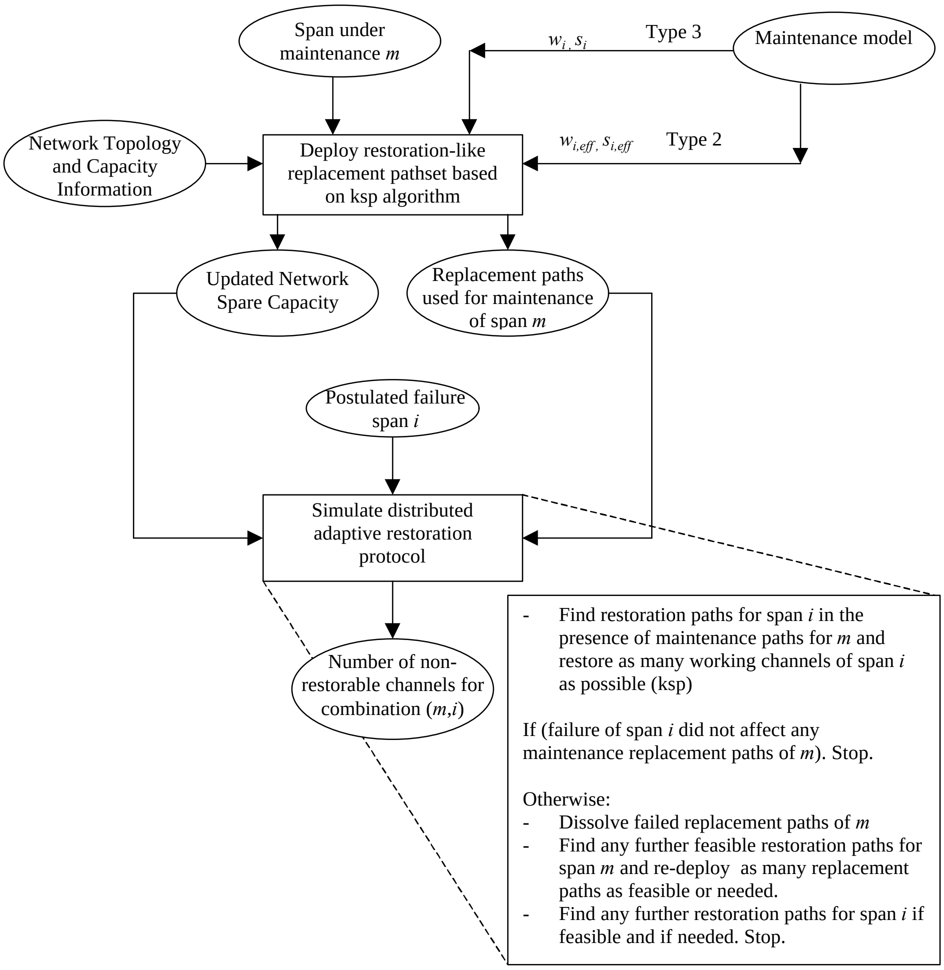

Figure 1. Algorithm for calculating theoretical risk fields

Fig.2(a) shows the risk field for Type 3 (worst-case) maintenance on span M. Fig. 2(b) shows the corresponding Type 2 maintenance result in the same network. Under a Type 2 maintenance model for this span 15 of the 32 working channels are rolled to protection on the same span and the remaining 17 working channels are re-routed via replacement paths through the network in the vicinity of M. Under Type 3 all 32 working channels are re-routed. The extent and magnitude of the risk field behaves as one might expect from first principles. In Fig. 2(a), the risk field contains 14 spans and a total of 292 working channels are non-restorable *if* they should fail (in a complete span cut)[5] during span M maintenance. Under Type 2 maintenance the extent of the field drops to 8 spans and the total magnitude to 157. If the inevitable margin of extra capacity present in a modular design is available for re-routing, Fig. 2(c) shows a further contraction of the risk field to 7 spans and a magnitude total of 129 channels.

[5] Any single-channel failure would remain fully restorable.

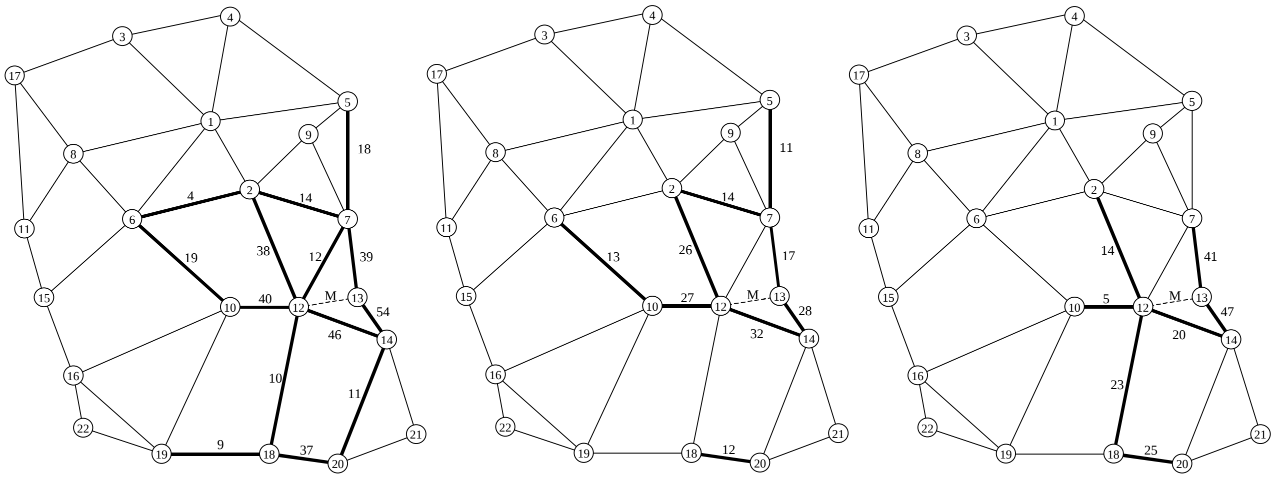

(a) Non-modular, Type 3, 14 spans, 292 ch's (b) Non-modular, Type 2, 9 spans, 157 ch's (c) Non-modular, Type 2, 7 spans, 129 ch's

Figure 2. Risk fields (in percent loss of restorability) for maintenance on span M in the network shown.

III. Experimental Effects of Modularity, Hop-limit and Bi-criterion Design

A. Test Networks

A series of experimental trials was designed to learn about the typical extent of risk fields due to maintenance in span-restorable mesh networks having a range of relevant design properties. Three basic topologies were used 11n21s1, 15n28s1, 19n37s1, the names indicating the number of nodes and spans in the topology. The 19n37s1 topology was obtained from [9]. The 15n28s1 is the well-studied "Bellcore" network model from [10] and 11n21s1 is a random transport graph selected from our library of test case models. All graphs were accompanied with randomly generated gravity-type demand patterns using nodal degree as a surrogate for the relative attraction effect of each node. For each basic topology and demand matrix five types of capacity design were produced for testing risk field properties. Three of them are "Joint Capacity Placement (JCP)" solutions. These are designs where the working path routing and spare capacity placement are jointly optimized under the formulation given in [8]. They are denoted JCP "H" where H is the hop-limit of the corresponding design[6]. The three (or in some cases four) JCP designs are complemented with a non-joint but modular capacity design at a hop limit of five, imported from the work in [8]. The modular capacity designs use the same demand routing from the corresponding JCP-5 non-modular designs but min-cost modular capacity decisions are made such that the working and spare total placement on each span is modular. Capacity modularities were 12, 24, and 48 channels. The relevance of the modular designs is that modularity should tend to reduce the extent of the risk fields. Finally, a corresponding design for each network was imported from a recent study of a bi-criterion optimization approach that allows a trade-off between total capacity and average restoration path lengths [4]. The bi-criterion networks tested here were those at which the greatest tightening of restoration paths has occurred before any increase in capacity investment is made in the method of [4] and are non-modular. The bi-criterion designs are also relevant to this study of risk fields as the "tightness" of the restoration paths should also affect risk fields.

B. Method and Results

Each span in each test network was tested for its effect as a maintenance span under Type 2 and Type 3 models for the maintenance process. For each span, in the context of a maintenance span, *m*, we compute the theoretical loss of restorability $L_m(i)$ for each other span *i* as a prospective failure span during the maintenance state. This is done for both maintenance models. We summarize the results for each test case and each maintenance model by computing the average total theoretical risk $\sum L_m(i)$ for each prospective

[6] The hop-limit as applied to working path routing in a joint design is an allowance for *additional routing hops* above the shortest path on the graph. For restoration paths it is a direct limit on the length of any restoration route. See [8].

maintenance scenario. These data are given as the absolute average number of channels in total that undergo a risk of restorability loss due to each maintenance action. For illustration $\sum L_m(i) = 100$ would be equivalent to a risk field of 10 spans in extent each exposed to a 10-channel restorability path shortage if a failure occurred on the span while in a maintenance state. The absolute values are not of primary significance since this depends on the total demand of the test network design, but it is not clear which case could be justified as a normalizer, so we present the absolute data for comparative inspection. A separate analysis of the span by span fractional restorability loss percentages (Figure 3) addresses the question of how significant a risk exposure these absolute totals represent. The second summary characteristic of each trial in Table 1 is the average logical extent of the risk field, $S(m)$. This is the average number of spans undergoing any non-zero risk over all maintenance span scenarios. The results are summarized in this way in Table 1.

Table 1. Average risk field extent $S(m)$ (spans) and total risk exposure $\sum L_m(i)$ (channels) in trials

Topology	Maint. Model	Capacity Design JCP 4		JCP 5		JCP 6		JCP 7		Modular		Bi-Criteria	
		$\sum L_m(i)$	$S(m)$	$\sum L_m(i)$	$S(m)$	$\sum L_m(i)$	$S(m)$	$\sum L_m(i)$	$S(m)$	$\sum L_m(i)$	$S(m)$	$\sum L_m(i)$	$S(m)$
11n21s1	Type 3	247.7	14.2	294.7	16	330.1	16.7	n/a		152.1	9.2	294.1	16
	Type 2	128.4	10.2	169	12.4	202.8	13.5			57	5.2	169.1	12.5
15n28s1	Type 3	352.4	14.5	401.8	15.7	406.2	16	490.4	18.3	263.2	11.7	401.1	16
	Type 2	172.5	9.3	209.1	10.6	215.8	11	274.2	13.5	99.7	6	208.9	10.6
19n37s1	Type 3	530.6	18	577.2	19.5	713	23.4	535.8	20	398.5	15.1	586.9	19.8
	Type 2	297.3	12.8	323.3	13.5	420.7	16.9	276.2	13.5	181.2	8.9	330.5	13.9

While Table 1 gives the extent and absolute magnitude averages for each risk field, Figure 3 shows the results from a converse standpoint which is the distribution of risk exposure on non-maintenance spans in the network due to all possible single-span maintenance actions in its network. The data in Figure 3 pools individual trial $L_m(i)$ values for the Modular H=5 designs over all networks for each maintenance type. The three test networks provide 21+28+37 = 86 instances of maintenance-spans to test for the risk imposed on each other span in the test case.

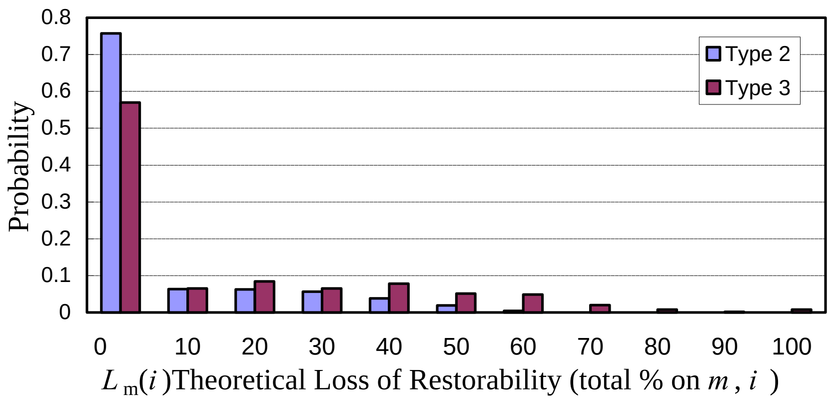

Figure 3. Histogram of individual span restorability risks over all Type 2 and Type 3 maintenance actions in the Modular H=5 network designs.

The total population of $L_m(i)$ risk exposure values for Fig. 3 is therefore 21(20) + 28(27)+ 37(36) = 2508 instances. The Modular designs are chosen for this exhibit of results because they are the most realistic family of capacity designs. To aid in interpretation of Fig. 3 it is a histogram which shows the probability of experiencing a restorability risk of the percentage magnitude shown in each x-axis category. As produced the results inherently assume all spans are equally likely maintenance candidates. Reading one point for illustration Figure 3 says, for instance, that *under Type 2 maintenance there is 75% chance that any given other span experiences zero risk*. Or, reading from, say, the sixth x-axis category as a further example, we would say that *under Type 3 maintenance there is 5% chance of any given other span being exposed to a risk of 30-40% loss of restorability.*

C. Discussion of results

First, and not surprisingly, Table 1 shows that in all trial networks the Type 3 maintenance model engenders much greater risk of restorability loss than with Type 2. Secondly, in comparing both risk field extent and total risk exposure as the design hop limit decreases from H = 6 down to H = 4 we see the expected contraction of extent and risk magnitude, due to changes in the design that increasingly keep restoration flows (and maintenance replacement paths) closer to home. Only the results for 19n37s1 JCP 7 break the trend and are more difficult to explain in that we see a lower risk field extent and magnitudes than JCP 6 for that network. Being jointly optimized capacity designs, we surmise that this must be a network-specific effect, i.e., that in that particular topology at a hop limit of seven some significant shift in the basic solution structure is enabled which is enough to overcome the otherwise general trend in $S(m)$ and $\sum L_m(i)$ as the design H increases. This is being looked into further as of the deadline for this paper.

Some important practical observations can be made from Fig. 3. First, under 2508 instances of spans at potential risk due to 86 instances of Type 2 span maintenance, in three different networks, there *was not one case of 100% risk exposure* as in rings. In fact, there was only a 25% chance of any risk and the risk exposure was less than 50% of the number of working channels on the span in 99.6% of the trial cases. Moreover under Type 2 no spans ever experience risk exposures over 70%. Even under Type 3 maintenance, 91.4% of the time the risk imposed on other spans is <50%. Clearly this means that *priority signal paths could be given an assurance, not possible in rings, of zero risk of restorability loss due to maintenance*. As expected there are more instances of higher-risk situations under Type 3 maintenance, including a small percentage of 100% restorability risks. The results contain exactly 20 instances of $L_m(i) = 1$ under Type 3 but every one of these is easily explained by its correspondence to a degree-2 node in the respective test network. Under Type 3 maintenance to one span at a degree-2 node, the opposite span at the node undergoes a 100% restorability risk, as in a ring, purely from topological considerations. However, this is only true for the mesh under Type 3 maintenance. Under Type 2 maintenance there are no instances of $L_m(i) = 1$. In contrast $L_m(i) = 1$ for either Type 2 or 3 maintenance models in a ring.

IV. Considerations for simultaneous maintenance actions

In this section we give some preliminary considerations about risk-related issues of conducting more than one maintenance action at a time on different parts of the network. A central concern is on ways to manage multiple maintenance actions so that the total risk exposure is not compounded to be more than the sum of the individual maintenance actions risk magnitudes. A simple solution of the span maintenance scheduling problem would be to schedule each maintenance one after the other so as to minimize the consequences of possible true span failures happening while these maintenance actions are being conducted but in a large continental or even metro network workforce efficiency and total maintenance intensity might both require that simultaneous maintenance actions be routinely performed.

A. The notion of maintenance orthogonality

Obviously if two maintenance actions generate completely disjoint risk fields then it should be possible to conduct them at the same time without running a higher risk of outage than if they are conducted one after the other. For example, if a span maintenance conducted in the north-west part of the city affects only spans in that part of the network and a span maintenance conducted in the south-east part of the city affects only the spans in that part of the network then doing the maintenance on the span in the south-east while the span in the north-west is also under maintenance will not aggravate the risk in the north-west region and vice-versa. But more generally the risk fields may not be disjoint. However, in any case where the total risk of the combined action is not greater than sum of the risks of the same actions in series then we say the two maintenance projects are mesh orthogonal. A way to define this approximate notion of orthogonality[7] is:

$$\text{Condition for mesh } \textit{maintenance orthogonality}: \sum_{i\in S} L_N(i)\cdot\Delta T_N = \sum_{k=1..N}\sum_{i\in S} L_k(i)\cdot d\,t_k \qquad (4)$$

where:

$L_N(i) =$ number of channels at risk on span i under N simultaneous maintenance actions

$\Delta T_N =$ elapsed time in composite maintenance scenario

$L_k(i) =$ number of channels at risk on span i under individual maintenance scenario k.

$d\,t_k =$ elapsed time in individual maintenance scenario k.

Conversely the difference between terms in Eq. (4) is a measure of the compounded risk of doing the actions in parallel. To compute the test of orthogonality as defined in Eq. (4) we need a method to compute the left-hand side, i.e., the composite risk of the simultaneous maintenance states. Note the role that elapsed time can play as well as the more subtle topological effects that are involved. If the parallel actions can be done in less time, compound risk may be effectively avoided. Obviously it is preferable to conduct parallel operations that are truly topologically orthogonal, however. Obvious extensions to the procedure in Figure 1 can be used to compute $L_{m_1,m_2...N}(i)$ values by conducting the test of span i restorability in the presence the set of N specific sets of maintenance replacement paths and span spare capacity withdrawals.

B. Managing avoidable risk for priority paths

Unless we were willing to pay the quite significant price for a complete dual-failure restorable capacity design [5] we have seen above that a certain risk due to maintenance is unavoidable. But what about trying to do so just for a subset of premium service customers? Let us here consider that in a network where much maintenance activity is required, it may be unavoidable that multiple simultaneous actions have to be undertaken without any assurance of orthogonality in the above sense. In such circumstances a reasonable viewpoint might be to attempt to at least ensure the avoidance of compound risk for priority paths. Thus, the idea here is to consider a selective measure of the total risk for a class of priority paths crossing each span. This would lead to a corresponding definition of maintenance orthogonality with respect to priority paths. Ensuring orthogonality of maintenance risks for a subset of premium services may be much more feasible than trying to achieve it for all working capacity.

A proposed measure of priority-services total risk caused by a maintenance action is the following:

$$\text{Priority Risk}(m) = \sum_{i\in S}\left(\max(0, P_m(i) - w_p(i)) + \max(0, P_m^*(i) - w_p(m)\right) \qquad (5)$$

where $P_m(i)$ is the number of paths found for restoration of channels on span i when span m is in the maintenance state, $P_m^*(i)$ is the number of paths found for replacement of channels on span m when it is under maintenance and span i fails and $w_p(k)$ is the number non-priority service paths crossing any span k.

Based on the above definition of the priority risk, and similarly to what has been introduced for the general definition of the risk, a total "serial risk" $\text{Risk}^+(m_1, m_2, ..., N)$ can be defined for multiple maintenance

---

[7] Here "orthogonality" is meant in an evocative sense. Mathematical orthogonality is tempting to claim for a pair of risk vectors as it would require only that $\vec{L}_{m_1}\cdot\vec{L}_{m_2} = 0$ which would be true any time that either vector has a zero in each position. However, careful reasoning about the risk field vectors will show that L vectors can change in the presence of each other in ways dependent on the topology. If one span, *m1*, is in maintenance in the presence of another span *m2*, the risk on another span i from span *m1*, is not necessarily a constant under *m2*. Constructions can be made to show examples where paths for *m2* could isolate span *i* from other capacity needed to form its restoration paths, even though span *i* is not in its risk field. Hence, we have to rely only on the more general condition of Eq. (4) as the *definition* of maintenance orthogonality and cannot give a strictly general stipulation of a sufficient condition.

actions when done individually in series as well as a total "parallel risk" $\text{Risk}^{\parallel}(m_1, m_2, ..., N)$ for multiple maintenance actions done in parallel. The difference in these two computed risk measures is a more specialized assessment of whether a proposed set of simultaneous maintenance actions will deleteriously expose priority paths to compound risk.

As seen in the previous sections, the risk values calculated for Type 2 are almost always below 50% and always below 70%. Consequently, if priority channels never represent more than 30% of the channels on any span then the priority risk for *single* maintenance actions is already null. The priority risk for parallel maintenance actions, however, might be non-zero for some maintenance combinations. Calculation of the total parallel risk $\text{Risk}^{\parallel}(m_1, m_2, ..., N)$ will identify such situation for avoidance.

V. Conclusion

This work has developed and applied methods for analysis of the risk to restorability engendered by maintenance actions of spans of a mesh-restorable network. It is now clear that while rings have a contained risk field of (S-1) other spans that are at 100% restorability risk, mesh networks exhibit a more extended risk field but with much lower risk of restorability loss on any one span. From the fact that over 90% of spans have restorability risks under 50% for either Type 2 or 3 maintenance models, it seems valid to conclude that mesh networks would be able to support a large fraction of priority services which can be *guaranteed zero risk to their normal restorability* due to the operators maintenance activities. The assumptions are only that maintenance actions are done one span at a time and that priority status is taken into account in restoration.

Acknowledgements

The authors thank Ofer Shai for developing the risk field calculation programs and John Doucette for his help in reviewing this paper. Both are with TR*Labs* Network Systems and University of Alberta.

References

[1] W.D. Grover, “Distributed Restoration of the Transport Network”, *in Network Management into the 21st Century*, editors T. Plevyak, S. Aidarous, IEEE / IEE Press, Chapter 11, Feb. 1994, pp. 337-417.

[2] W. D. Grover, “Self-organizing Broadband Transport Networks”, Proc. of the IEEE Special Issue on Communications in the 21st Century, vol. 85, no.10, pp. 1582-1611, October 1997.

[3] Herzberg and S. Bye, “An optimal spare-capacity assignment model for survivable networks with hop limits,” in *Proc. IEEE Globecom’94*, 1994, pp.1601-1607.

[4] J. Doucette, W. D. Grover, T. Bach, "Bi-criteria studies of mesh network restoration path-length versus capacity tradeoffs,” Proc. *OFC 2001*, Anaheim, Calif., March 17-22, 2001.

[5] M. Clouqueur and W. D. Grover, “Computational and design studies on the unavailability of mesh-restorable networks,” *Proc. DRCN 2000*, April 2000, Munich, Germany, pp. 181-186.

[6] D. A. Dunn, W. D. Grover, and M.H. MacGregor, “A comparison of k-shortest paths and maximum flow methods for network facility restoration”, *IEEE JSAC*, vol. 12, no. 1, pp. 88-99, Jan. 94.

[7] M. H. MacGregor and W. D. Grover, “Optimized k-shortest paths algorithm for facility restoration”, *Software-Practice & Experience*, Vol. 24(9), Wiley & Sons, pp.823-834, 1994.

[8] 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.

[9] A. Lardies, A. Aguilar, "Planning Methodology for SDH + Optical Networks*", First Int. Workshop on the Design of Reliable Comm. Networks, DRCN '98*, Belgium, 1998.

[10] *The Role of Digital Cross-connect Systems in Transport Network Survivability*, SR-NWT-002514, Issue 1, Bellcore Special Report, Jan. 1993.