Design of Reliable Communication Networks (DRCN) 2003, Banff, Alberta, Canada, October 19-22, 2003

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*
*Tellium, 2 Crescent Place, Oceanport, NJ 07757, USA
#Ciena Corp., 5965 Silver Creek Valley Road, San Jose, CA 95138, USA
E-mail: jlabourdette@tellium.com, ebouillet@tellium.com, sramamur@ciena.com, aakyamac@tellium.com

***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 requiring any 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 will 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 DWDM systems. The optical switches provide lightpath-based connectivity between client equipments such as routers, also shown in Figure 1.

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 [1].

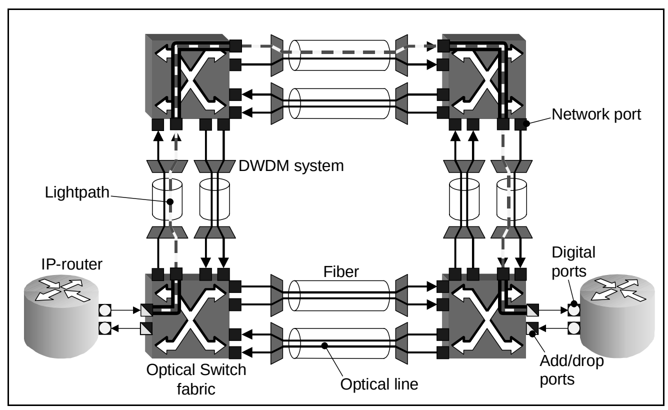

Figure 1: Optical mesh network

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 [2,3,4,5]. Path diversity guarantees that primary and backup lightpaths will not simultaneously succumb to a single failure. Unlike local span protection, back-up 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

[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 [1,3].

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

Dedicated or 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.

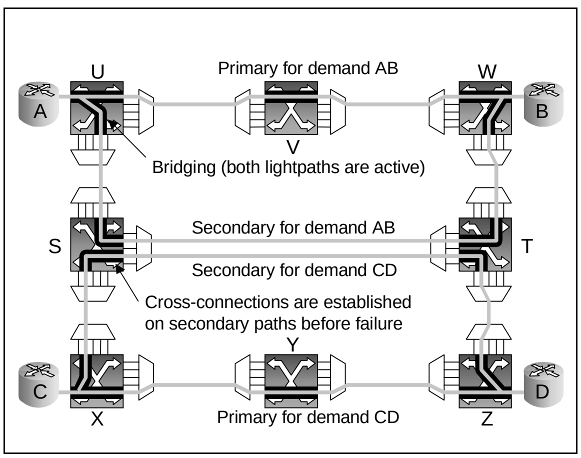

Figure 2. Dedicated mesh (1+1) protection

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. 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 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 [2,3] 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 back-up paths. See [12-14] for comparison of routing efficiency when only partial information is available.

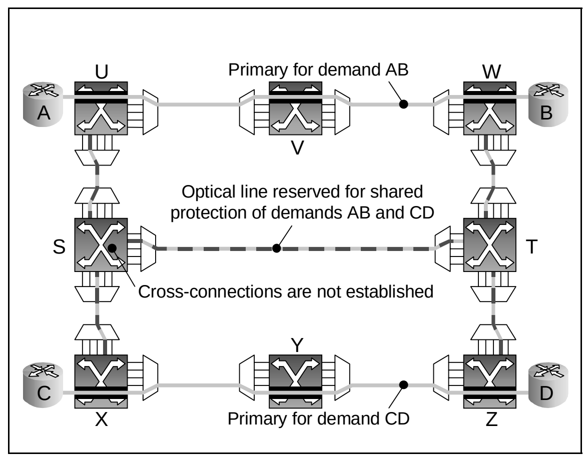

Figure 3. Shared mesh restoration: before failure

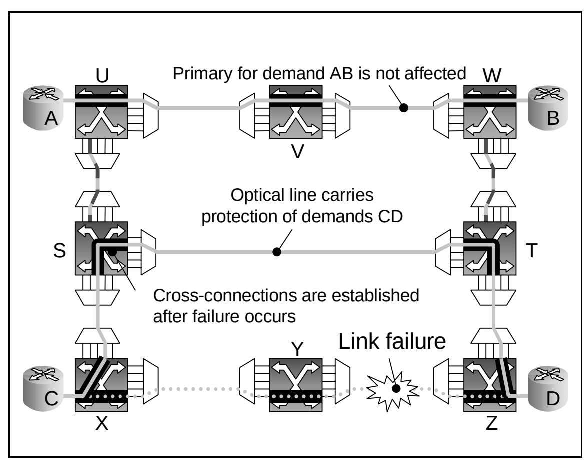

Figure 4. Shared mesh restoration: after failure

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) Optical Line Groups. 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 are default to one SRLG per link and one link per SRLG. We also assume no parallel links.

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 [8], we obtain (see Average Path Length in Appendix B):

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

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 back-up path in the case of a dedicated mesh protected lightpath, or (b) to maximize sharing in the case of a shared mesh restorable lightpath. 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 reapply our approximation of shortest path length on this new graph.

The computation of the average hop-length of the backup path ***a***-***z***, in the context of dedicated 1+1 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 ***b*** of ***a***, and removing node ***a*** and its δ adjacent edges, including edge (***a***,***b***).

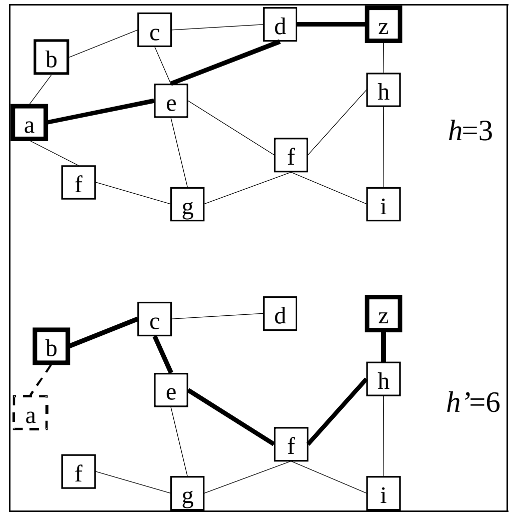

Figure 5. The backup path cannot traverse edges already used for its primary (*a*,*z*), and so these edges can be removed. Furthermore, the backup is at least as long as the primary and we assume no parallel edges, hence the backup is at least two hops long. This is represented by starting from any neighbor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

The purpose of the transformation is to determine the average degree δ' of this new graph (whose number of nodes is ***n***-1). Using this graph transformation approach, the new graph average degree is[2]:

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

where h is the average hop-length of the primary path as computed in (1). The average path length of the back-up path for dedicated mesh protected lightpath is then approximated by the length of the shortest path in the transformed graph[3] (using Eq (1)), plus one[4] as:

$$h' \approx \frac{\ln\left(\frac{(n-2)(\delta'-2)}{\delta-1}+1\right)}{\ln(\delta'-1)}+1 \qquad (3)$$

Figure 6 plots the approximation for h and h' against experimental path lengths computed in randomly generated networks (see Random Graphs in Appendix A) with average node degree 3.5, which is 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 exhibit the same behaviour.

In the case of shared mesh restoration we introduce ε, the cost of a shareable channel reserved for restoration [2,7,8,9]. The cost is actually the ratio of the cost of the same channel if it were not shareable (see [7,8,9] for details). The ratio ε ranges from 0 to 1.

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

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

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

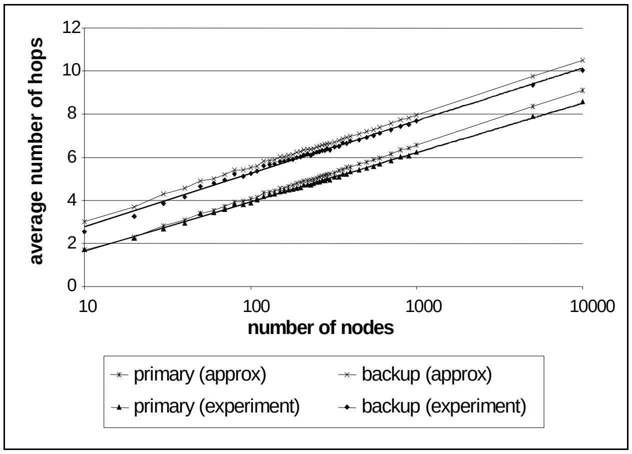

Figure 6: Comparison of path length approximations against path length computed in randomly generated networks of average degree 3.5

We express the average length of the back-up path for shared mesh restorable lightpath *h''* as:

$$h'' = h' + (1-\varepsilon) h_0 \qquad (4)$$

where h_o is determined experimentally (e.g., in the range of 0 to a few hops) or can sometimes be derived from the topology. For example, $h_o = 0$ 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 back-up path for shared mesh protected lightpath may be longer than the corresponding back-up length for dedicated mesh protected lightpath because sharing can be mined on different and therefore longer paths than the shortest path. Alternatively, the back-up 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 back-up 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 the case of shared mesh protection and dedicated mesh protection concerns 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:

$$\boldsymbol{R} = \frac{\textit{total number of protection channels}}{\textit{total number of working channels}}.$$

While the answer to this question is trivial in the case of 1+1 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 protection, the 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[5].

$$\boldsymbol{R_d} = \frac{h'}{h} \geq 1. \qquad (5)$$

In the case of shared mesh restoration, we cannot express 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 fill factor of a protection channel, that is the number of lightpaths whose protection path are using that channel. Then, the ratio R of protection to working capacity can be expressed as:

$$\boldsymbol{R_s} = \frac{h''}{h}\frac{1}{F} = \boldsymbol{R_d}\frac{1}{F} + (1-\varepsilon)\frac{h_0}{h}\frac{1}{F} \geq \frac{\boldsymbol{R_d}}{F}. \qquad (6)$$

Note that the formula for the ratio R does not assume that shared back-up channels are either pre-assigned to particular back-up paths or used as part of pool of shared channels [14]. The difference between the two approaches would be captured by different values of F. Note also that ***Rs***, contrary to ***Rd*** in the case of dedicated mesh (1+1) protection, is not independent of the number of lightapths as more lightpaths will provide for better sharing, thus increasing *F*, and therefore reducing ***Rs***. However, F, and ***Rs*** should become independent of the number of lightpaths when that number becomes large enough. Notice also that if F is fixed to one by capping the amount of sharing that is acceptable [13], ***Rs*** becomes the same as ***Rd*** as ε becomes one (no sharing possible). Finally, in the case where the back-up path is selected to be the same as for dedicated mesh protection, then $\boldsymbol{R_s} = \boldsymbol{R_d}/F$.

C. Sharing Analysis

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 ligpaths whose back-up 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 back-up channels required on the arbitrary link *l*, and F is simply the ratio of lightapths whose back-up path traverses *l* divided by the number of back-up channels required. The details of the analysis are given in Appendix C. Results comparing the value of ***Rs*** to this approximation are given here. Figure 7 and Figure 8 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.

[5] Assuming an un-capacitated network.

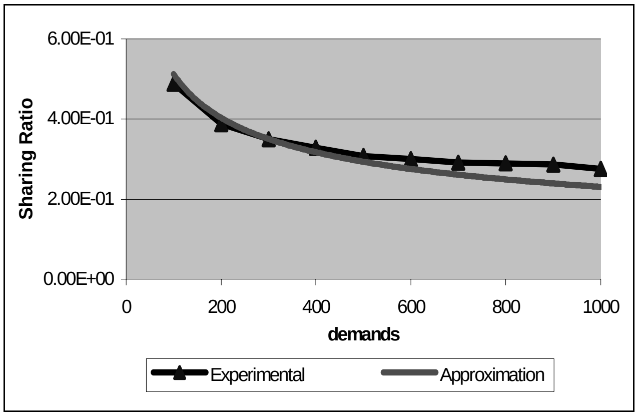

Figure 7 Sharing ratio, experimental versus approximation as demand increases on a 50 node, 75 link chordal ring network (degree = 3)

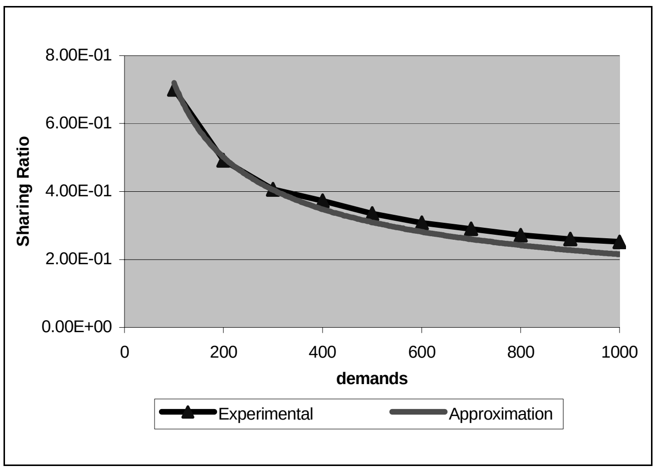

Figure 8 Sharing ratio, experimental versus approximation as demand increases on a 150 node, 300 link chordal ring network (degree = 4)

III. DIMENSIONING MESH OPTICAL NETWORKS

We introduce a node model with the following parameters:

S = size of switch

γ = switch loading, ratio of number of switch ports used to total number of ports (switch size) average over all switches

$\boldsymbol{R}$ = ratio of protection to working capacity

$\boldsymbol{T}$ = ratio of through working capacity to total working capacity

$\boldsymbol{Pr}$ = ratio of add/drop ports used for drop-side protection to add/drop ports used for service[6]

We model a node as shown in Figure 9. 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 a pool of A ports and a network-side port from a pool of W ports. The primary path of a through lightpath uses two ports from a pool of Th ports. A restoration channel on the back-up path of either a shared mesh restorable or a dedicated (1+1) mesh protected lightpath uses a port from a pool of P ports. The sizes of the pools verify the following conservation equations:

$$A + W + P + Th = \gamma S \quad (7)$$

$$A = W(1 + \boldsymbol{Pr}) \quad (8)$$

$$P = \boldsymbol{R}(W + Th) \quad (9)$$

$$Th = \boldsymbol{T}(W + Th) \quad (10)$$

Eq. (8) captures the fact that some of the drop-side ports are used for drop-side protection.

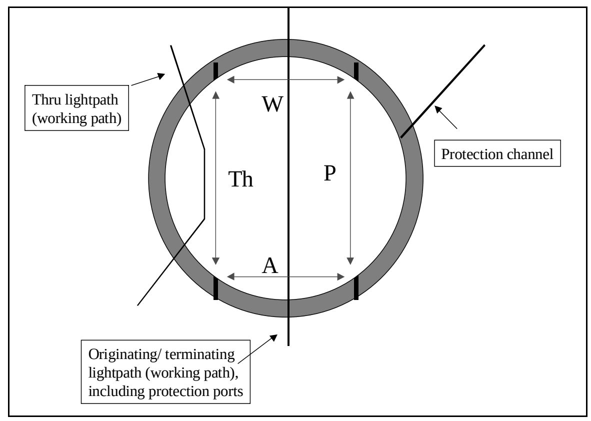

Figure 9: node model

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 Eqs (8), (9), and (10) into (7), we obtain after simplification:

$$A \approx \gamma S \frac{1 + \boldsymbol{Pr}}{1 + \boldsymbol{Pr} + (1 + \boldsymbol{R})h} \quad (11)$$

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

$$L_{network} = \frac{A * n}{2(1 + \boldsymbol{Pr})} \approx \frac{\gamma S}{2} \frac{n}{1 + \boldsymbol{Pr} + (1 + \boldsymbol{R})h} \quad (12)$$

Note that $L_{network} \propto o(n/h) = o(n/\ln n)$. From Eq. (12), and using $\delta=2m/n$, we can write the average number of lightpath per link in a maximally loaded mesh network as:

---

[6] 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. ***Pr*** is 0 if no drop-side protection is used; 1 if 1+1 drop-side protection is used; 1/N is 1:N drop-side protection is used.

$$L_{link} = \frac{L_{network} * h}{m} \approx \frac{\gamma S}{\delta} \frac{h}{1 + \boldsymbol{Pr} + (1 + \boldsymbol{R})h} \quad (13)$$

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

$$L_{node} = \frac{L_{network} * (1+h)}{n} \approx \frac{\gamma S}{2} \frac{(1+h)}{1 + \boldsymbol{Pr} + (1+\boldsymbol{R})h} \quad (14)$$

Note that $L_{node} \rightarrow \gamma S / [2(1+\boldsymbol{R})]$ when $n \longrightarrow \infty$, independent of *h*. The formula for the number of lightpaths in a maximally loaded network $L_{network}$ is a function of the protection ratio ***R***. ***R***, in the case of shared mesh protected lightpaths, depends in turn on the number of lightpaths in the network through the fill factor of shared back-up channels, F. Therefore, determining $L_{network}$ and $\boldsymbol{R_s}$ is equivalent to 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 ***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 protection. Also, operational networks are usually run around 70% utilisation. Finally, we assume here that ***Pr***=0. Figure 10 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 (***R***=0), shared-mesh protected demands (***R***=0.7), and dedicated mesh (1+1) protected demands (***R*** obtained from Eq. (5)). Two sets of curves are given for two different utilisation levels of switches of size 512, with γ=0.7 and 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 (and from that the geographical coverage of the network) given other characteristics such as average node degree and switch size.

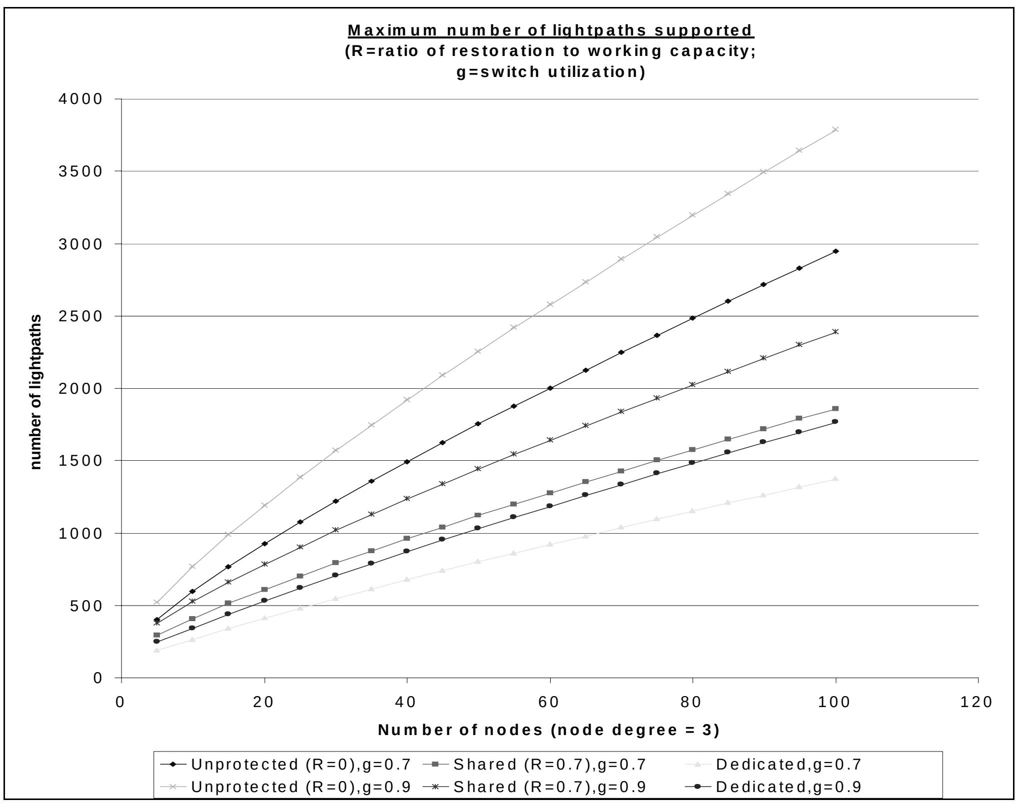

Figure 10: Max. number of lightpaths as a function of number of nodes for different switch utilisations g (switch size = 512)

In Table 1 and Table 2, we compare our approximate formulas with experimental results obtained for both a real and a random 50-node network. In the experiments, the networks were loaded with uniform demand, until one of the nodes reached capacity and five percents of new demands started to block[7]. The network characteristics and the usage achieved at the point where 5% of the uniform demand gets blocked are shown in Table 1, for both dedicated

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

and shared mesh routing. The resulting usage or network utilization ranges from 64% to 75%, which is characteristics of real networks. The resulting total number of demands or lightpaths carried in the experimental networks is shown in Table 2 ($L_{network}$), along with the average working and back-up path length, and the sharing ratio ***R***, for both dedicated and shared mesh routed demands[8].

Network	Type	n	m	Usage	Degree
Real	1+1	50	88	0.64	3.52
	Shared	50	88	0.68	3.52
Random	1+1	50	75	0.75	3.00
	Shared	50	75	0.75	3.00

TABLE 1: CHARACTERISTICS OF REAL AND RANDOM NETWORKS OF GIVEN SIZE (N, # OF NODES, M, # OF LINKS), DEMAND, USAGE[9], AND NODE DEGREE

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 2. For shared mesh routed demands[10], we considered two cases. In the first case, we computed the average working and back-up path lengths and the sharing ratio assuming uniform demand distribution. In the second case, we reflected 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 back-up path length, sharing ratio, and the maximum number of demands that can be supported. As can be seen in Table 2, in the random network case, the results are within a few percent for shared mesh routing to about 10% for dedicated mesh routing. In the real network case, the results are overestimated by 20-25% for dedicated mesh routing and shared mesh routing when using the experimental working path length. As expected, the approximations become less accurate as the network becomes less random. These results however are very encouraging in justifying the applicability of our approximation formulas, and further research into refining them.

V. RESTORATION TIME BEHAVIOUR

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 [12]. 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 [12]. 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:

$$T_r = T_0 + (L - 1)\,S \qquad (15)$$

where T_r is the final restoration latency for all lightpaths, T_0 is the restoration latency for the first lightpath and S is a parameter that represents the slope of the linear tail of restoration latency versus number of lightpaths restored. T_0 is obtained by assuming a single lightpath failure and analysing the restoration protocol (triggering, messaging, processing at switches 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 back-up paths are given by h and h'' in Equations (1) and (4). T_0 is a function of h and h'' as well as it depends on which link on the primary path fails. S is derived from experimental results.

[8] The experimental average path lengths and protection to working ratios are determined before blocking occurs.
[9] Usage refers to the number of ports used out of total numbers of ports per switch, averaged over all switches.
[10] Note that for shared mesh routed demands, the number of lightpaths is a function of the sharing ratio ***R***, which is itself a function of the number of lightpaths. The determination of R would thus normally require solving a fixed point equation. However, since this is an approximation, we computed R for 1000 lightpath demands, and assumed that its valued does not vary significantly around this point.

	Experim. h-work	Theoret. h-work	Experim. h-back	Theoret. h-back	Theoret. h-back using Experim. h-work	Experim. R	Theoret. R	Theoret. R using Experim. h-work	Experim. Lnetwork	Theoret. Lnetwork	Theoret. Lnetwork using Experim h-work
Real Net - 1+1	4.01	3.35	5.34	4.80	na	1.33	1.43	na	730	895	na
Real Net - Shared	4.01	3.35	5.34	4.80	4.83	0.40	0.16	0.15	1217	1768	1540
Random Net - 1+1	3.74	4.10	5.47	6.00	na	1.46	1.46	na	963	865	na
Random Net - Share	3.74	4.10	5.47	6.00	5.96	0.28	0.18	0.25	1629	1644	1692

TABLE 2: COMPARATIVE RESULTS FROM EXPERIMENTS AND APPROXIMATION FORMULAS FOR PATH LENGTH OF WORKING AND BACK-UP PATHS, SHARING RATIO *R*, AND TOTAL NUMBER OF DEMANDS SUPPORTED

Using the result for a single lightpath, the average number of lightpath per link (resp. 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 is *L* lightpaths failing for the same network configuration as above. The number *L* of lightpaths per link is obtained from Equation (13). 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.

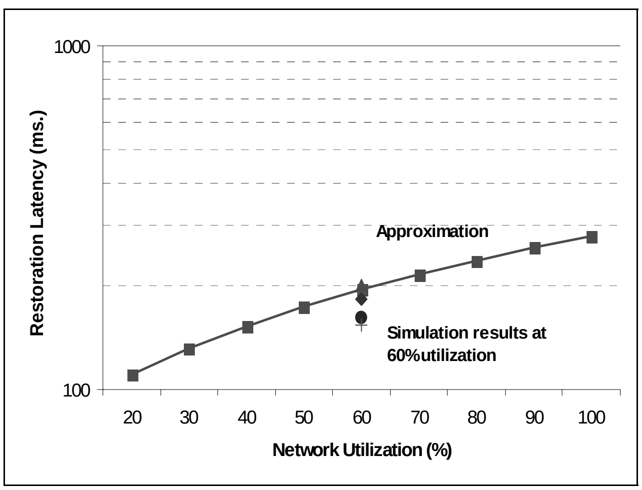

Figure 11: Analytical results vs. simulation results for hypothetical 50-node network.

Figure 11 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 11 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 [12]. In Figure 11, 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. This behaviour is 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 12 for two different size networks.

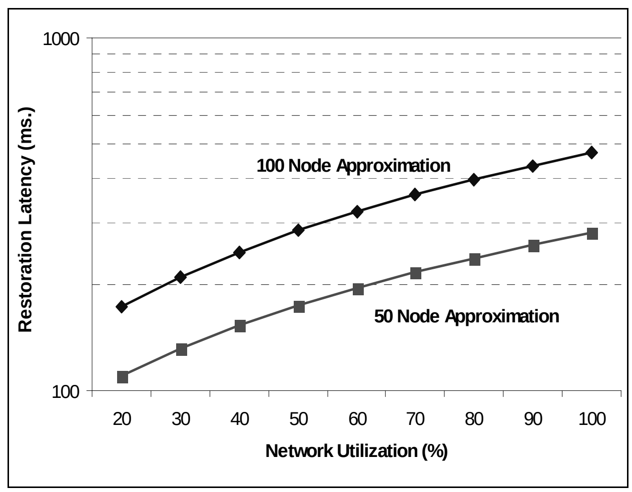

Figure 12: Restoration latency as a function of the total network utilization γ

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 restoration performance that can be expected without requiring any extensive simulation studies.

ACKNOWLEDGEMENT

The authors would like to acknowledge their colleagues George Ellinas, Chris Olszewski, and Sid Chaudhuri for many stimulating discussions.

REFERENCES

[1] J. Doucette and W. D. Grover, "Comparison of Mesh Protection and Restoration Schemes and the Dependency on Graph Connectivity", 3rd international Workshop on the Design of Reliable communication Networks (DRCN 2001), Budapest, Hungary, October 2001.

[2] J-F. Labourdette et al., "Routing Strategies for Capacity-Efficient and Fast-Restorable Mesh Optical Networks", special issue of Photonic Network Communications on "Routing, Protection, and Restoration Strategies and Algorithms for WDM Optical networks", July/Dec 2002 (invited paper).

[3] G. Ellinas et al., "Routing and Restoration Architectures in Mesh Optical Networks", Optical Networks Magazine, Jan/Feb 2003.

[4] Y. Liu, D. Tipper, and P. Siripongwutikorn, "Approximating optimal spare capacity allocation by successive survivable routing", Proc. IEEE INFOCOM, Anchorage, AL, April 24-28 2001.

[5] M. Kodialan and T.V. Lakshman, "Dynamic routing of bandwidth guaranteed tunnels with restoration", Proc. IEEE INFOCOM 2000, March 2000.

[6] B. Doshi et al., "Optical Network Design and Restoration", Bell-Labs Technical Journal, Jan-March 1999.

[7] R. Doverspike et al., "Fast Restoration in a Mesh Network of Optical Cross-Connects", OFC 1999, San Diego CA, Feb. 1999.

[8] N. Biggs; *Algebraic Graph Theory*, Cambridge Math. Library ISBN 0-521-45897-8 pbk, (1974, 1993 (2nd edition)).

[9] J. Doucette, W. Grover, and T. Bach: Bi-Criteria Studies of Mesh Network Restoration: Path Length vs. Capacity Tradeoffs, Proc. of OFC 2001, Anaheim, CA, March 2001.

[10] E. Bouillet, J-F. Labourdette, R. Ramamurthy, and S. Chaudhuri: Enhanced Algorithm Cost Model to Control Tradeoffs in Provisioning Shared mesh Restored Lightpaths, Proc. of OFC 2002, Anaheim, CA, March 2002.

[11] C. Qiao et al., "Novel Models for Efficient Shared-Path Protection", Proc. of OFC 2002, Anaheim, CA, March 2002.

[12] G. Li et al, "Efficient Distributed Path Selection for Shared Restoration Connections", IEEE-INFOCOM 2002, New York NY, Jul. 2002

[13] E. Bouillet et al., "Stochastic Approaches to Compute Shared Mesh Restored Lightpaths in Optical Network Architectures", Infocom 2002, June 23-27, NY, NY.

[14] C. Qiao et al, "Distributed Partial Information Management (DPIM) Schemes for Survivable Networks – Part I", IEEE-INFOCOM 2002, New York NY, Jul. 2002.

[15] A. A. Akyamac et al,. "Optical Mesh Networks Modeling: Simulation and Analysis of Restoration Performance", NFOEC 2002, sept. 2002, Dallas, TX.

[16] R. Ramamurthy et al., "Limiting Sharing on Protection Channels in Mesh Optical Networks", Proc. of OFC 2003, Atlanta, GA, March 2003.

[17] S. Datta et al., "Efficient Channel Reservation for Back-up Paths in Optical Mesh networks", Proc. of Globecom 2001, San-Antonio, TX, Nov. 2001.

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 13 with the ARPANET network.

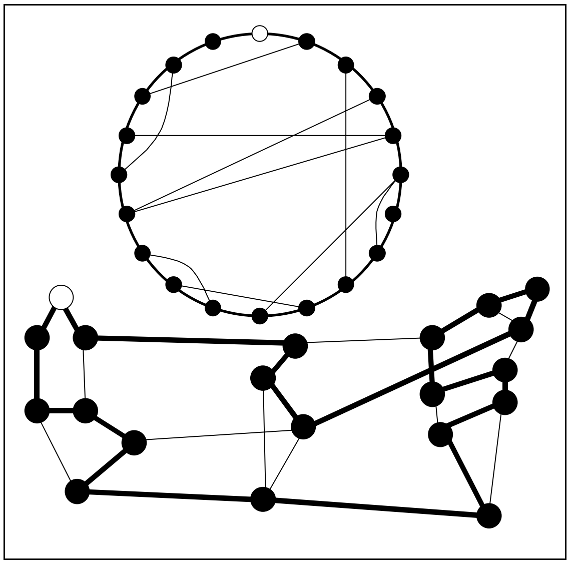

Figure 13: Chordal ring (top) embedded on Arpanet (bottom).

B. Average Path Length

In this appendix, we derive the average number of hops to reach two nodes in a graph of n nodes and degree δ. The well-known Moore bound [8] gives the maximum number of nodes in a graph of diameter D and maximum degree $\delta_{max} > 2$:

$$n \le 1 + \delta_{\max} \sum_{h=1}^{D} (\delta_{\max} - 1)^{h-1} = 1 + \delta_{\max} \frac{(\delta_{\max} - 1)^{D} - 1}{\delta_{\max} - 2} \quad (1)$$

As illustrated in Figure 14, 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 δ_{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 (1):

$$D_{\min} \ge \frac{\ln\left[1 + (n-1)\frac{\delta_{\max} - 2}{\delta_{\max}}\right]}{\ln(\delta_{\max} - 1)} \quad (2)$$

Equations (1) and (2) can be combined to determine the lower bound of the average hop-length:

$$h \ge \delta_{\max} \sum_{i=1}^{D_{\min}-1} i(\delta_{\max} - 1)^{i-1} + D_{\min}\left[n - 1 - \delta_{\max}\sum_{j=1}^{D_{\min}} (\delta_{\max} - 1)^{j-1}\right] \quad (3)$$

Equation (3) is a rather conservative lower bound of the average path length. Instead we replace δ_{max} by the average degree δ in equation (2) and obtain equation (4):

$$h \approx \frac{\ln\left[\frac{(n-1)(\delta-2)}{\delta} + 1\right]}{\ln(\delta - 1)} \quad (4)$$

This equation is exact for full mesh networks (5), it converges to proper limits for infinite size networks (6), and our experiments indicate that it gives a fair approximation of the average path length. However it is not appropriate for networks with average degrees lower than 3. In fact we can show that as the average degree tends to 2, which is equivalent to a ring, this equation converges to an average path length that is twice the average path length of the equivalent ring.

Convergence to full mesh connected:

$$\frac{\ln\left[\frac{(n-1)(\delta-2)}{\delta} + 1\right]}{\ln(\delta - 1)} \xrightarrow{n \to \delta+1} 1 \quad (5)$$

Convergence to infinity:

$$\frac{\ln\left[\frac{(n-1)(\delta-2)}{\delta} + 1\right]}{\ln(\delta - 1)} \xrightarrow{n \to \infty} \frac{\ln n}{\ln(\delta - 1)} \quad (6)$$

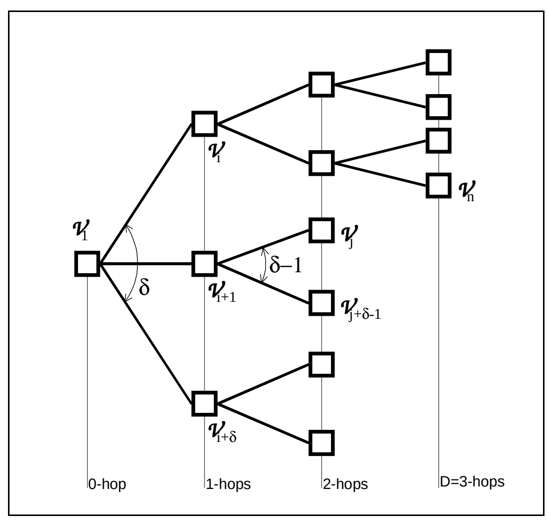

Figure 14. Moore Tree. First vertex v_1 is connected to δ vertices one hop away from v_1. Subsequent vertices are connected to δ-1 vertices one hop farther from v_1.

C. Ratio of Shared Protection To Working Capacity

In this appendix, we derive the average number of shared back-up 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 back-up channels across all failures, that is not pre-assigning channels to particular back-up 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 back-up path of a lightpath is h', given by Equation (3). We thus want to determine the average number of times an arbitrary link is traversed by all the primary paths whose back-up paths share a given common link. Under pooling of shared back-up channels, this is the average number of back-up channels needed on that back-up link to insure that all lightpaths that would be subject to the failure of this arbitrary link would restore on their back-up path across the back-up link considered. The number, in turn, tells us the average fill factor F of back-up channels by dividing the number of back-up channels by the number of protected lightpaths.

To analyze this problem, we map it to an equivalent urn problem. In the equivalent urn problem, the number of balls in the urn is the number of links in the network m. The number of balls picked is the number of links per path (h for a primary path, h' for a back-up path). And the number of experiments is the number of paths (number of back-up paths n_b which share a given link, or number of primary paths n_p whose back-up paths share a common link).

First, we need to determine the number n_b of back-up paths traversing a given link. To analyze this problem, we map it to an equivalent urn problem. In the equivalent urn problem, the number of balls in the urn is the number of links in the network m; the number of balls picked is the number of links per back-up path h'; and the number of experiments is the number of lightpaths L.

The equivalent urn problem is thus the following. Assume an urn of m balls, then, assume that h' balls are picked from the urn (without replacement). The balls are identical and equiprobable (since we assume that the end points of the lightpaths are uniformly distributed). This experiment is repeated independently a total of L times.

In one experiment in which h' balls are selected from the urn, the probability p' of selecting any given ball is found as follows:

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

Then, after L independent experiments, the probability $p'(x)$ that a given ball is selected exactly x times, $0 \le x \le L$, is given by:

$$p'(x) = \binom{L}{x} p'^{x} (1-p')^{(L-x)}$$

Therefore, a given link is selected as part of the back-up path of L lightpaths, on average:

$$n_b = \sum_{x=1}^{L} x \cdot p'(x) = L\frac{h'}{m}$$

Second, we determine the maximum number of times $n_p=n_b$ primary paths whose backup paths traverse a common link, traverse a given link**.** To analyse this problem, we again map it to an equivalent urn problem. In the equivalent urn problem, the number of balls in the urn is the average number of links in the network m' traversed by primary paths whose backup share a common link. The number of balls picked is the number of links per primary path h; and the number of experiments is the number of primary paths n_p whose back-up paths share a common link.

The equivalent urn problem is thus the following. Assume an urn of m' balls, then, assume that h balls 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 n_p times.

For n_p primary paths, we want calculate the probability 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. This would correspond to the probability that, for n_p 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 n_p. The expected value of the above probability is equal to the average number of times an arbitrary link is traversed by n_p primary paths.

In one experiment in which h balls are selected from the urn, the probability p of selecting any given ball is found as follows:

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

In order to determine m', note that each primary with its backup forms a ring that has an average diameter D equal to:

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

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

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

Then, after n_p independent experiments, the probability *p(x)* that a given ball is selected exactly x times, $0 \leq x \leq n_p$, is given by:

$$p(x) = \binom{n_p}{x} p^x (1-p)^{(n_p - x)}$$

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 n_p 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:

$$p(x) = [1 - \Pr(A)]\Pr(B)$$

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

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

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

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

Then, we find that:

$$\Pr(B) = [q(x)]^{m'}$$

Thus,

$$\mathrm{P}(x) = [q(x)]^{m'} \cdot \left[1 - (1 - p(x))^{m'}\right]$$

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

$$\mathrm{E}\{P\} = \sum_{x=1}^{n_p} x \cdot [q(x)]^{m'} \cdot \left[1 - (1 - p(x))^{m'}\right]$$

where n_p, *p(x)* and *q(x)* are calculated as shown above.

The average fill factor F of back-up channels is obtained by dividing the average number of protected lightpaths n_p, whose back-up path traverses a given link, by the average number of back-up channels needed, E{P}, on that link, yielding, as a function of the total number of lightpaths L:

$$F = \frac{n_p}{\sum_{x=1}^{n_p} x \cdot [q(x)]^{m'} \cdot \left[1 - (1 - p(x))^{m'}\right]}$$

With n_p, *p(x), q(x)* determined as shown above, and h and h' given in Equations (1) and (3).