Availability Analysis of Span-Restorable Mesh Networks

Matthieu Clouqueur, *Student Member, IEEE* and Wayne D. Grover, *Fellow, IEEE*

***Abstract*—The most common aim in designing a survivable network is to achieve restorability against all single span failures, with a minimal investment in spare capacity. This leaves dual-failure situations as the main factor to consider in quantifying how the availability of services benefit from the investment in restorability. We approach the question in part with a theoretical framework and in part with a series of computational routing trials. The computational part of the analysis includes all details of graph topology, capacity distribution, and the details of the restoration process, effects that were generally subject to significant approximations in prior work. The main finding is that a span-restorable mesh network can be extremely robust under dual-failure events against which they are not specifically designed. In a modular-capacity environment, an adaptive restoration process was found to restore as much as 95% of failed capacity on average over all dual-failure scenarios, even though the network was designed with minimal spare capacity to assure only single-failure restorability. The results also imply that for a priority service class, mesh networks could provide *even higher availability than dedicated* 1+1 APS. This is because there are almost no dual-failure scenarios for which some partial restoration level is not possible, whereas with 1+1 APS (or rings) there are an assured number of dual-failure scenarios for which the path restorability is zero. Results suggest conservatively that 20% or more of the paths in a mesh network could enjoy this ultra-high availability service by assigning fractional recovery capacity preferentially to those paths upon a dual failure scenario.**

***Index Terms*—Availability, mesh networks, network fault tolerance, network reliability, optical transport networks, protection and restoration, reconfiguration.**

I. Introduction

THE COMMON AIM of ring and mesh-based restoration or protection techniques is to provide real-time recovery of carrier signals against any single span failure. Such networks are said to be "fully restorable," but the term really refers only to single-failure scenarios. While methods for the capacity design of survivable networks have developed greatly in the last decade, the related problem of how this affects availability is of growing interest. A basic question is: "What does the investment in spare capacity for restorability mean *quantitatively* for the availability of service paths?" Clearly single-failure restorability is of great overall benefit to network integrity, but that qualitative recognition does not directly help a service provider know what service level agreements (SLA) can be safely offered. For ring-based networks, some recent work has provided relatively thorough and exact mathematical models for quantification of path availability [1]–[3], but for mesh-restorable networks the problem of availability analysis, and the relationship between availability and network capacity is not as well understood.

A. Restorability, Reliability, and Availability

There are several ways in which the "reliability" of a network or service can be measured. Depending on the service, the relevant measures may be availability, throughput, delay, probability of graph disconnection, dial-tone delay, service establishment times, cell-loss rate, error-rate, and so on. To users of a wavelength-managed optical transport network, the end-to-end availability of lightpaths is of paramount concern. This is the orientation that we take toward the general notion of reliability. In common language,[1] reliability is a qualitative perception of the network being predictable and dependable. Our focus, however, is on the specific concepts of *restorability* and *availability* and how they are interrelated in the capacity design of mesh-restorable networks.

We define *restorability* of a network as a whole as the average fraction of failed working span capacity that can be restored by a specified mechanism within the spare capacity that is provided in a network. The average is taken over some previously stipulated set of failure scenarios, most often the set of all complete single-span failures. The restorability of a specific failure scenario has a corresponding definition. For each single failure scenario i, some w_i units of working capacity, such as wavelength-links or SONET OC-n carrier signals, are disrupted by the failure. The fraction of w_i that is subsequently restored is the restorability $R_1(i)$ of the individual failure scenario. The network-wide ratio of restorable capacity to failed capacity over all single-failure scenarios is called the single-failure network restorability, R_1.

Our interest is in the availability of service paths in networks which are efficiently capacitated for $R_1 = 1$ but then subjected to double-failure scenarios. The restoration mechanism is "span restoration," to be described further. Note that restorability is a logical design property of a network. It is not in itself a probabilistic measure, so some further considerations will have to come into play to relate restorability to availability. Because restorable networks may operate at an STS, wavelength, wave-band or even fiber-managed layer, we will refer generically to the basic unit of capacity that is manipulated for restoration as a *link*, and the set of all working and spare links between

Manuscript received March 8, 2001; revised December 20, 2001.

The authors are with TR*Labs*, Edmonton, Alberta, Canada T5K 2P7 and also with the Department of Electrical and Computer Engineering, University of Alberta, Edmonton, Alberta (e-mail: mclouque@edm.trlabs.ca; grover@edm.trlabs.ca).

Publisher Item Identifier S 0733-8716(02)03060-3.

[1]As opposed to the specific technical meaning of the term *reliability* in the fields of reliability engineering [4], [5] and *network reliability* in computational graph theory [6]–[8].

adjacent cross-connecting nodes as a *span*.[2] A *path* is a specific concatenation of two or more links on a route over spans to bear a payload signal between its origin–destination (O–D) nodes.

Availability is the probability that a system (in this case a signal path) will be found in the operating state at a random time in the future. Availability inherently reflects a statistical equilibrium between failure processes and repair processes in maintained repairable systems that are returned to the operating state following any failure. In the technical sense of the word, *reliability* is not the same as availability. Reliability is the probability that a system or component will operate without any service-affecting failure for a period of time t. Reliability is a monotonically decreasing probability function of time, $R(t)$ [4], [5], and a specific reliability number always implies an assumed duration of time. Reliability is of concern to known how soon the next repair expenses might be incurred, etc., but reliability itself does not consider the repeated cycles of failure, repair time, and return to service which determine the availability of an ongoing service. Reliability is a more mission-oriented measure: for example the likelihood that a component will operate without any failure through the duration of a rocket boost phase is a reliability measure. In contrast availability asks: given the frequency of failures and the rate at which repairs are conducted, what is the average fraction of time that one will find the system in the operating state?

B. Span-Restorable Mesh Networks

In span restoration,[3] the rerouting for survivability occurs between the immediate end nodes of the break. This need not be via a single route, nor via only simple two-hop routes. The general idea of span restoration is illustrated in Fig. 1. The design methods incorporate a hop or distance limit, or (if routes are characterized by their optical pathloss) by a loss limit. Setting the hop or distance limit H allows a tradeoff between the maximum length of restoration paths and the total investment in spare capacity. As H is increased, more complex and finely resolved patterns of rerouting are permitted, resulting in greater sharing of spare capacity. At a threshold value of H, the theoretical minimum of spare capacity is reached [11]. Real-time mechanisms for distributed adaptive span restoration were initially developed in the Sonet era [12]–[16] and are being considered for adaptation to WDM transport applications. The optimal spare capacity design of span-restorable mesh networks is also well understood today [11], [17]–[19].

[2]The terms *span* and *link*—as used here—have their origins in the transmission networking community. As Bhandari [9] explains, the point is to distinguish between the logical links of higher service-layers, in this case the logical network of lightpath connectivity links and the physical transmission "spans" over which all end to end logical links are established. "… Spans are the set of physical transmission fibers/cables in the physical facility graph. Links (or edges) of the logical connectivity graph are built from spans. A given span can, thus, be common to a number of links." [9].

[3]It is important to note that historically some authors refer to mesh networks without implying *mesh-restorability* in the sense we consider. The former is a reference to the topology providing more than one route between nodes, instead of a spanning tree. "Restorable mesh" specifically implies an active restoration mechanism and an optimized distribution of spare capacity to support a target level of restorability. Recently, some authors [10] also refer to a form of fiber-level span-restoration as *loop-back recovery*.

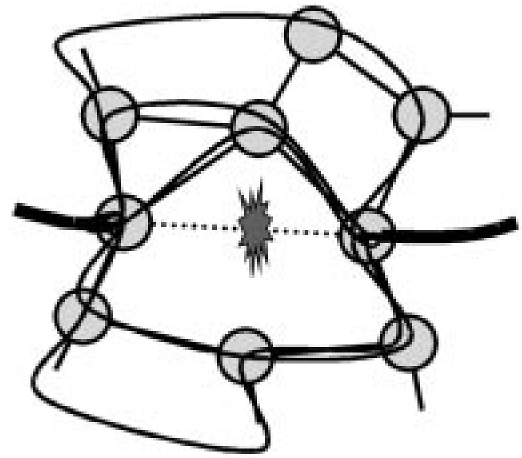

Fig. 1. General concept of span restoration.

C. Outline

Section II reviews prior literature related to the present work. Section III addresses the problem of computing the availability of a path through an actively reconfigured mesh-restorable network. Section IV applies these methods to a series of test networks to determine the availability of paths through a dual-failure analysis of failures in networks that are optimally designed only for $R_1 = 1$. Section V presents and discusses the results. One of the most important findings is that the dual-failure restorability is very high. In concluding, we explain the important implication this finding has: that a self-healing mesh network could provide an ultra-high availability for premium services. Such a service class would enjoy availability superior even to that provided by $1{+}1$ diverse automatic protection switching (APS) by virtue of withstanding almost all dual failures (as well as all single failures).

II. Prior Literature

The analysis of service path availability in ring-based networks has recently been looked at in some depth [1]–[3]. In contrast, the availability analysis of mesh-restorable networks is not as amenable to a detailed analytical approach and has consequently usually been approached with significant simplifying approximations. The problem is that the flexible nature of a mesh-restorable network, its routing adaptability and its extensive sharing of spare capacity, make it far less clear how to directly enumerate the outage-causing failure combinations.

Judging by name only, the closest body of literature to our present concern would appear to be that of "network reliability" [6]–[8]. But this field is concerned with various measures of the graph disconnection probability under the assumptions of edge failure probabilities that are very high compared to our case and with no limitation on the number of simultaneous edge failures. While this is a challenging theoretical problem, the purely topological existence of a route is not a sufficient condition for the availability of paths in a realistic transport network. Path availability also depends intimately on the working and spare capacitation of the network and the precise details of the restoration mechanism that applies. Network reliability in the sense of [6]–[8] makes no consideration of standby redundancy, active restoration schemes, and the repair of physical failures. Our approach is more along the lines advocated by Spragins [20], [21] where the actual transmission, switching and routing structures and specific fault recovery mechanisms are taken into account. The functional model and failure combinations are mechanistically detailed and constrained.

References [22] and [23] both provide studies that in some ways are precursors to the present work. Reference [22] shows the significant improvement brought to path availability from measures such as protection switching on transmission sections, random diversification of trunk routes, and cross-connect route-diversity switching. Rowe's study [22] used simulation of a long random sequence of failure and rerouting actions to predict availability. That work did not, however, address mesh-restorable networks in the sense that we consider them with a fully developed capacity design theory and detailed models for the span restoration rerouting process. Wilson [23] studied the impact of ring and mesh restoration mechanisms on service availability. His "point-to-point" mesh network model is the closest to our work but is limited to a special configuration that also does not reflect the complete diversity, topology dependence, and adaptability of the mesh-restorable networks that we consider.

Arijs *et al.* [24] compare ring and mesh architectures from a cost versus availability point-of-view. Their availability calculations for the mesh network are limited to a dedicated mesh *protection* model because of the "extensive simulations required" with the analysis of mesh *restoration*. This aspect is precisely one of the open issues of mesh availability analysis; how to determine availability in the presence of a dynamic adaptive mesh restoration scheme, as opposed to predetermined protection arrangements.

Cankaya [25] also found it necessary to simplify the mesh restoration model to address the availability analysis with Markov modeling methods. Solving the equilibrium equations gives a network-wide availability based on the assumption that below a certain functioning level the network "as a whole" is consider failed. Specific effects of the network topology, the capacity distribution, and restoration routing behavior are, however, completely lost in the assignment of global state transition probabilities to apply the Markov method.

Barezzani *et al.* [26] study the availability of a network with several traffic priority classes. Network availability is defined as the probability that the proportion of end-to-end connections in the up state at any time is above a certain level for each of several priority classes. The study also investigates the improvements brought to the availability of high priority traffic by allowing lower priority traffic to be dropped for its restoration. The same definition of network availability is used by Vercauteren in [27] in the context of multilayer restoration, in which restoration starts at a given layer of the transport network when the lower levels have exhausted all their restoration capability. Reference [27] recognizes that an availability definition based on a "network fully up" or "network fully down" does not convey the notion of how much traffic is actually lost. That is addressed with an expected loss of traffic metric for the whole network.

Lumetta *et al.* [28] study the impact of dual link-failures in a class of four-fiber "loop-back protected" networks where unprotected traffic is allowed to replace the role of protection capacity on certain links of the backup network. Such "generalized loop-back" networks were introduced in [35]. The basic idea is of 100% redundant networks which are like rings in that every span consists of a bidirectional working fiber pair and a matching pair of backup fibers, as in a four-fiber bidirectional line switched rings (BLSR). Unlike a BLSR, however, protection rerouting can take a generalized route over the equal-capacity backup network, rather than being restricted to following a particular ring structure. While not offering any efficiency improvement over ring-based networking, the removal of the ring-overlay construct is seen as an advantage in terms of flexibility in adding new links and in allowing better recovery levels to multiple failures and node failures than in ring-based networks. Generalized loopback networks are, thus, like BLSRs in all regards except that the ring structuring is lifted, creating in effect bidirectional line switched networks (BLSNs), where N stands for network instead of ring.

In [28], the efficiency of such networks is improved by defining a class of unprotected traffic and a corresponding subgraph of links which carry unprotected traffic on what would otherwise have been the protection fibers of that link. In this role, such links provide no protection capacity to the rest of the network, but also do not access protection on other links upon their failure for the unprotected part of their two traffic components. (This is truly unprotected working capacity, not preemptable extra traffic on protection.) Other links are comprised of the same four-fiber ring-like combination of protected working fibers and pure protection fibers. The study [28] characterizes how the percentage of all two-link failures from which the network can completely recover depends on the number of links put into the role of carrying unprotected traffic instead of providing protection.

In relation to the present work, it is implicit that any amount of unprotected traffic in the sense of [28] is always possible in the networks we consider. One can simply designate *any individual service path* as unprotected, without requiring that it be routable over any specific subset of fiber links having that attribute. Any number of unprotected paths may be present. They are simply excluded from access to spare capacity upon restoration, although they can also easily be given a best-efforts restoration treatment if desired, but in all cases are excluded from consideration in the spare capacity design (except, in a modular capacity design, for their straight-forward inclusion in the final span capacity totals). A somewhat related aspect in the present work is that it is implicit that all spare capacity can be used for preemptable "extra traffic" services, as it could also have been in [28], providing support for an even a greater volume of nonprotected traffic services. We make no further mention of either of these except to stress the ease of support for either unprotected or preemptable traffic in the class of networks we consider. Our network test cases, therefore, need to be understood as considering only the protected-traffic problem, it being implicit that there can be any additional amount of unprotected traffic that simply does not enter the restorability design problem. In addition, there can be preemptable traffic on spare capacity at any time in the networks we consider, it being essentially invisible to the consideration of dual-failure availability for protected-service paths. When spare capacity is needed any traffic on it is bumped without any added delay or consideration involved.

The other main difference vis-a-vis [28], [35] is that we consider true span-restorable networks in the sense well-established now by [11]–[19], [32]–[34], and others. In these networks span capacities can have an arbitrary breakdown of working and spare channel units, and all routing and cross-connecting functions are managed at the channel level, not the whole-fiber level. For example, if two bidirectional fiber systems are present on a span, each supporting 128 wavelength channels, then any partitioning of the 256 logical channels into working, spare and unequipped is possible in our designs. Moreover, some spans may only need two fibers, while others may have, say six, and so on: whatever capacity is required as part of an overall capacity-optimized demand routing and spare capacity placement plan. There is no constraint of requiring a completely four-fiber network or multiples of exactly four-fibers everywhere. This is far more general and efficient than having a purely four-fiber BLSR-like span capacitation.

III. Theory and Method for Mesh Network Availability Analysis

We now consider the problem of computing the availability of paths through a span-restorable mesh network that is designed for $R_1 = 100\%$. The approach draws on an established numerical approximation method often used in telecommunications availability analysis and on a new computational procedure to assess the effects of dual span-failures in the presence of different models of the restoration processes.

A. Basic Approach to Assessing Unavailability

The most common equation for availability is $A =$ MTTF/(MTTF $+$ MTTR) where A is the availability, MTTF is the mean time to failure, and MTTR is the mean time to repair. Related to this is the *unavailability* $U = 1 - A$. In the telecommunications industry, there is a well-validated framework for approximate availability analysis, based on the following assumptions.

1) A two-state "working–failed" model describes the status of all elements.
2) Elements fail independently aside from specific common-cause failure mechanisms that may be identified for specific consideration.
3) The in-service times (or times between failure) and repair times are independent memoryless processes with a constant mean.
4) The repair rate is very much greater than the failure rate. Equivalently, the MTTR is much smaller than the MTTF.

These assumptions are generally accepted for practical analysis when modeling a large number of system instances over a long operating time. Point 2) does not mean that we disregard known correlated-failure scenarios, such as if two cables share the same duct. In what follows, we would include that as a single failure whose impact is the simultaneous loss of both cables. Independence is only assumed between elements which are not obviously linked under a functional understanding of the system [7].

Point 4) is abundantly evident from experience and it is the main reason that network reliability (in the sense of graph disconnection) is not of greater practical concern. To illustrate, To and Neusy [1] give data for 100 miles of optical cable, the component of highest failure rate, for which MTTF $= 19\,000$ h, MTTR $= 12$ hours. Point 4) is also the basis for a simplified form of mathematical treatment because it implies that

$$\prod_{i=1\cdots N} A_i \approx 1 - \sum_{i=1\cdots N} U_i \tag{1}$$

where A_i is the availability of the ith element of a set of N elements in a *series* functional relationship, and U_i is the corresponding unavailability, $1 - A_i$. This says we can "add unavailabilities instead of multiplying availabilities" for elements in series. Freeman [29] illustrates the accuracy of this approximation with an example of six elements in series having unavailability from 10^{-5} to 10^{-3}. The accuracy of the approximation is better than 0.5% which "is typically far more precise than the estimates of element A_is" [29, p. 2073].

B. End-to-End Path Availability

If we first imagine a mesh network over which a path d is provisioned over n spans, but with no restoration mechanism, we would fairly accurately estimate

$$A_{\text{path}}(d) \simeq 1 - \sum_{i=1}^{n} U^p_{\text{link}}(i) \tag{2}$$

where $U^p_{\text{link}}(i)$ is the *physical unavailability* of the ith link in the path.

Therefore, one way of thinking about the action of span restoration is that it is a transformer of physical span unavailability to a lower *equivalent* unavailability of links on the span. This viewpoint argues that from the standpoint of an end-to-end path, there are two equally acceptable ways in which a link along the path can be in "up" state: either it is physically working, or, it is physically "down" but transparently replaced by a restoration path between its end nodes. Thus, if we define the *equivalent unavailability* of a link $U^*_{\text{link}}(i) = p(\text{link } i \text{ down} \cap \text{link } i \text{ not restored})$, then the path availability has the same form as (2) but is based on U^*_{link} not U^p_{link}. This line of reasoning reduces the problem of calculating path availability to determining the *equivalent unavailability* of links U^*_{link} in a span-restorable network based on the capacity in the network and the particulars of the restoration mechanism.

C. Determining the Equivalent Unavailability of Links in a Span-Restorable Network

Let us consider the first three orders of failure multiplicity $f = 1, 2, 3$ corresponding to single, double, and triple span failures. Our viewpoint is to determine the fraction of the *physical* unavailability of links on span i that is converted to *equivalent* unavailability. Clearly, with no restoration mechanism and no spare capacity, there is a 1 : 1 conversion: $U^*_{\text{link}} = U^p_{\text{link}}$. But with a restoration mechanism and a given distribution of spare

TABLE I
RELATIVE CONTRIBUTION OF FAILURE EVENTS TO THE UNAVAILABILITY OF LINKS ON A SPECIFIC SPAN i

Network State f	Description	Probability given span i failure	Time exposure	Capacity exposure	Contribution in example
1	Single span-failure, i	1	$\frac{\text{Av. Rest. Time}}{MTTR}$	0	1.38×10^{-8}
2	Dual span-failure, i and other span j	$(S-1)\cdot U^{p}_{link}$	$\frac{\text{Av. Rest. Time}}{0.5\times MTTR}\cdot R_2$	$1-R_2$	8.55×10^{-7}
3	Triple span-failure, i and other spans j, k	$\frac{(S-1)(S-2)}{2}\cdot(U^{p}_{link})^2$	$\frac{\text{Av. Rest. Time}}{0.33\times MTTR}\cdot R_3$	$1-R_3$	4.62×10^{-9}

capacity, the fraction of U^{p}_{link} that comes through to U^{*}_{link} depends on the failure states of other spans in the order-f failure scenario. Not all demands crossing span i may be restorable if there are coincident failures outstanding on other spans. It also depends in principle on the reconfiguration time for demands that are restorable but in practice this is a very small effect compared to outage due to multiple failure states that are not fully restorable. Conceptually, however, the restoration time for restorable demands may be longer for higher order failures. The action of the restoration mechanism within the spare capacity environment that survives the failure scenario can then be thought of as providing a mapping from physical to equivalent link unavailability. The mapping takes two effects into account:

1) First, a multiple failure state may or may not support the feasibility of restoration for all links on span i. In the absence of a priority scheme, each link on span i will have to share this exposure to a capacity-related risk of incomplete restoration. To characterize this, we define the multiple-failure restorability R_f. R_f is the fraction of the total failed working capacity that can be restored averaged over all f-order failure scenarios.
2) Secondly, we allow for a general reconfiguration outage-time for links that *are* restored. Although in practice, we assume techniques that reconfigure in a few seconds at the very most (making this factor insignificant), the general model allows that restoration time could become significant depending on the failure order. Thus, U^{*}_{link} is shielded from U^{p}_{link} on a span i through the following mapping as shown in (3) at the bottom of the page.

Table I provides expressions for the terms of (3) for $f = 1, 2, 3$. The domain of the exposure function is [0,1] and it expresses the extent to which links on a failed span are exposed to outage by virtue of incomplete restorability due to coincident failure states on other spans or, if restorable, the extent to which they are exposed to the restoration time. For example, if the network is designed for $R_1 = 100\%$ then for any $f = 1$ scenario the capacity exposure is zero and the time exposure is the restoration reconfiguration time. Table I gives the corresponding easily derived expressions for $f = 2, 3$. For lack of better data we assume a constant reconfiguration time in Table I. The time exposure values are the ratio of expected restoration time to the expected average time in the corresponding failure state. For instance the average time in an $f = 2$ state will be half the MTTR if the failures are independent.

Let us now use (3) and Table I to support an argument that in practice $f = 2$ failure scenarios will dominate the unavailability. This follows because by definition there is no capacity exposure in networks designed for $R_1 = 1$. Thus, for $f = 1$ we have only a time exposure to the restoration process. The next most likely failure scenarios are $f = 2$ states. For dual span-failures we may expect a significant capacity exposure in a network designed for $R_1 = 1$. Similarly for $f = 3$ scenarios, we expect that $R_3 \ll 1$ is likely. As an example, consider a 20-span network for which the restoration time is two seconds, the physical span MTTR is 12 hours, $U^{p}_{\text{link}} = 3 \times 10^{-4}$ and $R_2 = 0.5$, $R_3 = 0$. Based on results that follow, $R_2 = 0.5$ is conservative and $R_3 = 0$ is the worst case assumption for this argument. Under these assumptions, Table I shows the resulting contributions to U^{*}_{link}, indicating that dual span-failures are by far the main factor to consider.

As explained previously, the link equivalent unavailability values U^{*}_{link} represent the probability that any individual link (a single capacity unit) on a span is in the "failed and nonrestored" state at any point in time. As a consequence, there are a small number of situations where the unavailability of a *path* may be over estimated by (2). In other words, (2) is strictly a somewhat pessimistic estimate of the actual path availability. This arises in the rather specific circumstances of an (i, j) dual span failure scenario for a path that: i) crosses both spans i and

$$U^{*}_{\text{link},\,i} = U^{p}_{\text{link},\,i} \cdot \sum_{f=1,2,3} \underbrace{p(\text{state } f-1)}_{\substack{\text{physical unavailability}\\ \text{of other spans}}} \cdot \underbrace{\left[\underbrace{(\text{restoration time exposure}) \cdot R_f}_{\substack{\text{time exposure}\\ \text{(restorable fraction)}}} + \underbrace{1 \cdot (1 - R_f)}_{\substack{\text{capacity exposure}\\ \text{(unrestorable fraction)}}}\right]}_{\text{exposure function}} \tag{3}$$

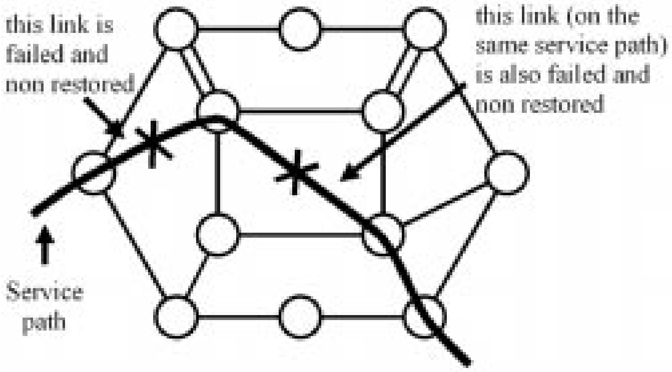

Fig. 2. Example of how a dual-span failure can be counted twice in the unavailability of a service path.

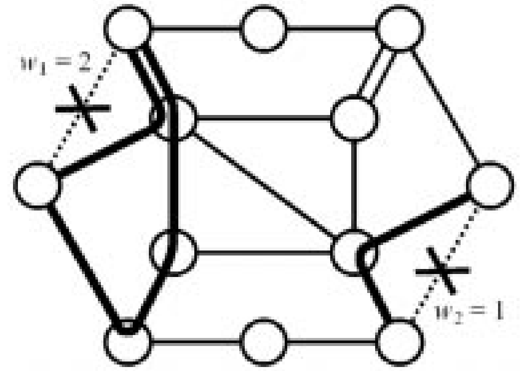

Fig. 3. Dual failure with no spatial interaction.

j when, ii) *both* of the spans i and j have less than complete individual restorability levels under the (i, j) failure scenario. Under these rather specific circumstances, the path availability will be slightly under-estimated because (2) will sum the individual link unavailabilities on each span above as independent contributors to the path outage whereas in reality they only contribute once to the unavailability arising as a single scenario. Fig. 2 shows an example of how a span failure could be counted twice in the unavailability of a given service path.

It is important to realize that this will not be a large numerical effect, however, and that any bias due to it is a pessimistic one. Of all dual failure combinations only those where both failures fall on the same path can possibly have this effect, *and only if* each of the spans in those scenarios individually has less that a full restoration level. If either individually has full restoration, then $U^*_{\text{link}} = 0$ for that span in the specific scenario, and the overestimation effect does not occur. For dual span failures several spans apart on the same path this is likely often to be the case due to the spatial separation of their individual restoration patterns. Where the over estimation will be more likely is for adjacent-span failures as shown in Fig. 2.

A final point is to also observe that the overestimation is based on assuming a random allocation of restored links on each span to paths transitting the span. In practice it could be advantageous to make a coordinated allocation of restored capacity on each span on a priority-path basis. In that case, a certain number of priority paths could effectively see $U^*_{\text{link}} = 0$ almost all the time. In fact, any time $R_2(i, j) > 0$, we could think of the top-most priority path being preferentially allocated the restored links, in which case its path availability would also be estimated by (2), but with $U^*_{\text{link}} = 0$ any time $R_2(i, j) > 0$. This is an observation to keep in mind when we see in the results just how extremely rare it is ever to see an $R_2(i, j) = 0$ result, and it is an important point to which we will return in closing the paper.

D. Analysis of Dual-Span Failure Scenarios to Obtain the Dual-Failure Restorability, R_2

Thus guided, we base the rest of our assessment of mesh availability on dual span-failure considerations. A specific dual-failure is denoted (i, j) naming the two spans involved and the order of failure. When considering a span-restorable mesh network there are four logical categories that describe dual-failure combinations:

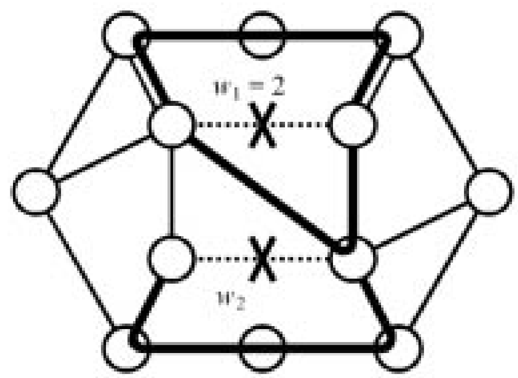

Fig. 4. Dual failure with possibility for spatial interaction depending on the value of w_2.

1) failures that are spatially independent (the restoration route-sets are disjoint);
2) failures with individual restoration route-sets that "contend" spatially;
3) cases where the second failure damages one or more restoration paths of the first failure;
4) cases where the two failures isolate a degree 2 node or effect a cut of the graph.

Figs. 3–5 illustrate the first three categories. In these illustrations, only the network of spare links is shown and working capacities are indicated only for the failed spans. The bold lines show the restoration paths formed to restore the indicated number of working links. Where relevant, the first failure to occur is associated with w_1.

Fig. 3 shows a case of no spatial interaction between the individual restoration route sets. Since both failures are fully restorable, $R_2(i, j) = 1$ for the scenario. Whether the restoration route-sets interact or not depends on the rerouting mechanism and the working and spare capacitation of the graph. Fig. 4 portrays a case in which there is spatial interaction between the routes of the respective restoration path sets. Depending on the spare capacities in Fig. 4, outage may or may not result from contention for available spare capacity between the two failures. As drawn, after restoration of w_1, only one restoration path is

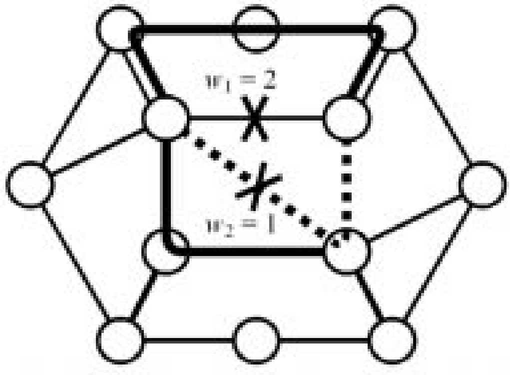

Fig. 5. Second failure hits first restoration path set.

feasible for the second failed span. Consequently, if $w_2 > 1$ the second failed span will not be fully restored. If $w_2 = 2$ in Fig. 4 we would say that for the failure of the particular *ordered pair* of spans (i, j) the restorability on span 2 is 1/2 and over the two failures together [$R_2(i, j)$ for the scenario] is 3/4. Note in general that if the order of the failed spans had been different, the restorability for each span could differ as well.

Fig. 5 is a case where the second failure damages restoration paths of the first failed span. In this case the number of restored links for the second failure depends on the remaining spare capacity after the first failure and depends on the "secondary" adaptability of the restoration process, i.e., whether the restoration mechanism is capable of (or allowed to) act again when a previously-deployed restoration path is severed by a second failure. The result also depends on whether sufficient spare capacity remains for the mechanism to repair the damage to its first path set by an updated restoration response.

Finally, a degree-2 node may be isolated by the failure of both its adjacent spans, creating an unrestorable situation, $R_2(i, j) = 0$. More generally if there is any cut of the network graph that contains only two edges, there are two ordered pairs that will disconnect the graph. In this case the amount of spare capacity and the adaptability of the restoration mechanism have no influence on the restorability and $R_2(i, j) = 0$ for any such dual-failure scenario.[4]

E. Determining the Network Average Dual-Failure Restorability, R_2

As alluded to in the prior section, Figs. 3–5 convey why it is very difficult to predict mesh availability analytically. The $R_2(i, j)$ of each failure scenario depends in detail on the specifics of the (i, j) pair, the failure sequence, the exact working and spare capacities, the graph topology, and the assumed restoration dynamics. To overcome this without approximation we use functional routing models of various restoration processes to explicitly determine the outcomes by computer-based experiments of all dual-failure restoration scenarios. In doing so we consider three technical models for the restoration process to see how varying levels of adaptability will affect R_2. Each of the models corresponds to different technical options for engineering the restoration mechanism. The restoration routing experiments are based on *k-shortest paths* (ksp) routing behavior for the basic single-failure response model [30], [31]. *ksp* means that each restoration path set is formed by first taking all paths feasible on the shortest route, followed by all paths feasible on the next shortest route not reusing any spare capacity already seized on the shortest route, and so on, until either all required paths are found, or no more *can* be found. This is known to be extremely close to maximum flow in typical transport networks [30] and can be computed in $O(n \log n)$ time [31]. ksp is also an accurate functional model for the self-organized restoration path-sets formed by the SHN™ protocol [16]. We now look at three levels of adaptability in restoration that can be modeled with ksp to determine R_2 by exhaustion of all $f = 2$ experimental trials. In all cases, the number of needed but infeasible restoration paths is denoted $N(i, j)$ for the (i, j) failure pair.

1) Static Restoration Preplans: The first model for restoration behavior is meant to represent restoration that is wholly based on centrally computed single-failure preplans. In this model, restoration of each span failure follows a predetermined plan, trying to restore both spans as if each was an isolated failure. If not enough spare capacity exists to support the superposition of both static preplans, restoration paths of the second failure are suppressed to conform to what is feasible within the spare capacity remaining after the first failure. When this is necessary, restoration paths for the second preplan that cross one of the spans, on which spare capacity is in shortage, are deleted in order of decreasing length.

2) First-Event Adaptive: The second model of restoration dynamics assumes that after a first failed span has been restored (but not repaired), the restoration mechanism of any second failure is aware of, and adaptive to, the changes in available spare capacity resulting from the first failure. Moreover, if any restoration path for the first failed span is routed over the second failed span the restoration mechanism will combine the requirements for the second span failure with the failed restoration paths of the first failure. However, new restoration paths are not sought between the end nodes of the first failed span. In other words, restoration paths affected by a second failure are referred into the second failure's w_i quantity.[5]

3) Fully Adaptive Behavior: The third model is of a completely adaptive restoration mechanism including both spare-capacity awareness following a first failure and re-restoration efforts for the first span from its original end-nodes for any damage to previously deployed restoration paths. The restoration mechanism will try to find new restoration paths between

[4]Clearly, degree-2 nodes are a logical "weak point" in networks where high availability is desired through restoration actions (as opposed to enhanced physical security of the spans). This is not particular to mesh networks, however; whether ring, mesh or APS-based, dual failures adjacent to a degree-two node have the same effect.

[5]The second model is based on the techniques of a distributed self-updating and self-organizing restoration protocol [16]. Such a protocol immediately gives "working" status to any spare link it uses in a restoration path. This inherently updates itself should it be triggered to act again to protect either new capacity or prior restoration paths in the event of a second failure. The awareness of the updated spare capacity environment following the first failure is implicit. For present work, the idea is that the complete range of realistically expected restoration responses is encompassed between the extremes of 1) and 3).

TABLE II
TEST NETWORK CHARACTERISTICS

Network	Nodes	Spans	Redundancy Non modular	Redundancy Modular	min cut < 3
Bellcore	11	23	55.9 %	92.1 %	3
EuroNet	19	37	50.6 %	78.3 %	5
6n14s	6	14	44.1 %	59.4 %	0
11n20s1	11	20	91.6 %	105.8 %	0
11n20s2	11	20	47.5 %	73.2 %	0
12n18s1	12	18	99.8 %	112.1 %	0
12n20s1	12	20	104.4 %	130.1 %	3
12n24s1	12	24	48.4 %	89.6 %	0
15n28s1	15	28	58.1 %	84.7 %	2
16n26s1	16	26	78.8 %	83.0 %	0
22n41s1	22	41	53.5 %	73.0 %	2

the end nodes of the first failed span when a second failure severs any of its initial restoration paths. This includes a release of surviving spare capacity on restoration paths of the first failure that were severed by the second, and repetition of the first span's restoration effort for the newly outstanding unrestored capacity, but in recognition of the spare capacity now withdrawn by the second failure.

IV. EXPERIMENTAL RESULTS FOR THE DUAL FAILURE RESTORABILITY

Using programs that implement each of the above restoration models, we conducted all possible ordered dual-failure experiments and recorded the number of nonrestorable working links $N(i, j)$ for each failure pair (i, j). The network-wide average dual-failure restorability is then

$$R_2 = 1 - \frac{\sum_{\forall (i,j), i \neq j} N(i, j)}{\sum_{\forall (i,j), i \neq j} (w_i + w_j)}. \quad (4)$$

R_2 restorability was calculated for several test networks which were capacitated with only the theoretical minimum of spare capacity for an assurance of $R_1 = 100\%$ in modular and nonmodular capacity environments. A summary of the test networks is given in Table II. The "$\min \text{cut} < 3$" column indicates the number of degree-2 nodes and other forms of weight-2 cuts in each test network.[6] The R_2 results, which average all individual $R_2(i, j)$ *include* the inevitable penalty of $R_2(i, j) = 0$ for the specific i, j failure combinations that this column identifies. "Redundancy" is the ratio of total spare capacity to the total amount of working capacity used in shortest-path routing of the working demands. Redundancy is of relevance because if the R_2 trials were repeated with arbitrarily high redundancy, correspondingly higher R_2 levels would result when adaptive restoration mechanisms are involved. The spare capacity available in these test networks is, however, only that arising from optimal solutions for integer and modular capacity variants of Herzberg's formulation [11] with a hop limit of five. Thus, each test case is very stringent and very specifically capacitated: it has only the minimum amount of spare capacity sufficient to yield $R_1 = 1$. Demand patterns were generated using the gravity-based demand model as in [32]–[34]. The modular capacity designs use module sizes of 12, 24, and 48 capacity units. There is a significant redundancy increase associated with modular capacity placement. This is a realistic side effect of modular capacity design however and is, therefore, a practical effect which is important to reflect in assessing the R_2 level of these R_1-designed networks.

TABLE III
TEST RESULTS FOR R_2 LEVELS IN R_1-DESIGNED NETWORKS WITH INTEGER CAPACITY

Network	R_2 Fully static	Partly adaptive	Fully adaptive
Bellcore	0.670	0.719	0.730
EuroNet	0.757	0.796	0.806
6n14s	0.596	0.611	0.616
11n20s1	0.792	0.806	0.826
11n20s2	0.592	0.612	0.619
12n18s1	0.665	0.695	0.713
12n20s1	0.703	0.734	0.758
12n24s1	0.633	0.665	0.675
15n28s1	0.729	0.766	0.775
16n26s1	0.754	0.764	0.769
22n41s1	0.800	0.820	0.830

Table III shows the R_2 results for the networks with nonmodular capacities, including the effects of degree-2 node disconnections, where they arise, in the averages. R_2 values range by network from 0.59 to 0.80 for the fully static behavior, 0.61 to 0.82 for the partly adaptive behavior, and 0.62 to 0.83 for the fully adaptive behavior. With these nonmodular minimum-capacity networks there seems to be relatively little benefit to the more adaptive restoration behaviors. It is remarkable nonetheless that these networks, designed only to withstand single span

[6]Only one network (12n20s1) has a weight-2 cut that is not associated with a degree-2 node.

TABLE IV
TEST RESULTS FOR R_2 LEVELS IN R_1-DESIGNED NETWORKS WITH MODULAR CAPACITY

Network	R_2		
	Fully static	Partly adaptive	Fully adaptive
Bellcore	0.791	0.891	0.914
EuroNet	0.812	0.917	0.935
6n14s	0.695	0.919	0.989
11n20s1	0.793	0.821	0.832
11n20s2	0.746	0.796	0.819
12n18s1	0.699	0.740	0.763
12n20s1	0.727	0.768	0.797
12n24s1	0.794	0.868	0.895
15n28s1	0.837	0.873	0.889
16n26s1	0.790	0.805	0.813
22n41s1	0.859	0.895	0.905

failures, can inherently also restore about 70% of failed capacity in double-failure scenarios.

Table IV shows R_2 results in the same networks when designed with modular minimum-total capacities. Now the more adaptive behaviors show R_2 values that are about 10% higher on average compared to the static behavior. R_2 for the fully adaptive behavior in the modular capacity environment is, however, often 90% and over.

Fig. 6 considers the EuroNet test network as a sample case for looking more closely at the effect of the restoration model on the distribution of individual $R_2(i, j)$ values in the nonmodular design. The histogram shows that even though the R_2 is not much higher with the fully adaptive behavior on average, the adaptive effects are seen to be more active and significant in terms of raising the $R_2(i, j)$ for the worst cases of low individual $R_2(i, j)$ under the static model. Whereas 30% of dual failure scenarios had $R_2(i, j)$ levels of 50%–60% under static restoration, only 5% are that low under fully adaptive restoration. In fact Fig. 6(b) implies that if all dual failures were equally likely for planning purposes, *95% of the time the network could support more than 60% dual-failure restorability*. It is important to also note in Fig. 6 that the entries in the 0%–10% bins correspond to dual adjacent span cuts at the five degree-2 nodes that this network model contains. No protection or restoration scheme can recover from these few scenarios because they entirely isolate the degree-2 node from the rest of the network.

Let us now illustrate how to interpret these results in terms of path availability. Consider a path composed of 4 hops in a 20-span network having an average dual failure restorability $R_2 = 0.7$ (which is relatively conservative given results in Tables III and IV), and a physical span unavailability $U^p_{\text{link}} = 3 \times 10^{-4}$. In a nonrestorable network the equivalent link unavailability is no different from U^p_{link} so the unavailability of the path would be $4 \cdot U^p_{\text{link}} = 1.2 \times 10^{-3} \approx 10.5$ hours/year. In comparison, with $R_1 = 1$ and $R_2 = 0.7$, the capacity exposure term of the equivalent link unavailability is: $U^*_{\text{link}} = (3 \cdot 10^{-4})^2\,(20-1)(1-0.7) = 5.13 \times 10^{-7}$ for an end-to-end unavailability of $U_{\text{path}} = 4 \cdot U^*_{\text{link}} = 2.1 \times 10^{-6} \approx 1$ min/year. Thus, in this example, the improvement brought on by the basic investment to achieve $R_1 = 100\%$ is a *reduction by approximately 585 times* in the average unavailability of a four-hop path.[7]

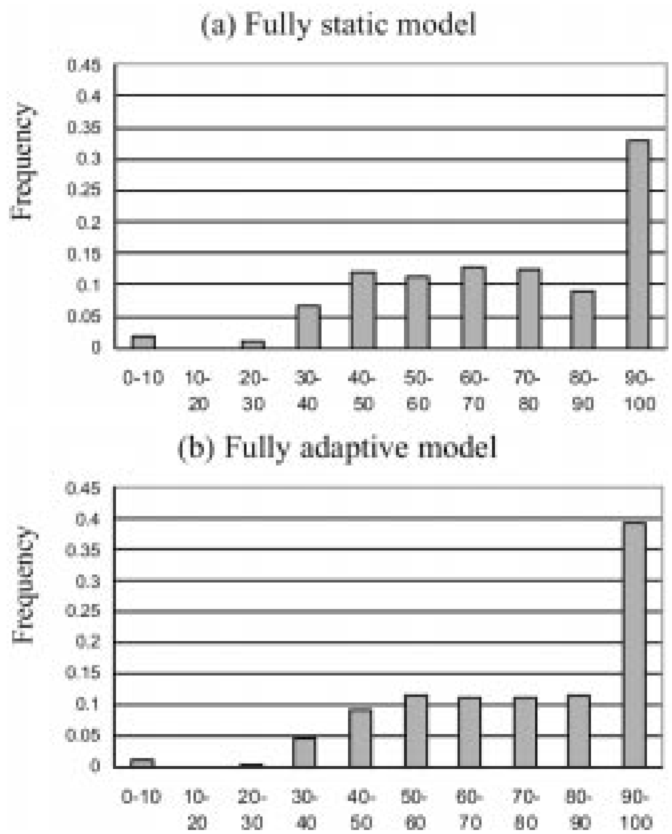

Fig. 6. Histogram of individual $R_2(i, j)$ levels per dual failure pair in Euronet.

Fig. 7 presents the R_2 level of each network plotted against its corresponding network redundancy. For the nonmodular and modular design cases, each test network is represented by its design redundancy on the x axis and the R_2 levels corresponding to the extremes of fully static and fully adaptive restoration models for that network on the y axis. The purpose of this scatterplot view of the data is to test the extent to which the dual-failure restorability correlates with the relative amount of spare capacity in the network. While both plots show a slight tendency toward greater R_2 with increasing redundancy, it is by no means a systematic progression. We interpret the almost flat general nature of the scatter of the test case designs as meaning that the availability depends more on the individual network and demand pattern than on the simple bulk redundancy of the network. In other words, some networks can have high redundancy but still be significantly less able to withstand dual failures than another network with significantly less percentage of spare capacity. This is probably an important factor to keep in mind when comparing ring networks to mesh based networks. Rings will generally have a greater bulk redundancy but this will not automatically imply higher availability because that redundancy is specifically arranged and locked-up in rings with limited network-wide access for the restoration.

[7]The example also gives a check on the numerical validity of disregarding the restoration time for restorable single failures if the latter is in the one to two second range (compared to one minute).

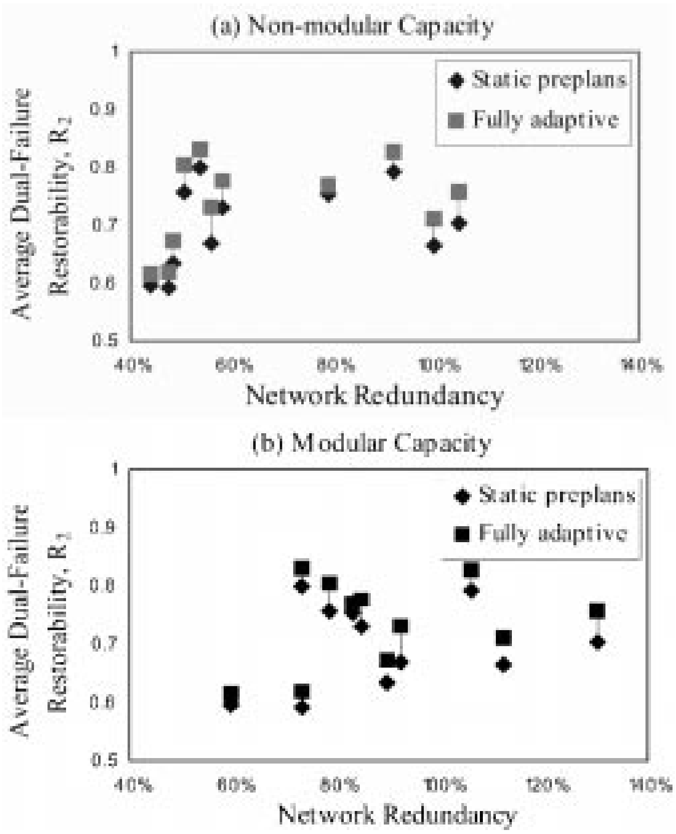

Fig. 7. Scatter plots of test network R_2 versus test network redundancy for modular and nonmodular capacity designs.

V. Concluding Discussion

A. Implications for a Mesh-Based Ultra-High Availability Class of Service

There is a very striking implication that can be realized from the data of Fig. 6. Let us state it as follows, then give the supporting explanation:

Proposition: A self-healing mesh network will provide higher availability for premium services than is possible with either ring or dedicated 1+1 APS.

This may initially seem to be a very surprising claim because dedicated 1+1 APS is usually considered the top end in terms of providing a high availability service. And yet when considering dual failures, it is obvious that any dual failure combination that hits both arms of the 1+1 path causes outage. In the language of this study, the 1+1 APS service has an assured exposure of $R_2 = 0$ for a specific set of $H_A H_B$ combinations of physical span failures, where $H_A H_B$ are the number of spans in the A and B signal feeds of the 1+1 APS arrangement.

But now consider the implications of Fig. 6, from the standpoint of a premium-grade service path in the mesh-restorable network. Fig. 6(b) shows that over *all* dual failure combinations, the restorability of total failed span capacity was over 20% in all cases that did not involve a dual failure next to a degree-2 node. In fact, for ~90% of the dual failures the R_2 level was over 50%. But for argument sake, let us assume that it was only 20%. Now if this minimum of 20% was always allocated preferentially to priority service paths, it would mean *that it would take a triple failure to affect such premium services*. Another way to view this, and to see why the service availability would exceed even that of 1+1 APS, is that it is the effect of the adaptive recovery effort in the dual failure case.

Imagine a side-by-side comparison of a 1+1 APS setup and a path through the mesh that has the same route as the "A" feed of the 1+1, illustrated in Fig. 8. A first-failure hits that path. The 1+1 APS switches to its "B" signal feed. The mesh deploys its first failure restoration reaction, for which $R_1 = 1$. So far so good. Both services survive. Now, let a second failure hit the "B" feed of the 1+1 APS setup, and, as a worst case in the mesh, assume the same failure hits the mesh restoration paths of the first failure. Now, the 1+1 service is out. But if the $R_2(i, j)$ of the dual failure scenario in the mesh is even 20%, the priority service in the mesh continues without outage. In fact, up to 20% of the affected paths can go on without outage in the mesh case. Because a path through any type of ring is also "down" in any state where there is a failure on the working path and on the reverse direction through the ring [1]–[3], the same argument serves as a existence proof that: *Given mesh minimum* $R_2 > 0$, *a certain number of priority services in a mesh can always enjoy higher availability than in a corresponding ring or 1+1 APS-based network*. In practice, the data of Fig. 6 suggest that a commercially significant fraction of all service paths could actually be given this form of ultra-high availability. Such "platinum" service customers could be guaranteed that their service path would only be affected by a *triple span failure*, taking their availability guarantees one order of magnitude higher. The only exception would be where the origin or destination node site is degree-2, which of course no scheme can protect against adjacent dual failures.

B. An Integrated Strategy for First-Failure "Protection" and Second Failure "Restoration"

Much is made these days of the distinction between "protection" and "restoration," it being often asserted in a very general way that protection is fast and restoration is slow. The view is too simplified, however, because there are really three categories of scheme to consider and the perception of how fast each can be depends whether it is assumed that path finding is time consuming or if it is path cross connection that is assumed to be slow. Moreover, it misses the always-present relationship between any restoration scheme and a corresponding preplanned protection scheme, which is derivable through distributed preplanning.

The three basic possibilities are:

1) schemes where the protection routes are known in advance and cross connection is not required to use the backup path(s), e.g.; UPSR, BLSR, 1+1 APS;
2) schemes where the protection routes are known in advance but cross connection is required in real time to use the backup path(s), e.g., MPLS shared backup path protection, or ATM Backup VP protection schemes;
3) schemes where the routes are found adaptively based on the state of the network at the time of failure, and the

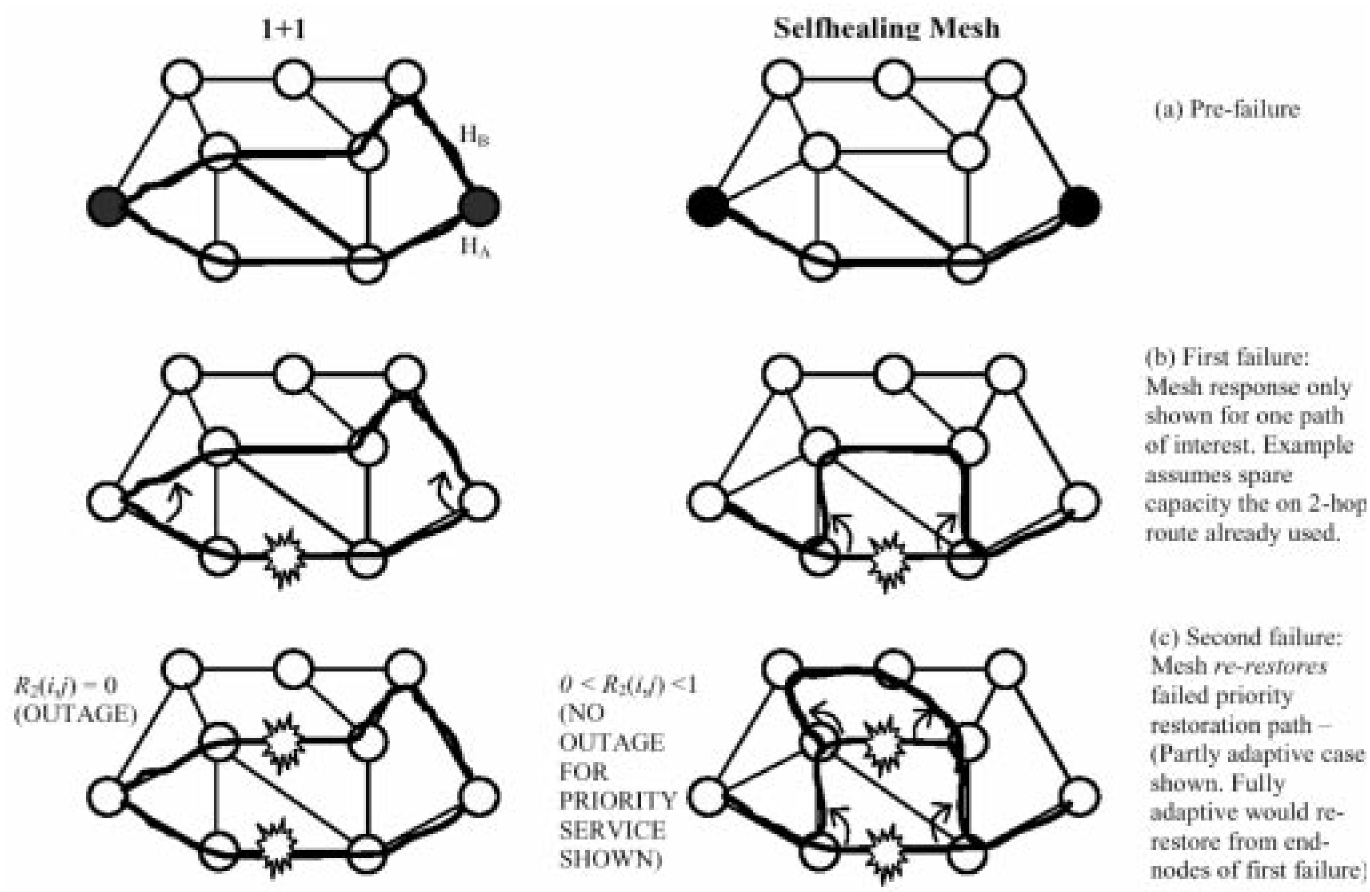

Fig. 8. Example showing how in general the nonzero $R_2(i,\,j)$ levels of mesh-restorable networks, can be allocated to priority paths in a way that lets them survive failures that 1+1 APS (or rings) would not survive.

cross connections to put the restoration routes into effect are also made in real time, e.g., self-healing networks, distributed self-organizing restoration schemes.

Schemes of type 3) are only slower than type 2) schemes if it is assumed that the restoration path-finding process is time consuming. But in such cases, distributed preplanning can create (and frequently update) a corresponding type 2) scheme where the protection routes are known in advance of failure. This is done by distributed preplanning with mock-failure trials responded to by the embedded restoration protocol. The concept is described more fully in [15] or [16]. It is quite a simple technique that retains all of the generality and database freedom of a distributed restoration algorithm, but provides a "protection" scheme of the type 2) above.

Consequently any type 3) scheme, which tend to be called restoration schemes, needs to be seen as actually providing the option of both a self-planning protection scheme *and* an on-demand dynamic adaptive real-time restoration scheme. The relevance to dual-failure recovery is that the response to a first failure can be based on a preplanned protection reaction, and only in the event of a second failure, is the truly adaptive but possibly slower restoration protocol itself executed directly in real time. Seen in this light, an adaptive restoration protocol is ideally suited for both requirements in a mesh network. First, it can serve as the engine for constant background preplanning of a fast "protection" reaction against single failures. Secondly, in the event of a second failure (more generally, any time the protection level is not 100%), it then executes directly in real time where its completely adaptive nature is exactly what is required to produce the highest possible overall recovery level.

C. Summary

We have described a partly theoretical and partly computational approach to the problem of availability analysis in mesh restorable networks. The overall method is practical to use but does not simplify-out the important details of irregular topology, capacity distribution, and restoration mechanism. The framework takes the view that in mesh networks that are fully restorable to single span failures, the major contribution to unavailability is exposure to incomplete restoration levels for dual span failures, as opposed to reconfiguration times or higher-order failures. The results showed that R_1-designed mesh restorable networks inherently enjoy high levels of dual failure restorability (R_2). The R_2 level can be over 90% for the combination of a fully adaptive restoration algorithm in a modular capacity design. These findings tend to counter the qualitative expectation in some quarters that mesh-restorable networks may not give as high a service availability as ring based networks because of their lower redundancy. What we see, however, is that despite the lower redundancy of the mesh, the generality of a highly adaptive routing process leads to an extremely high level of dual-failure restorability and, hence, to the high availability of paths realized over spans of the network. The fact that $R_2(i, j)$ is almost never zero, and usually at least 20% or more, also has striking implications for the high quality of premium-path service that a mesh-restorable network can provide, at no loss of assured R_1 restorability for all other services, and with no additional spare capacity other than that required for R_1 restorability itself.

REFERENCES

[1] M. To and P. Neusy, "Unavailability analysis of long-haul networks," *IEEE J. Select. Areas Commun.*, vol. 12, pp. 100–109, Jan. 1994.

[2] W. D. Grover, "High availability path design in ring-based optical networks," *IEEE/ACM Trans. Networking*, vol. 7, pp. 558–574, Aug. 1999.

[3] ——, "Resource management for fault tolerant paths in SONET ring networks," *J. Networks and Syst. Manage. (Plenum Publishing)*, vol. 7, pp. 373–394, Dec. 1999.

[4] R. Billinton and R. N. Allan, *Reliability Evaluation of Engineering Systems*, 2nd ed. New York and London: Plenum, 1992.

[5] P. O'Connor, *Practical Reliability Engineering*, 3rd ed. New York: Wiley, 1991.

[6] S. Rai and D. Agrawal, Eds., *Distributed Computing Network Reliability*. Piscataway, NJ: IEEE Press, 1990.

[7] D. D. Harms, M. Kraetzl, C. J. Colbourn, and J. S. Devitt, *Network Reliability*. Boca Raton, FL: CRC Press, 1995.

[8] C. J. Colbourn, *The Combinatorics of Network Reliability*. Oxford, U.K.: Oxford Univ. Press, 1987.

[9] R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*. Norwell, MA: Kluwer, Nov. 1998.

[10] M. Médard, S. Finn, and R. Barry, "WDM loop-back recovery in mesh networks," in *Proc. INFOCOM'99*, Mar. 1999, pp. 752–759.

[11] M. 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.

[12] W. D. Grover, B. D. Venables, M. H. MacGregor, and J. H. Sandham, "Development and performance verification of a distributed asynchronous protocol for real-time network restoration," *IEEE J. Select. Areas Commun.*, vol. 9, pp. 112–125, Jan. 1991.

[13] C. H. Yang and S. Hasegawa, "FITNESS: A failure immunization technology for network service survivability," in *Proc. IEEE GLOBECOM'88*, 1988, pp. 1549–1554.

[14] H. Komine, T. Chujo, T. Ogura, K. Miyazaki, and T. Soejima, "A distributed restoration algorithm for multiple-link and node failures of transport networks," in *Proc. IEEE GLOBECOM'90*, 1990, pp. 459–463.

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

[16] ——, "Self-organizing broad-band transport networks," in *Proc. IEEE: Special Issue Communications in the 21st Century*, vol. 85, Oct. 1997, pp. 1582–1611.

[17] H. Sakauchi, Y. Nishimura, and S. Hasegawa, "A self-healing network with an economical spare channel assignment," in *Proc. IEEE GLOBECOM'90*, 1990, pp. 438–443.

[18] B. D. Venables, W. Grover, and M. H. MacGregor, "Two strategies for spare capacity placement (SCP) in mesh restorable networks," in *Proc. IEEE ICC'93*, Switzerland, May 1993, pp. 267–271.

[19] W. D. Grover, T. D. Bilodeau, and B. D. Venables, "Near optimal spare capacity planning in a mesh restorable network," in *Proc. IEEE GLOBECOM'91*, 1991, pp. 2007–2012.

[20] J. D. Spragins, J. C. Sinclair, Y. J. Kang, and H. Jafari, "Current telecommunication network reliability models: A critical assessment," *IEEE J. Select. Areas Commun.*, vol. SAC-4, pp. 1168–1173, Oct. 1986.

[21] J. D. Spragins, "Telecommunication network reliability models based on network structures," in *Proc. Telecommunication Systems, Modeling, and Analysis Conf.*, Nashville, TN, Feb. 1993, pp. 79–90.

[22] J. W. Rowe, "A study of path availability improvements in a large network for various route diversity protection strategies, using failure event simulation," in *Proc. 2nd IEE National Conf. on Telecommunications*, 1989, pp. 167–172.

[23] M. R. Wilson, "The quantitative impact of survivable network architectures on service availability," *IEEE Commun. Mag.*, pp. 122–126, May 1998.

[24] P. Arijs, W. Van Parys, and P. Demeester, "Cost and availability comparison of WDM mesh and ring network architectures," in *Proc. 2nd Int. Workshop Design of Reliable Communication Networks, DRCN 2000*, Munich, Germany, Apr. 2000, pp. 33–38.

[25] H. C. Cankaya and V. S. S. Nair, "Reliability and availability evaluation of self-healing SONET mesh networks," in *Proc. Globecom 1997*, Nov. 1997, pp. 252–256.

[26] M. Barezzani, S. Pupolin, and M. Zorzi, "A new definition of transmission network availability with applications," *Euro. Trans. Telecommun.*, vol. 3, pp. 349–357, July–Aug. 1992.

[27] D. Vercauteren, P. Demeester, J. Luystermans, and E. Houtrelle, "Availability analysis of multilayer networks," in *Proc. 3rd Int. Conf. Telecom. System Modeling and Analysis*, Nashville, Mar. 1995, pp. 483–493.

[28] S. Lumetta, M. Médard, and Y.-C. Tseng, "Capacity versus robustness: A tradeoff for link restoration in mesh networks," *J. Lightwave Technol.*, vol. 18, pp. 1765–1775.

[29] R. L. Freeman, *Telecommunication System Engineering*, 3rd ed. New York: Wiley, 1996, pp. 444–447.

[30] 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 J. Select. Areas Commun.*, vol. 12, pp. 88–99, Jan. 1994.

[31] M. H. MacGregor and W. D. Grover, "Optimized k-shortest paths algorithm for facility restoration," in *Software—Practice & Experience*. New York: Wiley, 1994, vol. 24, pp. 823–834.

[32] J. Doucette and W. D. Grover, "Influence of modularity and economy-of-scale effects on design of mesh-restorable DWDM networks," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1912–1923, Oct. 2000.

[33] A. Lardies and A. Aguilar, "Planning methodology for SDH + optical networks," in *Proc. 1st Int. Workshop Design Reliable Communication Networks, DRCN'98*, Belgium, 1998.

[34] M. Bettin, G. Ferraris, and G. Pignari, "Comparison of protection and restoration schemes for SDH networks," in *Proc. 1st Int. Workshop Design Reliable Communication Networks, DRCN'98*, Belgium, 1998.

[35] M. Médard, S. Finn, and R. A. Barry, "WDM loop-back recovery in mesh networks," in *Proc. Infocom'99*, 1999, pp. 752–759.

Matthieu Clouqueur (S'99) is a graduate from Ecole Nationale Supérieure de l'Electronique et de ses Applications (National School of Electronics and Electrical Engineering), France. He is currently working toward the Ph.D. degree at the University of Alberta, Edmonton, AB, Canada.

Since he joined TRLabs, Edmonton, Alberta, Canada, in 1998, he has worked on blind equalization for nonlinear channels and on transport network survivability. He is now a Research Engineer at TRLabs in the Network Systems group. His research is on the availability of service in mesh-restorable networks.

Wayne D. Grover (S'74–M'76–SM'90–F'02) received the B.Eng. degree from Carleton University, Ottawa, ON, Canada, the M.Sc. degree from the University of Essex, England, and the Ph.D. degree from the University of Alberta, Edmonton, AB, Canada, all in electrical engineering.

He had ten years experience as Scientific Staff and Management at BNR, Ottawa and Edmonton, Canada, on digital networks, fiber optics, switching systems, digital radio, and others before joining TRLabs, Edmonton, Alberta, Canada, as its founding Technical Vice President, in 1986. In this position, he was responsible for the development of the TRLabs research program and contributed to the development of the TRLabs sponsorship base and he saw TRLabs through its early growth as a start-up to over the 100-person level. He now functions as Chief Scientist - Network Systems, at TRLabs and as Professor, Electrical and Computer Engineering, at the University of Alberta. He has patents issued or pending on 26 topics to date.

Dr. Grover has received two TRLabs Technology Commercialization Awards for the licensing of restoration and network design related technologies to industry. He is a recipient of the IEEE Baker Prize Paper Award for his work on self-organizing networks, as well as an IEEE Canada Outstanding Engineer Award, an Alberta Science and Technology Leadership Award, and the University of Alberta's Martha Cook-Piper Research Award. From 2001–2002, he is the recipient of the prestigious NSERC E.R.W. Steacie Memorial Fellowship. He is a P.Eng. in the Province of Alberta, Canada.