Fast Approximate Dimensioning and Performance Analysis of Mesh Optical Networks

Jean-François Labourdette, *Senior Member, IEEE* Eric Bouillet, *Member, IEEE*, Ramu Ramamurthy, Ahmet A. Akyamaç, *Member, IEEE*

***Abstract*— This paper presents a collection of approximation formulas that allow a network planner to quickly estimate the size of a mesh optical network with limited inputs. In particular, it provides a set of equations that relate number of sites, average fiber connectivity, demand load and capacity for various mesh protection architectures. These results can be used to easily and quickly estimate the amount of traffic that can be carried over a given network, or, conversely, given the traffic to be supported, to assess the characteristics of the topology required (in terms of number of nodes, connectivity). Finally, this analysis can be used to estimate the restoration performance that can be expected without resorting to extensive simulation studies.**

***Index Terms*— Optical networks, mesh networking, restoration, performance analysis.**

I. Introduction

While investigating, designing, or even negotiating a data-transport network it is always valuable to quickly anticipate a realistic gross estimate of its dimensions and cost. Very often the available information and/or time are insufficient to proceed with a full-scale study of the network. The task is further hindered by increasingly complex protection architectures. For instance, in shared mesh restoration, additional capacity is reserved to secure for every demand an alternate route that serves as backup in case of failure occurrence along its primary route. Since not all demands will be affected by a single failure, the reserved capacity can be shared among multiple demands. The amount of sharing and average time to re-establish services after any failures are difficult to estimate. The objective of this paper is to provide the framework and the formulas to estimate fundamental network characteristics and performance within an acceptable range of reality without having to resort to advanced network planning and modeling tools. These tools would then be used in a second phase when more detailed designs are required. The model and formulas presented are also very useful in understanding some fundamental behavior of networks as they capture and highlight the key relationships between different network characteristics (size, node degree, switch size, utilization,...) and traffic demand characteristics, and network performance (capacity, restoration times,...). We apply our model and techniques to optical mesh networks shown in Figure 1 made of optical switches connected to each other over inter-office Dense Wavelength Division Multiplexing (DWDM) systems. The optical switches provide lightpath-based connectivity between client equipments such as routers, also shown in Figure 1. This line of research is relatively new, but some early work can be found in [1], [2].

There exist several schemes for providing protection and restoration of traffic in networks. They range from protecting single links or spans to protecting traffic end-to-end at the path level. In addition, the protection capacity can be assigned in advance to pre-computed back-up routes or those routes can be computed in real-time after the failure. Different schemes achieve different trade-offs between restoration speed and capacity efficiency [3].

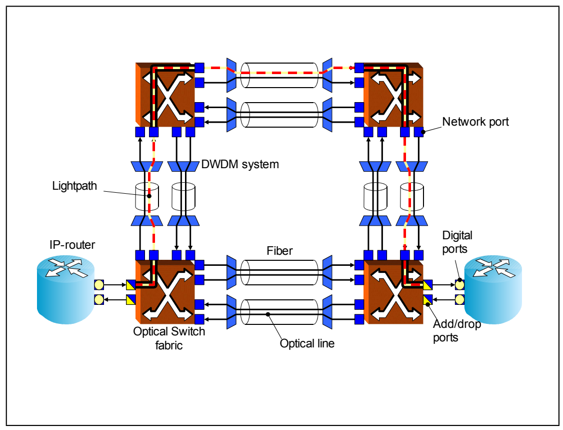

Fig. 1. Optical Mesh Network.

In end-to-end or path protection, the ingress and egress nodes of the failed optical connection attempt to restore the signal on a predefined backup path, which is SRLG[1]-disjoint, or diverse, from the primary path [5], [4], [6], [7]. Path diversity guarantees that primary and backup lightpaths will not simultaneously succumb to a single failure. Unlike local span protection, backup paths are provisioned with the working paths and thus the restoration does not involve further real-time path computations. Another aspect of path protection is that the restoration processing is distributed among ingress and egress nodes of all the lightpaths involved in the failure, compared to local span protection where a comparable amount of processing is executed by a smaller set of nodes, those adjacent to the failure. In the following we will only consider the cases where the protection path is failure-independent and is thus the same for all types of failures. By way of this restriction, the restoration paths may be computed and assigned before failure occurrence. There are two subtypes of

Jean-François Labourdette is with Verizon Enterprise Solutions Group. Eric Bouillet is with IBM T. J. Watson Research Center. Ramu Ramamurthy is with Cisco Optical Networking Group. Ahmet Akyamaç is with VPI Systems.
This work was conducted at Tellium, Inc.

[1]The concept of Shared Risk Link Group (SRLG) is used to model the failure risk associated with links riding the same fiber or conduit [3], [4]

path protection: (1) dedicated mesh (or 1+1) protection, and (2) shared mesh restoration.

Dedicated (1+1) mesh protection is illustrated in Figure 2. The network consists of four logical nodes (A to D) and two demands (AB and CD) accommodated across an eight node optical network (S to Z). The provisioning algorithm of this architecture computes and establishes simultaneously the primaries and their SRLG-disjoint protection paths. During normal operation mode, both paths carry the optical signal and the egress selects the best copy out of the two. In the example of Figure 2, all the optical channels on primary and secondary paths are active. In particular, the configuration reserves two optical channels between nodes S and T for protection. This is the fastest restoration scheme since for every lightpath one device is responsible for all the necessary failure detection and restoration functions. But it is also the most exigent in terms of resource consumption. If protection against node failure is also desired, then primary and backup paths must be node disjoint in addition to SRLG-disjoint.

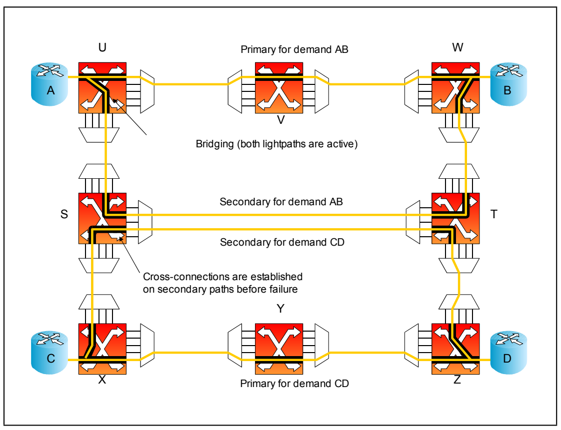

Fig. 2. Dedicated Mesh (1+1) Protection.

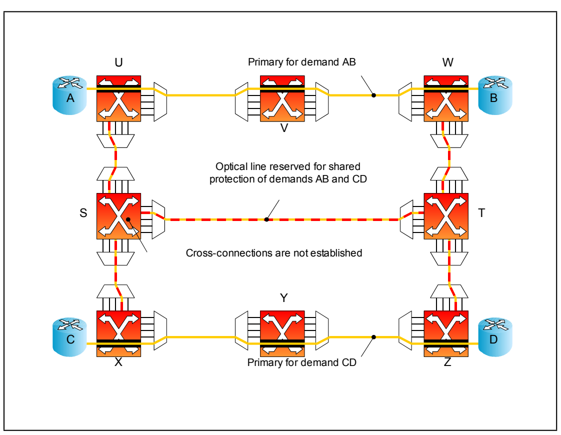

Fig. 3. Shared Mesh Restoration: Before Failure.

As in dedicated protection, in shared mesh restoration protection paths are predefined, except that the cross-connections along the paths are not created until a failure occurs (see Figure 3 and Figure 4). During normal operation modes the spare optical channels reserved for protection are not used. We refer to such channels as reserved (for restoration) channels.

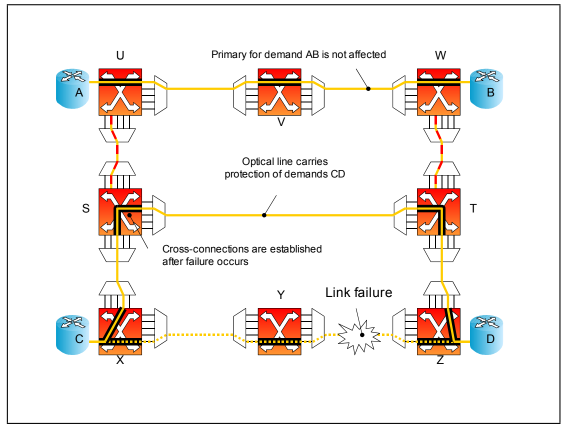

Fig. 4. Shared Mesh Restoration: After Failure.

Since the capacity is only ”soft reserved”, the same optical channel can be shared to protect multiple lightpaths. There is a condition though that two backup lightpaths may share a reserved channel only if their respective primaries are SRLG-disjoint, so that a failure does not interrupt both primary paths. If that happened, there would be contention for the reserved channel and only one of the two lightpaths would be successfully restored. Shared mesh restoration involves slightly more processing to signal and establish the cross-connections along the restoration path. There is thus an evident trade-off between capacity utilization and recovery time. In shared mesh restoration, node-diversity between primary and backup paths does not guarantee full protection against node failures. Additional sharing restrictions are required to guarantee restoration in case of node failure (for those lightpaths that did not terminate or originate at the failed node). Protection against node failure generally requires more bandwidth. See [5], [4] for further details and experimental results.

The procedure to route a lightpath consists of two tasks: (1) route selection, and (2) channel selection. Route selection involves computation of the primary and backup paths from the ingress port to the egress port across the mesh optical network. Channel selection deals with selecting individual optical channels along the primary and backup routes. The problems of selecting a route together with selecting channels on the route are closely coupled and if an optimal solution is sought both problems should be solved simultaneously. In this paper, we assume that routing computation is done with access to the complete network information, and that a k-shortest path approach is used for both the primary and backup paths. See [8], [9], [10] for comparison of routing efficiency when only partial information is available.

The outline of the paper is as follows. In Section II, we provide an analysis for approximating the path length and the protection capacity in mesh restorable networks. In Section III, we derive dimensioning formulas that approximate the number of lightpaths that can be carried in a maximally loaded network of given size and connectivity. We validate these approximations and give some examples in Section IV. In Section V, we use the analysis to estimate the restoration performance in mesh restorable networks. We conclude the paper in Section VI.

II. Approximate Path Length & Protection Capacity Analysis

In what follows, we represent a WDM network as a graph. Vertices (or nodes) represent the optical switches in the network, and edges represent (bi-directional) WDM links. We use n and m to denote respectively the number of vertices and edges. We call *degree* of a vertex the number of edges terminating at this vertex. The average vertex degrees of a graph is denoted δ. It is easily shown that $\delta = 2m/n$. In the remainder of this paper, we assume that all SRLGs default to one SRLG per link and one link per SRLG. We also assume no parallel links. Furthermore, we assume that traffic is uniform and equally distributed among all source-destination pairs[2].

A. Path Length Analysis

We are interested in the average path length of the primary or working path for a lightpath. We assume that it is equal to the average length of the shortest path. Assuming that the degree of the graph is greater than 2 (a reasonable assumption) and using a variation of the Moore bound [11], we obtain (see *Average Path Length* in Appendix B):

$$h \approx \frac{\ln[(n-1)\frac{\delta-2}{\delta}+1]}{\ln(\delta-1)} \tag{1}$$

Note that this is an approximation as one may want to take a longer working path than the shortest path either (a) to be able to find a diverse backup path in the case of a dedicated mesh protected lightpath, or (b) to maximize sharing in the case of a shared mesh restorable lightpath. This is a limitation of our current approach. However, it is our experience that shortest path length gives a very good approximation of working path length in both cases of dedicated and shared mesh protected lightpaths. In the case of dedicated mesh protection, we use a graph transformation technique (Figure 5) that essentially removes the source node (one less node) and its adjacent edges (δ less edges), as well as edges used by the working path (h less edges), to obtain a new graph. We re-apply our approximation of shortest path length on this new graph.

The computation of the average hop-length of the backup path $\mathbf{a}-\mathbf{z}$, in the context of dedicated mesh protection, is derived from a transformation of the graph as shown in the example of Figure 5. The transformation consists of (1) removing the h edges on the primary path, and (2) because we assume no parallel edges, selecting any neighbour $\mathbf{b}$ of $\mathbf{a}$, and removing node $\mathbf{a}$ and its δ adjacent edges, including edge $(\mathbf{a}, \mathbf{b})$. The purpose of the transformation is to determine the average degree δ' of this new graph (whose number of nodes is $n-1$). The backup path must be diverse from the primary path, and therefore, we should remove h primary edges from the graph before applying the formulae. However we also know that the second node after the source on the backup path cannot be the same as the second node after the source on the primary path. Otherwise either both primary and backup paths would traverse a common edge from the source to this second node, or there would be parallel edges between the two nodes which is contrary to our assumption. In order to understand this, suppose that we compute the length of a backup path from $\mathbf{a}$ to $\mathbf{z}$. The backup must consist of an edge (not primary) from the source $\mathbf{a}$ to any neighbor node $\mathbf{b}$, plus the length from neighbor $\mathbf{b}$ to the destination $\mathbf{z}$ of the path. Note that since we assume no parallel edges, and that the length of the backup path is equal to or longer than the length of the primary path, then $\mathbf{b}$ cannot be the same as $\mathbf{z}$. In order to compute the backup length from $\mathbf{b}$ to $\mathbf{z}$ we must remove: (1) all the primary edges, and (2) node $\mathbf{a}$ and all the edges adjacent to it, since we know that $\mathbf{a}$ cannot be part of the path from $\mathbf{b}$ to $\mathbf{z}$. In the transformed graph, the length of the backup is then one plus the length of the shortest path from $\mathbf{b}$ to $\mathbf{z}$.

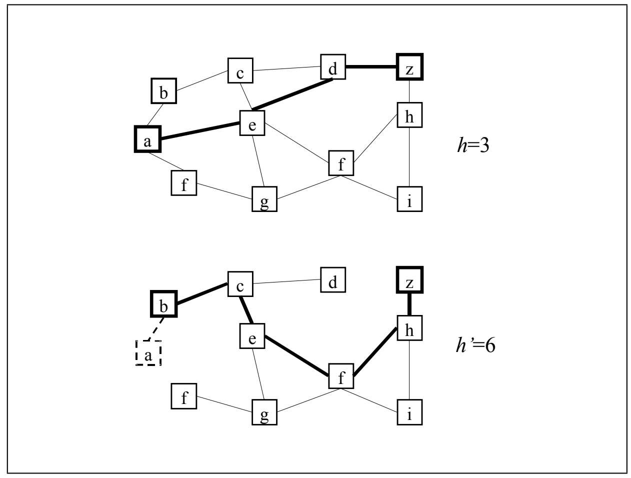

Fig. 5. The backup path cannot traverse edges already used for its primary $\mathbf{a}-\mathbf{z}$, and so these edges can be removed. Furthermore, the backup is at least as long as the primary and we assume no parallel edges, hence the backup is at least two hops long. This is represented by starting from any neighbor $\mathbf{b}$ of $\mathbf{a}$, other than $\mathbf{z}$, and adding one hop to the length of the backup path $\mathbf{b}-\mathbf{z}$.

Using this graph transformation approach, the new graph average degree is[3]:

$$\delta' = \frac{2[m-(\delta+h)+1)]}{n-1} \tag{2}$$

where h is the average hop-length of the primary path as computed in Equation (1). The average path length of the backup path for dedicated mesh protected lightpath is then approximated by the length of the shortest path in the transformed graph[4] (using Equation (1)), plus one[5] as:

$$h' \approx \frac{\ln[(n-2)\frac{\delta'-2}{\delta-1}+1]}{\ln(\delta'-1)} + 1 \tag{3}$$

Note that Equation (3) contains both δ and δ'. This is because $\delta-1$ in Equation (3) stands for the degree of the vertex origin

[2] A discussion of how to account for some non-uniformities can be found in Section VI

[3] Removing δ and h edges remove one edge too many because one edge is counted both as adjacent to the source node and part of the primary path, so the term $+1$ in the numerator.

[4] This assumes that there are no parallel edges.

[5] The term $+1$ is needed because the shortest path in the transformed graph starts one hop away from the source node.

of the path, as can be seen from the derivation of the Moore bound (see *Average Path Length* in Appendix B). Its degree is the average degree of the un-modified graph (used to compute the length of the working path) minus one to account for the working path, hence $\delta - 1$.

Figures 6 and 7 plot the approximations for h and h' against experimental path lengths computed in randomly generated networks (see *Random Graphs* in Appendix A) with average node degree 3 and 3.5, which are typical of real telecommunications networks. As seen from the plots, there is a very good match between the experiments and the approximation formulas for h and h'. Experimentation on similar networks with varying degree exhibits the same behavior.

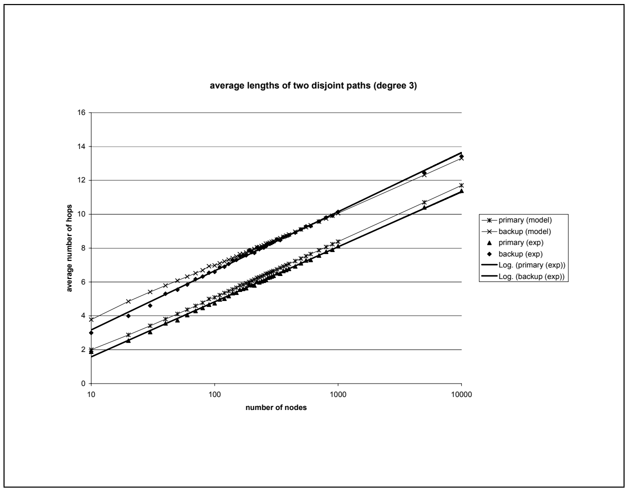

Fig. 6. Comparison of path length approximations against path length computed in randomly generated networks of average degree 3.

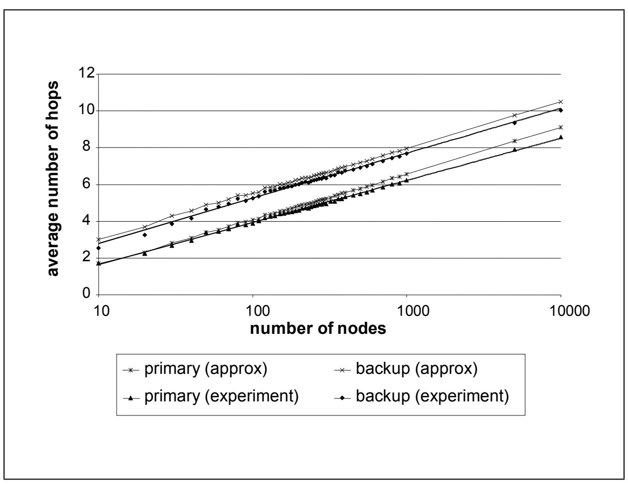

Fig. 7. Comparison of path length approximations against path length computed in randomly generated networks of average degree 3.5.

In the case of shared mesh restoration we introduce ϵ, the cost of a shareable channel reserved for restoration [5], [12], [13], [14]. The cost is actually the ratio of the cost of the same channel if it were not shareable to the cost when shareable (see [12], [13], [14] for details). The parameter ϵ ranges from 0 to 1. We express the average length of the backup path for shared mesh restorable lightpath h'' as:

$$h'' = h' + (1 - \epsilon)h_0 \tag{4}$$

where h_0 is determined experimentally (e.g., in the range of 0 to a few hops) or can sometimes be derived from the topology. For example, $h_0 = 0$ for a ring topology, or more generally for any topology where there is only one diverse path from the primary path. In general, the length of the backup path for shared mesh protected lightpath may be longer than the corresponding backup path length for dedicated mesh protected lightpath because sharing can be mined on different and therefore longer paths than the shortest paths. Alternatively, the backup path may be selected to be the same as in the case of dedicated mesh protection so as to minimize the combined length of the primary and backup paths. Then, sharing is determined after the paths have been selected, yielding $h_0 = 0$ and $h'' = h'$.

B. Protection Capacity Analysis

A recurrent question in shared mesh restoration and dedicated mesh protection network architectures addresses their respective overbuild measured in terms of capacity reserved for protection. This figure of merit is often expressed as the ratio of protection capacity to working capacity, with lower ratio meanings more capacity efficient protection. We define this ratio as follows:

$$\mathbf{R} = \frac{\textit{total number of protection channels}}{\textit{total number of working channels}} \tag{5}$$

While the answer to this question is trivial in the case of dedicated mesh protection with the leverage of Equations (1) and (3), it requires more thought for the case of shared mesh restoration. In the case of dedicated mesh protection, the average ratio of protection to working capacity is also the ratio of the average protection path length to the average working path length, and is independent of the number of lightpaths in the network[6].

$$\mathbf{R_d} \approx \frac{h'}{h} \geq 1 \tag{6}$$

where h and h' are given by Equations (1) and (3), respectively, as we assume uniform traffic demand.

In the case of shared mesh restoration, we cannot express $\mathbf{R}$ as the ratio of average protection path length to average working path length because some of the channels on a protection path can be shared between several lightpaths. We thus introduce a new parameter F, which represents the sharing factor of a protection channel, that is the average number of lightpaths whose protection path are using that channel. Then, the average ratio $\mathbf{R}$ of protection to working capacity can be expressed as:

$$\mathbf{R_s} \approx \frac{h''}{h}\frac{1}{F} = \mathbf{R_d}\frac{1}{F} + (1-\epsilon)\frac{h_0}{h}\frac{1}{F} \geq \frac{\mathbf{R_d}}{F} \tag{7}$$

where h and h'' are given by Equations (1) and (4), respectively, as we assume uniform traffic demand.

[6]Assuming an un-capacitated network.

Note that the formula for the ratio $\mathbf{R}$ does not assume that shared backup channels are either pre-assigned to particular backup paths or used as part of pool of shared channels [15]. The difference between the two approaches would be captured by different values of F. Note also that $\mathbf{R_s}$, contrary to $\mathbf{R_d}$ in the case of dedicated mesh protection, is not independent of the number of lightpaths as more lightpaths will provide for better sharing, thus increasing F, and therefore reducing $\mathbf{R_s}$. However, F, and $\mathbf{R_s}$ should become independent of the number of lightpaths when that number becomes large enough. Note that, in the case where the backup path for shared mesh restoration is selected to be the same as for dedicated mesh protection (i.e., $\epsilon = 1$), then $\mathbf{R_s} = \mathbf{R_d}/F$. If in addition F is fixed to one by capping the amount of sharing that is acceptable (no sharing allowed) [16], $\mathbf{R_s}$ becomes the same as $\mathbf{R_d}$.

C. Sharing Analysis

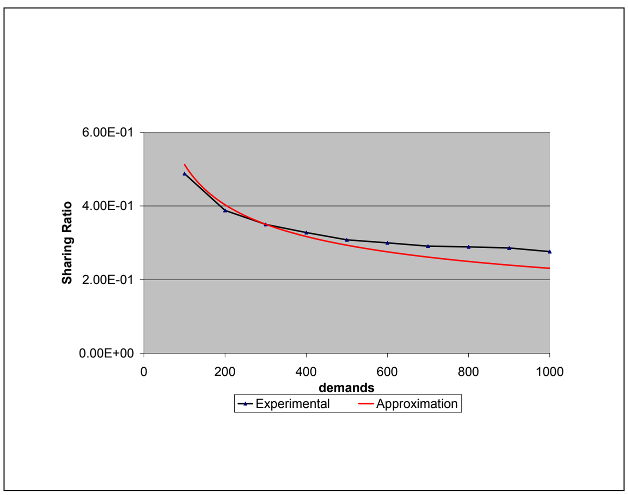

Fig. 8. Sharing ratio $\mathbf{R_s}$, experimental versus approximation as demand increases on a 50 node, 75 link chordal ring network (degree 3).

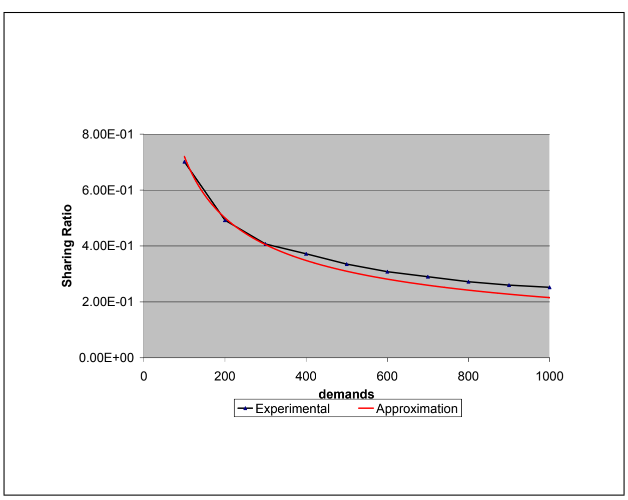

Fig. 9. Sharing ratio $\mathbf{R_s}$, experimental versus approximation as demand increases on a 150 node, 300 link chordal ring network (degree 4).

The sharing analysis consists of determining the relationship between F and the number of lightpaths, or demands, in a network. The analysis first determines the number of lightpaths whose backup path traverses an arbitrary link l, and then the largest number of corresponding primary paths that traverse any given link. That number is the number of backup channels required on the arbitrary link l, and F is simply the ratio of lightpaths whose backup path traverses l divided by the number of backup channels required. The details of the analysis are given in *Ratio of Shared Protection to Working Capacity* in Appendix C. Results comparing the value of $\mathbf{R_s}$ to this approximation are given here. Figure 8 and Figure 9 compare the approximation of the sharing ratio against experimental sharing ratios computed in random chordal ring graphs of respectively 50 nodes, 75 links, and 150 nodes, 300 links. The approximation is rather accurate but tends to over-estimate the amount of sharing achieved as the number of demands increases.

III. DIMENSIONING MESH OPTICAL NETWORKS

We introduce a node model with the following parameters:

S = size of switch
N_a = number of add/drop ports
N_{ot} = number of network-side ports used by originating/terminating working path
N_{th} = number of network-side ports used by through working path
N_p = number of protection ports (dedicated mesh and shared mesh)
γ = switch loading
P_r = ratio of add/drop ports used for LAPS protection to add/drop ports used for service[7]
T = ratio of through to working capacity[8]
$\mathbf{R}$ = ratio of protection to working capacity

We model a node as shown in Figure 10. The ports on a switch are categorised as either add/drop ports that are facing towards the outside of the network and connected to client equipment, or network-side ports, that are facing towards the inside of the network and that support trunks connecting the nodes to each other. The primary path of an originating/terminating lightpath uses one or more add/drop ports from the pool of N_a ports and a network-side port from the pool of N_{ot} ports. The primary path of a through lightpath uses two ports from the pool of N_{th} ports. A restoration channel on the backup path of either a shared mesh restorable or a dedicated mesh protected lightpath uses a port from the pool of N_p ports. The sizes of the pools verify the following conservation equations:

$$N_a + N_{ot} + N_{th} + N_p = \gamma S \tag{8}$$

$$N_a = N_{ot}(1 + P_r) \tag{9}$$

[7] Drop-side protection refers to ports on the drop-side of a switch (as opposed to the network side) that are dedicated to provide protection to working drop-side ports. P_r is 0 if no drop-side protection is used; 1 if 1+1 drop-side protection is used; $1/N$ if $1:N$ drop-side protection is used

[8] Working capacity includes ports used for through working paths.

$$N_{th} = T(N_{ot} + N_{th}) \tag{10}$$

$$N_p = \mathbf{R}(N_{ot} + N_{th}) \tag{11}$$

Equation (9) captures the fact that some of the drop-side ports are used for drop-side protection.

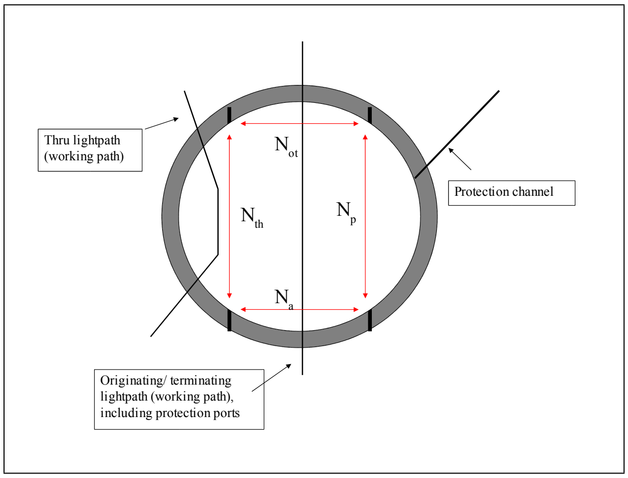

Fig. 10. Node Model.

Given a path of length h, the path traverses $2(h-1)$ ports at through or intermediate switches (two per switch) while the path uses two additional ports on the network side of the originating and terminating switches, yielding $T = 2(h-1)/2h = (h-1)/h$. Rewriting as $1 - T = 1/h$, and plugging along with Equation (9), Equation (11), and Equation (10) into Equation (8), we obtain after simplification:

$$N_a = \gamma S \frac{1 + P_r}{1 + P_r + (1 + \mathbf{R})h} \tag{12}$$

In the case of lightpaths with no drop-side protection ($P_r = 0$) and no network-side protection ($R = 0$), we have $N_a = \gamma S/(1 + h)$, as expected. The average number of lightpaths in a maximally loaded mesh network (dedicated and shared mesh) can then be derived as:

$$L_{network} = \frac{N_a \times n}{2(1 + P_r)} = \frac{\gamma S}{2} \frac{n}{1 + P_r + (1 + \mathbf{R})h} \tag{13}$$

Note that $L_{network} \propto o(n/h) = o(n/\ln n)$. From Equation (13), and using $\delta = 2m/n$, we can write the average number of lightpath per link in a maximally loaded mesh network as:

$$L_{link} = \frac{L_{network} \times h}{m} = \frac{\gamma S}{\delta} \frac{h}{1 + P_r + (1 + \mathbf{R})h} \tag{14}$$

Note that $L_{link} \rightarrow \gamma S/[\delta(1 + \mathbf{R})]$ when $n \rightarrow \infty$, independent of h. The average number of lightpaths per node in a maximally loaded mesh network is:

$$L_{node} = \frac{L_{network} \times (1 + h)}{n} = \frac{\gamma S}{2} \frac{1 + h}{1 + P_r + (1 + \mathbf{R})h} \tag{15}$$

Note that $L_{node} \rightarrow \gamma S/[2(1 + \mathbf{R})]$ when $n \rightarrow \infty$, independent of h. The formula for the number of lightpaths in a maximally loaded network is a function of the protection ratio $\mathbf{R}$. $\mathbf{R}$, in the case of shared mesh protected lightpaths, depends in turn on the number of lightpaths in the network through the sharing factor of shared backup channels, F. Therefore, determining $L_{network}$ and $\mathbf{R_s}$ requires solving a fixed point equation.

IV. Dimensioning Examples

Let us demonstrate how these formulas can be used to dimension optical mesh networks. For reasonable size networks, we have measured $\mathbf{R}$ to be in the range of 1.2 to 1.5 for dedicated mesh protection and 0.4 to 0.8 for shared mesh restoration. Also, operational networks are usually run around 70% utilisation. Finally, we assume here that $P_r = 0$. Figure 11 plots the maximum number of lightpaths as a function of the number of nodes (with average node degree three) for the case of unprotected demand ($\mathbf{R} = 0$), shared-mesh restorable demands ($\mathbf{R} = 0.7$), and dedicated mesh protected demands ($\mathbf{R}$ obtained from Equation (6)). Two sets of curves are given for two different utilisation levels of switches of size 512, with $\gamma = 0.7$ and $\gamma = 0.9$ utilisation levels. From these curves, it is easy to determine the maximum number of lightpaths that can be supported for a given network size and at a given network utilisation. Inversely, given a certain amount of traffic that needs to be carried, it is easy to estimate the number of nodes given other characteristics such as average node degree and switch size.

In Table I and Table II, we compare our approximate formulas with experimental results obtained for both a real network and several random networks of different sizes. In the experiments, the networks were loaded with uniform demand, until one of the nodes reached capacity and five percents of new demands got blocked[9]. The network characteristics and the usage[10] achieved at the point where 5% of the uniform demand gets blocked are shown in Table I, for both dedicated and shared mesh routing. The resulting usage or network utilization ranges from 64% to 88%, which is characteristic of real networks. The resulting total number of demands or lightpaths carried in the experimental networks is shown in table II ($L_{network}$), along with the average working and backup path length, and the sharing ratio $\mathbf{R}$, for both dedicated and shared mesh routed demands[11].

Our approximations were then used to estimate the total number of lightpaths that could be carried in networks with the same characteristics as the experimental networks (in terms of number of links and nodes), and for the same network utilization. These results are reported in Table II. In all networks studied and for results obtained through our theoretical formulas, the lengths of the primary and backup paths of shared mesh restorable lightpaths are the same as

[9] We do not consider any limitation on the number of wavelengths per link, only ports per node. Furthermore, in the context of opaque, or OEO, network, there is no limitation from wavelength conversion capability.

[10] Usage refers to the number of ports used out of total numbers of ports per switch, averaged over all switches.

[11] The experimental average path lengths and protection to working ratios are determined before blocking occurs.

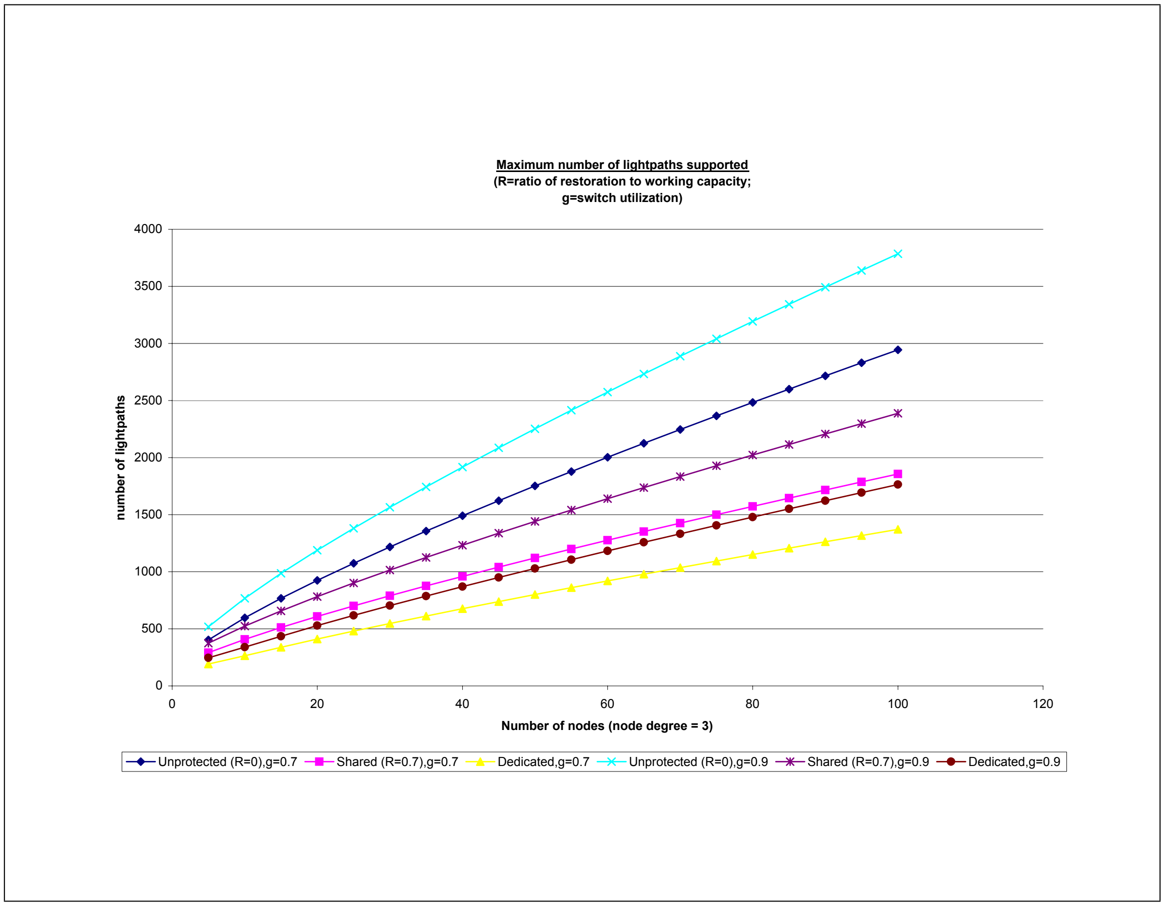

Fig. 11. Maximum number of lightpaths as a function of number of nodes for different switch utilisations γ (switch size = 512).

those of the primary and backup paths of dedicated mesh protected lightpaths (that is $\epsilon = 1$ and $h'' = h'$), and backup channels are pooled. While backup channels are also pooled in the experimental cases, primary and backup path lengths may vary between shared mesh restorable and dedicated mesh protected lightpaths because the routing algorithm used iterates over many combinations of paths. For shared mesh routed demands[12], we considered two cases. In the first case, we computed the average working and backup path lengths and the sharing ratio, assuming uniform demand distribution. In the second case, we attempted to reflect more accurately the distribution of the demand used in the experimental network by re-using the average working path length from the experimental results. From it, we computed the average backup path length, sharing ratio, and the maximum number of demands that can be supported.

One first notices from Table II that the theoretically-derived value of $\mathbf{R}$ is in general slightly lower than the experimental value for dedicated mesh protection, and more so for shared mesh restoration in random networks. That reflects that fact that our approximations are under-estimating the backup path length and are rather optimistic with respect to the amount of sharing that can be achieved, as can be observed in Table II.

A second interesting observation is that, when plugging the experimental values of h and $\mathbf{R}$ in the approximation formula for the total number of demands supported $L_{network}$, we obtain results that are within a couple of percents of the experimental results for random networks, and within 8% for the real network. This validates the accuracy and the applicability of the node model and traffic conservation equations developed in Section III.

As can be seen in Table II, in the case of *random networks*, the theoretical and experimental results for the maximum number of demands $L_{network}$ that can be supported are within 11% of each other for dedicated mesh protection and within 28% (but only 15% if we exclude random network 2) for shared mesh restoration. Plugging the experimental average working path length h into the approximation formulas has some impact, reducing to 7% and 23% (15 % excluding random network 2) the difference for dedicated mesh protection and shared mesh restoration, respectively.

In the case of a *real network*, the theoretical model overestimates by 20% and 38% for dedicated mesh protection and shared mesh restoration respectively, and 13% and 24% when plugging the experimental working path length into the theoretical formulas. Here again, using the experimental working path length has a non-negligeable effect.

[12]Note that for shared mesh routed demands, the number of lightpaths is a function of the sharing ratio $\mathbf{R}$, which is itself a function of the number of lightpaths. The determination of $\mathbf{R}$ would thus normally require solving a fixed point equation. However, since this is an approximation, we computed $\mathbf{R}$ for 1000 lightpath demands, and assumed that its value does not vary significantly around this point.

	exp h	th h	exp h', h''	th h' h''	th h' w exp h	exp R	th R	th R w exp h	exp L_{net}	th L_{net}	th L_{net} w exp h
Real network - dedicated mesh	4.01	3.35	5.34	4.80	4.83	1.33	1.43	1.20	730	895	833
Real network - shared mesh	4.01	3.35	5.34	4.80	4.83	0.40	0.16	0.15	1217	1781	1551
Random network 1 - dedicated mesh	3.74	4.10	5.47	6.00	5.96	1.46	1.46	1.59	963	865	897
Random network 1 - shared mesh	3.74	4.10	5.47	6.00	5.96	0.28	0.18	0.25	1629	1644	1692
Random network 2 - dedicated mesh	3.33	3.22	5.43	4.79	4.81	1.63	1.49	1.44	439	476	469
Random network 2 - shared mesh	3.39	3.22	5.53	4.79	4.82	0.66	0.24	0.24	622	820	787
Random network 3 - dedicated mesh	4.10	4.14	6.15	5.61	5.60	1.5	1.35	1.36	895	941	945
Random network 3 - shared mesh	4.09	4.14	6.09	5.61	5.60	0.37	0.22	0.22	1500	1629	1645
Random network 4 - dedicated mesh	4.67	4.61	6.78	5.98	5.98	1.45	1.30	1.28	1191	1276	1269
Random network 4 - shared mesh	4.59	4.61	6.64	5.98	5.97	0.36	0.13	0.13	2023	2350	2359
Random network 5 - dedicated mesh	4.91	4.95	7.07	6.26	6.26	1.44	1.26	1.27	1510	1614	1620
Random network 5 - shared mesh	4.86	4.95	6.98	6.26	6.25	0.29	0.13	0.13	2736	3028	3076
Random network 6 - dedicated mesh	2.43	2.34	3.46	3.49	3.50	1.42	1.49	1.44	806	815	803
Random network 6 - shared mesh	2.44	2.34	3.51	3.49	3.5	0.35	0.32	0.32	1304	1377	1334
Random network 7 - dedicated mesh	3.01	2.90	4.16	4.01	4.02	1.38	1.38	1.33	1340	1392	1371
Random network 7 - shared mesh	3.01	2.90	4.13	4.01	4.02	0.27	0.16	0.16	2284	2522	2451
Random network 8 - dedicated mesh	3.43	3.30	4.67	4.39	4.40	1.36	1.33	1.28	1706	1790	1761
Random network 8 - shared mesh	3.41	3.30	4.70	4.39	4.4	0.25	0.17	0.17	3058	3318	3232
Random network 9 - dedicated mesh	3.60	3.53	4.78	4.61	4.61	1.33	1.31	1.28	2297	2353	2335
Random network 9 - shared mesh	3.60	3.53	4.79	4.61	4.61	0.21	0.09	0.09	4099	4542	4471

TABLE II

COMPARATIVE RESULTS FROM EXPERIMENTS AND APPROXIMATION FORMULAS FOR PATH LENGTH OF WORKING AND BACKUP PATHS, SHARING RATIO $\mathbf{R}$, AND TOTAL NUMBER OF DEMANDS SUPPORTED.

As expected, the approximations become less accurate as the network becomes less random. But these results based on first order characteristics are very encouraging in justifying the applicability of our approach and the resulting approximation formulas, and certainly warrant further research into refining them.

V. RESTORATION TIME BEHAVIOR

Shared mesh restoration studies using simulation tools show that restoration times are mainly influenced by the number of failed lightpaths processed by a switch during restoration [17]. In particular, the worst case occurs when all lightpaths terminate at the same two end switches rather than at switches distributed throughout the network. Furthermore, simulation studies have shown that, for a given topology and a given set of primary and backup routes, the restoration time increases roughly linearly as the number of lightpaths simultaneously failed is increased [17]. Thus, a coarse analytical approximation can be constructed which assumes the worst-case scenario involving the maximum number of lightpaths that are processed by the same number of end nodes. The analytical approximation assumes a linear dependency between the restoration time and number of lightpaths restored. The average restoration latency can then be approximated using the worst case assumption that L lightpaths with h-hop primary paths and h''-hop backup paths all terminate at the same two switches and that a failure occurs in the middle link of the primary path (in terms of number of hops). The analytical approach uses a linear model to approximate the average restoration latency as follows:

$$T_r = T_0 + (L-1)S \tag{16}$$

where T_r is the final restoration latency for all lightpaths, T_0 is the restoration latency for the first lightpath (or a single lightpath) and S is a parameter that represents the slope of the linear tail of restoration latency versus number of lightpaths restored. T_0 can be obtained using a modeling tool, and/or real-life testbed or field results. T_0 can also be determined analytically based on the details of switch and restoration protocol architectures. The restoration time for one lightpath can be computed analytically based on fault detection/triggering times, messaging times, processing and cross-connect execution times at switches, propagation delays, etc. T_0 can thus either be measured or estimated as the time to restore a single lightpath in a network of size n and degree δ. The average lengths of the primary and backup paths given by h and h'' in Equations (1) and (4) can be used to estimate the propagation delays. S is the cross-connect time, and can be obtained by taking into account the switching fabric cross-connect time as well as intra-node communication and messaging times involved in setting up a cross-connect during

Network type	*n*	*m*	*degree*	*usage for ded. mesh*	*usage for shared mesh*
Real network 1	50	88	3.52	0.64	0.68
Random network 1	50	75	3.00	0.75	0.75
Random network 2	25	36	2.88	0.67	0.64
Random network 3	50	74	2.96	0.79	0.77
Random network 4	75	113	3.01	0.77	0.76
Random network 5	100	152	3.04	0.77	0.78
Random network 6	25	49	3.92	0.87	0.88
Random network 7	50	101	4.04	0.86	0.86
Random network 8	75	150	4.00	0.81	0.84
Random network 9	100	202	4.04	0.84	0.86

TABLE I

CHARACTERISTICS OF REAL AND RANDOM NETWORKS OF GIVEN SIZE (n NUMBER OF NODES, m NUMBER OF LINKS), NODE DEGREE, AND USAGE.

a restoration event. Typically, S would be of the order of a few milliseconds.

Using the result for a single lightpath, the average number of lightpaths per link (respectively node) in a maximally loaded network, and some modeling, we can derive results for average restoration time in a maximally loaded network. We now consider the case where there are L lightpaths failing for the same network configuration as above. The number L of lightpaths per link is obtained from Equation (14). We further assume that all those lightpaths originate and terminate at the same end nodes. Therefore, the restoration requests are going to contend for resource at the end nodes to perform bridge and switch cross-connects. These are very conservative assumptions in a worst case scenario. Most likely, under more realistic assumptions, we would observe shorter restoration times.

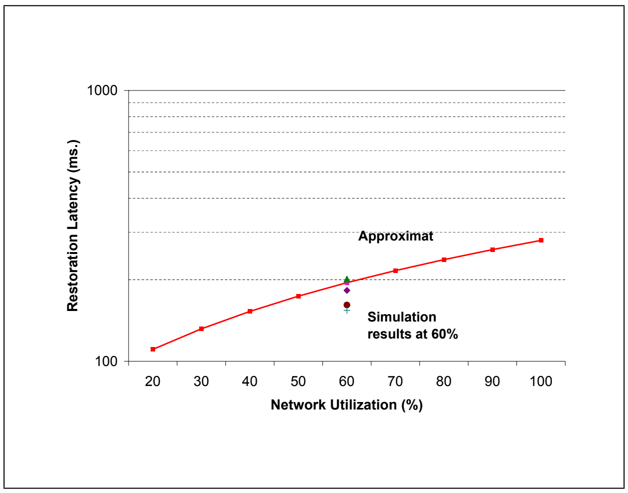

Fig. 12. Analytical results vs. simulation results for hypothetical 50-node network.

Figure 12 compares the coarse analytical results obtained for a 50 node network with the simulation results. The analytically calculated restoration latency curve is shown in Figure 12 versus the network utilisation. We conducted an actual 50 node network study with a utilisation of 60% (where $L = 36$ lightpaths failed for the analytical approximation) and a backup channel sharing ratio of 0.46 [17]. In Figure 12, we superimpose the simulation results for five single failure events affecting the most number of lightpaths at 60% utilisation. As can be seen from this figure, the analytical approximation yields a restoration latency which is within the same order or magnitude of the results obtained using simulation.

Figure 13 compares the coarse analytical results obtained for a 43 node network with the simulation results. The analytically calculated restoration latency curve is shown in Figure 13 versus the network utilisation. We conducted an actual 43 node network study with a utilisation of 43% (where $L = 18$ lightpaths failed for the analytical approximation) and a backup channel sharing ratio of 0.92. In Figure 13, we superimpose the simulation results for five single failure events affecting the most number of lightpaths at 43% utilisation. As can be seen from this figure, the analytical approximation again yields a restoration latency which is within the same order

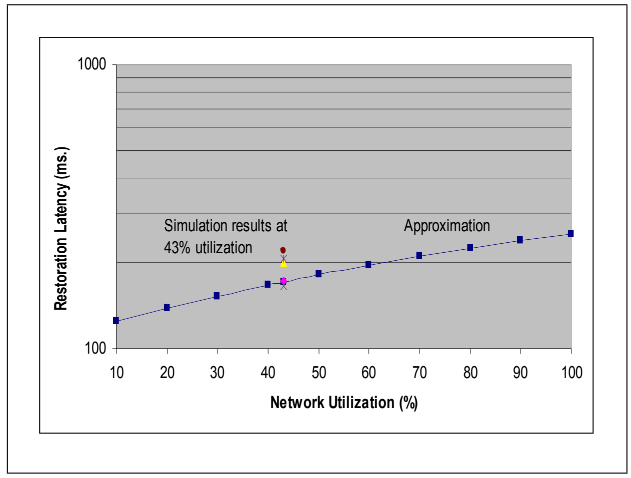

Fig. 13. Analytical results vs. simulation results for hypothetical 43-node network.

or magnitude of the results obtained using simulation. These behaviours are typical of similar studies we have performed for different networks.

Having validated the basic model and parameters, we can now use our approximation formulas for different networks and estimate the restoration times one could expect for a maximally loaded network at different utilization levels. This is shown in Figure 14 for two different size networks.

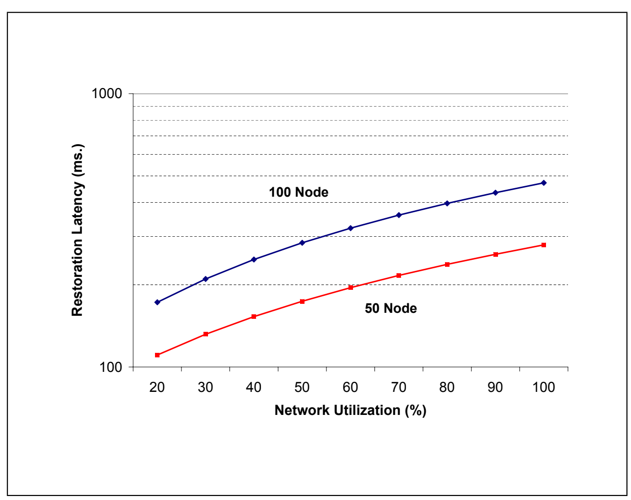

Fig. 14. Restoration latency as a function of the total network utilization γ

VI. Conclusion & Future Work

This paper presented a collection of approximation formulas that allow a network planner to quickly estimate a network size with limited inputs. In particular, it provides a set of equations that relate number of sites, average fiber connectivity, demand load and capacity for various protection architectures. These results can be used to easily and quickly estimate the amount of traffic that can be carried over a given network, or, inversely, given the traffic to be supported, to assess the characteristics of the topology required (in terms of number of nodes, connectivity). Finally, this analysis can be used to estimate the expected restoration performance without requiring any extensive simulation studies.

This work assumed uniform traffic demand but it can be extended to non-uniform traffic demand by deriving the path length h from certain charateristics of the traffic matrix. For example, for one-hop traffic demand, we would use $h = 1$, or the traffic demand could be routed along a shortest path and weighted according to the amount of point-to-point demand to derive h. Further extension of this work would index traffic demands and paths by source-destination node-pairs, thus allowing for arbitrary traffic patterns.

Acknowledgments

The authors would like to acknowledge their current or past colleagues George Ellinas, Chris Olszewski, and Sid Chaudhuri, as well as the reviewers for their valuable comments.

References

[1] J-F. Labourdette et al., "Fast Approximate Dimensioning and Performance Analysis of Mesh Optical Networks", 4^{th} International Workshop on the Design of Reliable Communication Networks (DRCN), Banf, Canada, October 2003.
[2] S. Korotky, "Network Global Expectation model: A Statistical Formalism for Quickly Quantifying Network Needs and Costs", Journal of Lightwave Technology, March 2004.
[3] J. Doucette and W. D. Grover, "Comparison of Mesh Protection and Restoration Schemes and the Dependency on Graph Connectivity', 3^{rd} International Workshop on the Design of Reliable communication Networks (DRCN), Budapest, Hungary, October 2001.
[4] G. Ellinas et al., "Routing and Restoration Architectures in Mesh Optical Networks", Optical Networks Magazine, January-February 2003.
[5] J-F. Labourdette, E. Bouillet, R. Ramamurthy, G. Ellinas, S. Chaudhuri and K. Bala, 'Routing Strategies for Capacity-Efficient and Fast-Restorable Mesh Optical Networks", Photonic Network Communications, Special Issue on Routing, Protection and Restoration Strategies and Algorithms for WDM Optical Networks, July-December 2002.
[6] Y. Liu, D. Tipper and P. Siripongwutikorn, "Approximating optimal spare capacity allocation by successive survivable routing", Proceedings of IEEE INFOCOM 2001, Anchorage, AL, April 2001.
[7] M. Kodialan and T.V. Lakshman, "Dynamic routing of bandwidth guaranteed tunnels with restoration", Proceeding of IEEE INFOCOM 2000, Tel Aviv, Israel, March 2000.
[8] G. Li et al., "Efficient Distributed Path Selection for Shared Restoration Connections", Proceeding of IEEE INFOCOM 2002, New York, NY, July 2002.
[9] E. Bouillet et al., "Stochastic Approaches to Compute Shared Mesh Restored Lightpaths in Optical Network Architectures", Proceeding of IEEE INFOCOM 2002, New York, NY, June 2002.
[10] C. Qiao et al., "Distributed Partial Information Management (DPIM) Schemes for Survivable Networks - Part I", Proceeding of IEEE INFOCOM 2002, New York, NY, July 2002.
[11] N. Biggs, "Algebraic Graph Theory", Cambridge Math., 1974.
[12] J. Doucette, W. Grover and T. Bach, "Bi-Criteria Studies of Mesh Network Restoration: Path Length vs. Capacity Tradeoffs", Proc. IEEE/OSA Optical Fiber Commun. Conf., Anaheim, CA, March 2001.
[13] E. Bouillet, J-F. Labourdette, R. Ramamurthy and S. Chaudhuri, "Enhanced Algorithm Cost Model to Control Tradeoffs in Provisioning Shared mesh Restored Lightpaths", Proc. IEEE/OSA Optical Fiber Commun. Conf., Anaheim, CA, March 2002.
[14] C. Qiao et al., "Novel Models for Efficient Shared-Path Protection", Proc. IEEE/OSA Optical Fiber Commun. Conf., Anaheim, CA, March 2002.
[15] S. Datta, S. Sengupta, S. Biswas and S. Datta, "Efficient Channel Reservation for Backup Paths in Optical Mesh networks", Proceeding of IEEE GLOBECOM 2001, San Antonio, TX, November 2001.
[16] R. Ramamurthy et al., "Limiting Sharing on Protection Channels in Mesh Optical Networks", Proc. IEEE/OSA Optical Fiber Commun. Conf., Atlanta, GA, March 2003.
[17] A. A. Akyamac aet al., "Optical Mesh Networks Modeling: Simulation and Analysis of Restoration Performance", Proc. NFOEC, Dallas, TX, September 2002.
[18] R. Bhandari, "Survivable Networks: Algorithms for Diverse Routing", Kluwer Academic Publishers, 1999.

VII. Appendix

A. Random Graphs

All the randomly generated graphs used in our experiments consist of rings traversed by chords connecting randomly selected pairs of nodes. Very often it is possible to embed such a ring on a real network, as demonstrated in Figure 15 with the ARPANET network.

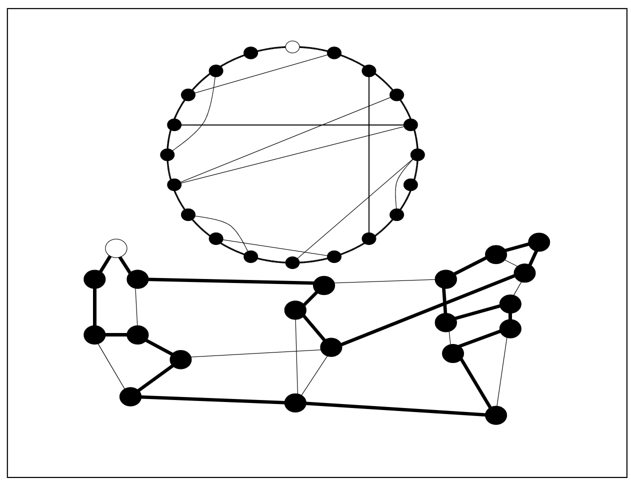

Fig. 15. Chordal ring (top) embedded on Arpanet (bottom).

B. Average Path length

In this appendix, we derive the average number of hops to reach two nodes in a graph of n nodes, m links, and degree $\delta = 2m/n$. The well-known Moore bound [18] gives the maximum number of nodes in a graph of diameter D and maximum degree $\delta_{max} > 2$:

$$n \leq 1+\delta_{max} \sum_{i=1}^{D} (\delta_{max}-1)^{i-1} = 1+\delta_{max} \frac{(\delta_{max}-1)^{D}-1}{\delta_{max}-2} \quad (17)$$

As illustrated in Figure 16, the Moore bound results from the construction of a tree whose root is the parent of δ_{max} vertices and each subsequent vertex is itself the parent of $\delta_{max} - 1$ vertices. The underlying idea is to pack as many vertices in D generations (hops) as is possible with respect to δ_{max}. The bound implies the existence of one such tree growing from every vertex and embedded in the graph, and is thus difficult to attain. It is nevertheless achievable for rings with odd number of vertices and for fully connected graphs. Reciprocally, given the number of nodes n, and degree δ, the lower bound D_{min} on the graph's diameter is easily obtained from Equation (17):

$$D_{min} \geq \frac{\ln[(n-1)\frac{\delta_{max}-2}{\delta_{max}}+1]}{\ln(\delta_{max}-1)} \tag{18}$$

Equations (17) and (18) can be combined to determine the lower bound of the average hop-length h:

$$\begin{aligned}(n-1)h \quad &\geq \quad \delta_{max}\sum_{i=1}^{D_{min}-1} i(\delta_{max}-1)^{i-1} \\ &+ \quad D_{min}\left(n-1-\delta_{max}\sum_{j=1}^{D_{min}-1}(\delta_{max}-1)^{j-1}\right)\end{aligned} \tag{19}$$

Equation (19) is a rather conservative lower bound of the average path length. To it we will prefer Equation (1) reproduced below and obtained by replacing δ_{max} by the average degree δ in Equation (18).

$$h \approx \frac{\ln[(n-1)\frac{\delta-2}{\delta}+1]}{\ln(\delta-1)}$$

The determination of the average path length using the Moore bound was obtained in part using a mathematical analysis, and in part empirically. First, assuming a uniform degree, we derived from the Moore bound an approximation of the average path length h. Even though the Moore bound tends to overestimate the size of the network, the model is remarkably accurate for large network sizes and large degrees δ. The reason for this is that the average path length is of the order of $\log_\delta(n)$, and it is therefore not very sensitive to variations in n. For instance, for $\delta = 3$, a one order of magnitude variation of n would correspond to only a $+$ or -1 variation of the average path length. We then refined the model so that when δ converge to $n-1$ (fully connected) or when n converges to infinity, the average path length converges to the expected value. We then validated empirically the model for the more general case when the degree is non-uniform.

Our experiments demonstrate that the equation is still valid for degree 3, but becomes inaccurate for degree less than 3. The reason for this is that as the degree converges to 2, the average path length converges to a linear function of the number of nodes, instead of following the logarithmic form of our model. In fact when the degree is 2, in which case the graph is a ring (it could also be a chain had we not assumed that the graph is 2-connected), the average path length is approximatively $n/4$, where the exact value depends on whether n is odd or even. A more accurate model that converges to the expected behavior for all known degrees and/or network sizes, may be the topic of future research.

Note that this is an approximation as one may want ot take a longer working path than the shortest path either (a) to be able to find a diverse backup path in the case of dedicated mesh protected lightpaths (and shared mesh restorable ones as well), or (b) to maximize sharing in the case of shared mesh restorable lightpaths. However, it is our experience that shortest path length gives a very good approximation of working path length for both dedicated mesh protected and shared mesh restorable lightpaths.

This approximation is exact for complete mesh networks (Equation (20)), and it converges to the proper limit for infinite size networks (Equation (21)). The approximation formula for h does not however converge to $n/4$ as should be expected but $n/2$ instead when the degree converges to 2 (i.e. when the topology becomes a ring).

$$\frac{\ln[(n-1)\frac{\delta-2}{\delta}+1]}{\ln(\delta-1)} \overset{n\to\delta+1}{\to} 1 \text{ (full mesh connectivity)} \tag{20}$$

$$\frac{\ln[(n-1)\frac{\delta-2}{\delta}+1]}{\ln(\delta-1)} \overset{n\to\infty}{\to} \frac{\ln n}{\ln(\delta-1)} \text{ (infinite network size)} \tag{21}$$

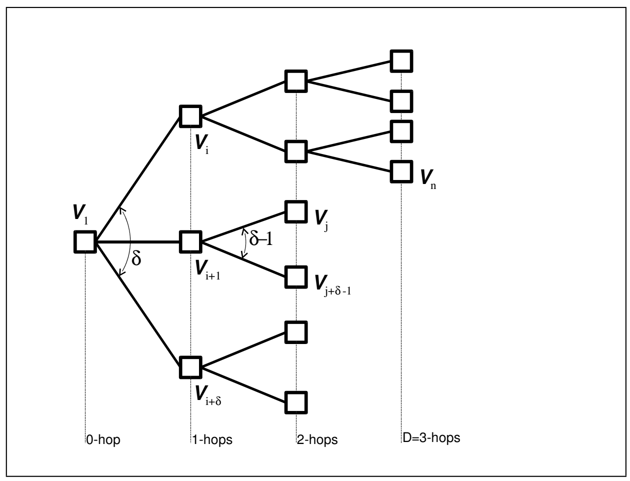

Fig. 16. Moore Tree. First vertex v_1 is connected to δ vertices one hop away from v_1. Subsequent vertices are connected to $\delta-1$ vertices one hop farther from v_1.

C. Ratio of Shared Protection to Working Capacity

In this appendix, we derive the average number of shared backup channels required on a link, as a function of the number of lightpaths L in the network, to guarantee restoration against single link failure. The restoration architecture used is that of pooling backup channels across all failures, that is not pre-assigning channels to particular backup paths. We consider a network with n nodes and m links. The average node degree is $\delta = 2m/n$. The average length of the primary path of a lightpath is h, given by Equation (1). The average length of the backup path of a lightpath is h', given by Equation (3). We want to determine the maximum number of times primary paths of lightpaths whose backup paths share a given link l traverse a common link. Under pooling of shared backup channels, this is the number of backup channels needed on link l to insure that all lightpaths that would be subject to the failure of a common link can be restored. The sharing factor is then simply the ratio of lightpaths whose backup paths traverse link l divided by the number of backup channels required.

First, we determine the number w of lightpaths whose backup paths traverse a given link l. A simple counting argument yields:

$$w = L\frac{h'}{m} \tag{22}$$

Second, we determine the maximum number of times the w primary paths whose backup paths traverse the given link l, traverse a common link. Let m' be the number of links traversed by those w primary paths. In order to determine m', note that each primary with its backup forms a ring that has an average diameter D equal to:

$$D = \frac{h+h'}{2} \tag{23}$$

Application of the diameter D and the degree $\delta = 2m/n$ to the Moore bound, gives us an estimate of m':

$$m' = min\left(m, \delta\left(\frac{(\delta-1)^{D-1}-1}{\delta-2}\right)\right) \tag{24}$$

The Moore bound is used to estimate the size of the sub-network comprising all the edges that are at most $D = (h+h')/2$ edges apart from the common link traversed by the w backup paths. Note however that with the exception of a few graphs, the Moore bound is known to overestimate the real size of the graph. The network has m edges, and the maximum number of edges available to the primary path is thus the minimum between m, and the size of the sub-network determined by way of the Moore bound.

To analyze the maximum interference among the w primary paths, we consider the following equivalent urn problem. Assume an urn of m' balls, then, assume that h balls (one for each link on a primary path) are picked from the urn (without replacement). The balls are identical and equiprobable (since we assume that the end points of the primary paths are uniformly distributed). This experiment is repeated independently a total of w times (one for each primary path).

We want to calculate the probability $P(x)$ that there is at least one link which is part of exactly x primary paths and that no other link is part of more than x primary paths among the w primary paths. This corresponds to the probability that, in w experiments, there is at least one ball that is selected exactly x times and there is no ball that is selected more than x times. Here, x can vary between 0 and w.

After w independent experiments, the probability $p(x)$ that a given ball is selected exactly x times, $0 \leq x \leq w$, is given by:

$$p(x) = \binom{w}{x} p^x (1-p)^{(w-x)} \tag{25}$$

where $p = \frac{h}{m'}$.

Now, we define two events A and B, where A is the event that no ball is selected x times and B is the event that no ball is selected more than x times during the w independent experiments. Then, the probability $P(x)$ that there is at least one ball that is selected exactly x times and there is no ball that is selected more than x times is given by:

$$P(x) = (1 - Pr(A))Pr(B) \tag{26}$$

where $Pr(A)$ and $Pr(B)$ are the probabilities of events A and B, respectively. $Pr(A)$ is readily found as follows:

$$Pr(A) = (1 - p(x))^{m'} \tag{27}$$

For a given ball, the probability $q(x)$ that the ball is selected less than or equal to x times is given by:

$$q(x) = \sum_{i=0}^{x} p(i) \tag{28}$$

Then, we can write that:

$$Pr(B) = q(x)^{m'} \tag{29}$$

Thus,

$$P(x) = q(x)^{m'}\left(1 - (1-p(x))^{m'}\right) \tag{30}$$

Finally, on average, the maximum number of times an arbitrary link is traversed by w primary paths is given by the following:

$$E\{P\} = \sum_{x=1}^{w} xP(x) = \sum_{x=1}^{w} xq(x)^{m'}\left(1-(1-p(x))^{m'}\right) \tag{31}$$

where w, $p(x)$ and $q(x)$ are calculated as shown above. The average sharing factor F of backup channels is obtained by dividing the average number of protected lightpaths w, whose backup paths traverse a given link, by the average number of backup channels needed, $E\{P\}$, on that link, yielding, as a function of the total number of lightpaths L:

$$F = \frac{w}{\sum_{x=1}^{w} xq(x)^{m'}\left(1-(1-p(x))^{m'}\right)} \tag{32}$$

with w, $p(x)$, $q(x)$ determined as shown above, and h and h' given in Equations (1) and (3).

PLACE
PHOTO
HERE

Jean-François Labourdette Dr. Jean-Francois Labourdette is regional architect manager with Verizon Enterprise Solutions Group. Before joining Verizon in March 2004, he was manager of system engineering at Tellium, responsible for all network element and network management system engineering activities. He had joined Tellium in October 2000 as manager of network routing and design, responsible for Tellium's routing architecture and algorithms, network dimensioning, and customer network design activities. Previously, he was a manager of multiservice globalization planning at AT&T, where he was responsible for network, service, and operations planning for AT&T international data services (frame relay, ATM, and MPLS IP VPN). Earlier at AT&T, he was a system engineer in the routing planning group, working on dynamic call routing for AT&T's switched network and facility routing and rearrangement for the AT&T transport network. Jean-Francois received his undergraduate degree in electrical engineering from l'Ecole Nationale Suprieure des Tlcommunications in Brest, France. He holds a Ph.D. from Columbia University, where he was the recipient of the 1991 Eliahu I. Jury Award for best dissertation. He is a senior member of the Institute of Electrical and Electronics Engineers (IEEE) and a member of the Optical Society of America. He has more than 40 publications in international conferences and journals (http://edas.info/S.cgi?author=labourdette).

PLACE
PHOTO
HERE

Eric Bouillet Dr. Eric Bouillet is currently at IBM where he works on data modeling and test data generation. Before joining IBM, Eric has worked at Tellium on the design of optical networks and optimization of lightpath provisioning and fault restoration algorithms, and before that in the Mathematical Sciences Research Center in Bell Labs/Lucent Technologies on routing and optimizations of telecommunication networks. Eric holds an M.S. and a Ph.D. in electrical engineering from Columbia University. He also holds a joint degree from l'Ecole Nationale Suprieure des Tlcommunications ENST Paris and EURECOM Sophia Antipolis. He has authored and co-authored several journal and conference papers and has a number of U.S. and international patents pending in the area of optical networking.

PLACE
PHOTO
HERE

Ramu Ramamurthy Ramu Ramamurthy received his B.Tech from IIT Chennai, and M.S. and Ph.D degrees from UC-Davis. He has worked on the design and development of routing algorithms and control plane protocols for optical networks most recently at Cisco, and prior to that at Ciena, Tellium and Telcordia.

PLACE
PHOTO
HERE

Ahmet Akyamac Ahmet A. Akyamac (M '00) received his B.S. degree in Electrical and Electronics Engineering from Bilkent University, Ankara, Turkey in 1993, and his M.S. and Ph.D. degrees in Electrical and Computer Engineering from North Carolina State University, Raleigh, NC in 1995 and 1999, respectively. He is currently a Technical Consultant at VPI Systems in Holmdel, NJ, working on data and transport network design tools. Previously, he was a Senior Network Architect at Tellium in Oceanport, NJ, where he worked in the fields of optical network design, routing and modeling, focusing on restoration performance and architectures, prior to which he was a Member of Technical Staff at Make Systems in Cary, NC (now part of Opnet Technologies), where he worked in the fields of IP and ATM network routing, MPLS traffic engineering and modeling. His research interests include the design, modeling and simulation, analysis and performance evaluation of telecommunication systems and networks. Ahmet is a member of the Eta Kappa Nu, Phi Kappa Phi. His email address is: akyamac@ieee.org.