

# Stochastic Approaches to Compute Shared Mesh Restored Lightpaths in Optical Network Architectures

Eric Bouillet, Jean-François Labourdette, Georgios Ellinas, Ramu Ramamurthy, Sid Chaudhuri

Tellium, Inc

185 Route 36, Building E, West Long Branch 07764

**Abstract:** In this paper we assess the benefits of using statistical techniques to ascertain the shareability of protection channels when computing shared mesh restored lightpaths. Current deterministic approaches require a detailed level of information proportional to the number of active lightpaths, and do not scale well as traffic demands and network grow. With the proposed approach we show that less information, independent of the amount of traffic demand, is sufficient to determine the shareability of protection channels with remarkable accuracy. Experiments also demonstrate that our approach yields faster computation times with no significant penalty in terms of capacity usage.

**Index terms:** Optical networks, Optical Switching, Mesh Protection, Stochastic Algorithm, Performance Analysis.

## A. INTRODUCTION

Dense Wavelength Division Multiplexed (DWDM) mesh network infrastructures that route optical connections (lightpaths) using intelligent optical cross-connects (OXC) have emerged as the technology of choice to implement next generation data[1]. In these architectures a single piece of equipment is capable of transferring tens of terabits per second. This equipment is continuously exposed to multifarious risks of breakdown, either due to human-induced mishaps, or to equipment malfunctions. In order to guarantee service persistence in such circumstances it is common for a carrier to reserve spare bandwidth on alternate paths, so that a service affected by a failure along its primary lightpath can be rapidly restored using the reserved bandwidth. Among the possible schemes for provisioning backup paths, dedicated protection and mesh restoration seem to be the most appropriate approaches in the context of DWDM networks [2][3][4].

In dedicated protection, the lightpath provisioning algorithm computes and establishes simultaneously the primaries and their protection paths. During normal operation mode, both paths carry the optical signal and the egress selects the best copy of the two. The concept of Shared Risk Group (SRG) was introduced to select the paths so that they will not be affected by a single failure[5][6]. An SRG expresses the relationship that associates optical lines (or possibly other optical components) with a single failure. It may consist of all the optical lines in a single fiber, or the optical lines through all the fibers wrapped in the same cable, or all the optical lines traversing the same conduit. Since a fiber can traverse several conduits, an optical line may belong to several SRGs. It suffices that a primary and its backup path are SRG disjoint to ensure that at least one path survives any single failure.

As in dedicated protection, shared mesh restored paths are predefined, except that the cross-connections along the paths are not created until a failure occurs. During normal operation modes the spare optical lines 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 line 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 mutually Shared Risk Group (SRG) disjoint, so that a failure does not interrupt both primary paths. If that happened, there would be contention for the reserved channel and at most only one of the two lightpaths would be successfully restored. Two lightpaths, or their protection, are said to be mutually compatible, if they are not affected by the same failure. If not, they are incompatible. Figure 1 (for normal mode) and Figure 2 (for restoration mode) illustrate an example of mesh restoration. The network consists of four client nodes (A to D) and two demands (AB and CD) accommodated across an eight node optical network (S to Z). The dashed lines represent channels reserved for protection. Using the routing of Figure 1, demands AB and CD are compatible with respect to SRG-failures and thus their protection share a single optical line in link S-T, one less than would be required in dedicated protection. Upon failure as depicted in Figure 2, the egress and egress nodes of the disconnected paths (X and Z) emit a request to the switches along the protection paths (S and T) to establish the cross-connections for that path. Once the cross-connections are established, each ingress and egress node restores the connection to the new path. This architecture requires fewer resources than in dedicated protection, but the restoration involves more processing to signal and establish the cross-connections along the restoration path.

They are two policies to assign reserved channels to restoration paths[7]: A failure dependent strategy assigns the reserved channels in real time after failure occurrence on a first-come first-serve basis depending on availability. A proper spare channel-provisioning scheme reserves enough channels so that all lightpaths can be restored for every type of failure. A failure independent strategy assigns the reserved channels at the time of lightpath provisioning prior to failure occurrence. One advantage of the failure independent strategy over the failure dependent is that during lightpath restoration the switches on the protection paths immediately and individually cross-connect to predetermined channels, based uniquely on the identifier of the lightpath being restored. For the scope of this study the failure independent strategy is used.

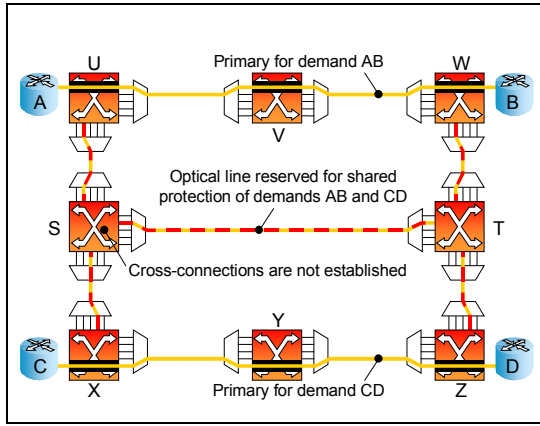


Figure 1. Mesh Restoration, before failure




Figure 2. Mesh Restoration, upon failure of fiber Y-Z

Now, consider the online problem of provisioning a mesh-restored lightpath using a centralized Route Computation Module (RCM.) Since this problem is proved to be NP-complete if minimization of the total capacity usage (working plus protection) is sought, a possible approach is to enumerate a list of  $K$  minimum cost primary paths and for every one of them compute the corresponding minimum cost restoration path. The RCM then returns the pair of paths with the lowest aggregated cost. The cost of a pair is the cost of the channels along both paths, excluding the cost of (preexisting) shareable reserved channels along the backup path. Given a primary path, we compute the minimum cost backup-path by: (i) setting the cost of the fibers (SRGs) traversed by the primary path to  $\infty$ , (ii) setting the cost of fibers with shareable channels to a constant  $\epsilon \ll 1$ , (iii) run a shortest path algorithm using the modified fiber cost metric. Step (i) and (ii) respectively ensure that primary and backup paths are SRG-diverse, and that the minimum cost backup path is found using shareable reserved channels whenever possible. In the following we are interested in step (ii), which consists of identifying shareable reserved channels. We show in particular that the time-complexity of this operation, if deterministic, is proportional to the total number of reserved channels, and thus does not scale well when the number of lightpaths established in the network becomes large. We then present a stochastic approach to

execute this operation with a certain probability of accuracy. We show that by trading a deterministic TRUE or FALSE statement for a PERHAPS statement with a measurable likelihood that PERHAPS is TRUE, the operation can be made independent of the number of reserved channels. The benefits of this substitution are: (1) reduction of the path computation time and (2) reduction of the amount of information necessary to compute the paths, with no penalty or small penalty in terms of capacity efficiency.

Section B develops the complexity of the deterministic approach to identify shareable channels. Section C describes the details of the stochastic approach. Section D describes an algorithm to compute mesh restored paths using the stochastic approach. Section E compares the results for realistic topologies using both stochastic and deterministic based algorithms, and Section F concludes this paper.

## B. TIME COMPLEXITY OF DETERMINISTIC APPROACH

In what follows, a list of SRGs protected by a given reserved channel consists of all distinct SRGs traversed by all the primary lightpaths whose respective protection paths are assigned the reserved channel. Thus a reserved channel can be reused to protect a primary path if no SRG traversed by the primary path appears in the list of SRGs already protected by the channel. We denote by  $h$  the average primary path length expressed in number of edges,  $m$  the number of edges in the network, and  $x$  the number of reserved channels per edge. We also assume the typical case where the average number of protected SRGs per reserved channel is on the order of  $O(m)$ . In order to identify shareable reserved channels in the network the algorithm must verify for each reserved channel in each edge that the list of SRGs protected by the channel does not intersect with the list of SRGs traversed by a primary path  $p$ . Therefore, the complexity of identifying all the edges with shareable reserved channels in the network is  $O(hxm^2)$ . This complexity is simplified to  $O(hxm)$  if each reserved channel maintains a fixed length array in which each entry indicates whether an SRG is used or not. The number of reserved channels per edge is a function of  $g$ , the number of lightpaths in the network, and can be approximated by  $x=O(gh'/m)$ , where  $h'$  is the average length of a backup path ( $h \geq h'$ .) Substituting  $x$ , the complexity of this operation is  $O(ghh')$ . Our primary concern here is the dependence of this complexity on the number of lightpaths established in the network. We thus propose to substitute this time consuming deterministic approach for a stochastic approach whose complexity remains constant with respect to the number of lightpaths.

## C. STOCHASTIC APPROACH

In what follows, we assume that the RCM has an up-to-date knowledge of the state of the network. In particular this information must include for each fiber (i) the number of reserved channels in the fiber, and (ii) for every SRG in the

network the number of reserved channels in the fiber that the SRG appears in.

We now describe the technique used to quickly compute the probability that a reserved channel is shareable with respect to a given primary path and based on the information available to the RCM. We will first introduce a simple combinatorial problem, solve it, and show the analogy between this problem and the one that we are interested in.

### 1st. A Simple Problem of Combinations

The problem: We are given  $N$  bags tagged from 1 to  $N$ , filled with marbles. Bag  $j$  ( $j \in \{1, \dots, N\}$ ) contains  $n_j$  marbles. All marbles in any given bag have the same color, but marbles in different bags have different colors, so that there is a one-to-one mapping between bags and colors. We are also given  $M$  bins tagged from 1 to  $M$ . We assume for the moment that the bins have infinite capacity. Next we empty the bags in the bins, so that not two marbles of the same bag (or same color) fall in the same bin. The questions are:

1. How many differentiable combinations of marbles to bins are possible? Assume that we cannot distinguish between two marbles of the same color.
2. Out of all combinations computed in (1), how many of them have empty bins left?
3. What is the probability that at least one bin is empty? Assume that probability of occurrence is the same for all combinations computed in (1) and (2), i.e. the marbles are uniformly distributed in the bins.




Figure 3. Bins and Bags problem

### The answers:

In the following we use  $C(p,q)=q!/p!(q-p)!$  to denote the unordered combinations of  $p$  out of  $q$  elements.

1. First note that a solution exists iff  $M \geq n_j \forall j \in \{1, \dots, N\}$ . The marble arrangement of each bag into the  $M$  bins is not conditional to other bags' arrangements. For each bag  $j$  there are  $C(n_j, M)$  possible ways to arrange the  $n_j$  marbles into an ordered set of  $M$  bins. There are thus  $Q = \prod_{j \in \{1, \dots, N\}} C(n_j, M)$  arrangements.
2. First note that if  $\exists j \in \{1, \dots, N\} / M = n_j$  then there is no such combination, and the answer is  $D=0$ . Note also that if  $M > \sum_j n_j$  then we cannot fill all the bins and the answer is  $D=Q$ . If  $M = 1 + \max_j \{n_j\}$ , the problem is equivalent to (i), except that now arrangements are confined to  $M-1$  bins, the last bin being left empty.

They are thus  $D = C(M-1, M) \prod_j C(n_j, M-1)$  possible arrangements. Otherwise, if  $M > 1 + \max_j \{n_j\}$ ,  $D = r(M-1)$ , where  $r(k)$  is the recursion over integers in  $k \in \{\max_j \{n_j\}, \dots, M-1\}$  such that  $r(\max_j \{n_j\} - 1) = 0$ , and  $r(k) = \{C(k, M) \prod_j C(n_j, k)\} - r(k-1)$ . (See Appendix in Section G)

3.  $P = D/Q$ , although the computation of  $D$  and  $Q$  may be tedious. We thus show here a mean to approximate this probability. First observe that the probability that at least one bin is empty is complement to the probability that all bins are non-empty. And the probability that a bin is non-empty is the complement to the probability  $p$  that this bin is empty. Although  $p$  is conditional to the probability of other bins being empty we assume that it is independent and identical for all bins. Therefore, given a bin the probability  $p$  that the bin is empty is the product of independent probabilities that all marbles of each bag are in other bins, that is  $p = \prod_j (1 - n_j/M)$ . Based on our observations and assumption, the probability that at least one bin is empty is  $P = 1 - (1-p)^M = 1 - (1 - \prod_j (1 - n_j/M))^M$ . The complexity of computing  $P$  (or its complement  $1-P$ ) involves computing  $N$  products and an  $M^{\text{th}}$  power. It is realizable in  $O(N + \log M) \approx O(N)$  time.

### 2nd. Analogy with SRG arrangement into a set of reserved channels

Assume that the  $M$  bins of the problem presented in 1st are the reserved channels in a given fiber. And assume that the  $N$  bags represent a list of  $N$  SRGs traversed by the primary path for which a reserved channel is sought. The  $n_j$  marbles denote the number of times each SRG of the list is protected (through pre-established paths) by the reserved channel set. Evidently the same SRG cannot be protected multiple times by the same reserved channel otherwise contention would exist through their respective primaries if the SRG fail. This restriction is expressed in the problem formulation by the fact that two identical marbles (same SRG) cannot fall into the same bin (reserved channel). Thus, the problem above deals with computing the probability that there is at least one shareable reserved channel, i.e. a reserved channel that does not contain any of the  $N$  SRGs. We have shown that this probability is approximated in  $O(N)$  time, where  $N$  is the number of SRGs on the primary path. Typically  $N$  is the average path length  $h$ . Therefore, the complexity of identifying all the edges with shareable reserved channels in the network is  $O(hm)$ . This complexity is to be compared with  $O(ghh')$  of the deterministic approach. Remember that in the computation of these probabilities we have made two simplifying assumptions: (i) the probability of a reserved channel being shareable is pairwise independent of other reserved channels, and (ii) SRGs are uniformly distributed across reserved channels. The effect of the first assumption is easy to quantify by way of simulations (see Section E.) The effect of the second assumption on the other hand is subtler because it depends on the policy used for allocating reserved channels. For instance a "First Fit" or "Max Fit" policy tends to pack

(protect) more SRGs in some reserved channels than others within the same fiber. As it turns out, a First Fit policy increases the probability that a reserved channel is available compared to a uniformly randomized allocation.

#### D. ALGORITHM

We describe here in details an algorithm that implements the stochastic approach, and compare it with the equivalent deterministic algorithm.

Given: a topology represented as a graph  $G(V,E)$  where vertices represent optical cross-connects (OXC) and edges represent fiber strands between OXCs. A network state database, that indicates for each edge (fiber) the number of channels available, the number of reserved channels, and the number of times each SRG in the network is protected by a reserved channel in that edge. The latter information is stored into an array. The array's indices correspond to SRGs and each entry in the array counts the number of reserved channel cross-connections that would occur in the edge if the corresponding SRG fails.

Input: a pair of nodes A-Z

Output: a pair of bi-directional lightpaths from A to Z, primary and secondary with minimum cost, excluding restoration channels that are shared with pre-established backup paths.

##### Algorithm:

1. Compute k-shortest paths. Sort the paths by length and denominate them  $w_1$  to  $w_k$ .
2. Set  $S = \emptyset$
3. For each shortest path  $w_i$ , do:
  - (i) To each edge that shares a SRG with  $w_i$  or has neither available channel nor reserved channel, assign infinite weight
  - (ii) For each edge without a reserved channel, set weight to cost of edge
  - (iii) For each edge with reserved channel, set weight to cost of edge times the probability that no reserved channel is shareable (by way of the approach presented earlier in section C of this document.)
  - (iv) Compute the shortest path  $s_i$  using the metric defined in parts (i) to (iii), and set  $S \leftarrow S \cup \{w_i, s_i\}$
4. Select the minimum cost path pair  $\{w_k, s_k\} \in S$ .
5. Use signaling to set-up the path. Reserved channels are assigned locally, using some local optimization algorithm for instance. If no path can be found in 4 return NO\_PATH.

The algorithm is self-explanatory. It differs from the deterministic algorithm only in step 3(iii). In the deterministic algorithm the weight of an edge is set to the edge cost times  $\epsilon \ll 1$  if it contains a shareable reserved channel and edge cost if it does not. In the stochastic algorithm this weight is replaced by the cost of the edge times the probability that no reserved channel is shareable in


the edge. Note that the deterministic approach requires additional information to compute the routes. In particular it needs to know whether each SRG is protected or not for every reserved channel. Whereas in the stochastic approach, only the number of times an SRG is protected in every edge by any reserved channels of that edge needs to be known. Finally, note that in step 5 we separated lightpath provisioning from routing, and channel assignment is performed in a distributed way after the lightpaths are selected by the RCM. The objective of the RCM is to compute the paths so that sharing is maximized during channel assignment. Even though an edge may be erroneously tagged as having a shareable channel during path computation, the channel assignment procedure during path setup will guarantee that they are no sharing violation. In order to guarantee this, the scheme used for channel assignment requires the same information as for the deterministic approach, however this information can be distributed across the nodes in the network: it suffices that each node maintains a local database of all the reserved channels terminating into it.

For the experiments presented in the next section we used both the deterministic and stochastic implementations of the algorithm. A great care was taken in optimizing the deterministic implementation for speed. The stochastic code was then derived from the deterministic code by modifying step 3(iii) as described above.


#### E. EXPERIMENTS AND RESULTS

##### *3rd. Accuracy and Distributions of Probability Functions*

In the following we first measure the quality of the estimated probability that an edge contains a shareable reserved channel given the information on the number of times each SRG traversed by the primary path is restored in that edge. The experiment consists of simulating a large number of arbitrary instances of the problem presented in section 1st. For each instance of the problem, we simulate several millions of random arrangements, and compute the ratio of combinations with available reserved channels to the total number of combinations (i.e. estimate 1st.3) We then compare the difference between each experimental probability and the corresponding exact and approximate probabilities obtained by computation. The results are shown in Figure 4 and 5. Figure 4 demonstrates the error distribution of the estimate probabilities minus experimental probabilities obtained over the range of problem instances. We observe that the estimate probability has a tendency to underestimate the experimental probability, but it is accurate within 0.05 for 85% of the time, which is quite remarkable given the simplicity of the computation. In comparison, the simulation exhibits an accuracy within 0.01 of the exact probability, and a closer look even indicates that 70% of the time the difference is within  $5 \times 10^{-4}$  (Figure 5).



**Figure 4.** Error distribution of estimate sharing probability versus simulation



**Figure 5.** Error distribution of exact sharing probability versus simulation

#### 4th. Comparison of Deterministic versus Stochastic Weight Functions on Real Networks

In the next set of experiments we consider two scenarios inspired from real life networks. NetA is a 100- node, 137-edge network, with one unit of demand between every pair of node (4950 demands). NetB is 220-node, 300-edge network, also with one unit of demand between every pair of node (24090 demands.) For the sake of simplicity we assume here that every edge costs one unit of currency and corresponds to one SRG (i.e. one SRG per edge and one edge per SRG). We then route the demands on each network using the deterministic and the stochastic algorithms. We are interested here in the processing time to complete each algorithm, and the quality of the solutions expressed in total number of channels required (used for primaries and reserved for backups.) Table 1 summarizes the results. For NetA (resp. NetB) we observe that the stochastic approach is 6.78 time faster (resp. 19.7 time faster) than the deterministic approach while the penalty is only 2% (resp. 3%) more capacity. Also important is the amount of information the RCM needs to compute the routes. The stochastic based RCM only require one array per edge, where each entry indicates the number of times the SRG is

protected in the edge by any reserved channel. For instance in the NetB problem, they are 300 such arrays (one per edge) of 300 entries each (one per SRG). For comparison. The deterministic approach needs an array for each reserved channel, where each entry corresponds to an SRG and indicates whether the SRG is protected or not by the reserved channel. In the solution of the NetB problems, 213052 of the channels are reserved for protection, thus 213052 arrays of 300 entries would be required in the deterministic method.

	Time to complete (sec.)		
	Determ.	Stoch.	Ratio.
<b>NetA</b>	156	23	6.78
<b>NetB</b>	9885	501	19.7

	Usage (# channels)		
	Determ.	Stoch.	%
<b>NetA</b>	61312	62716	102%
<b>NetB</b>	520771	536343	103%

**Table 1.** Deterministic versus stochastic algorithm, summary of results

Finally Figure 6 plots the distributions of sharing probabilities as computed in step 3(iii) of the stochastic algorithm during the provisioning of the demand in NetA and NetB. The distributions are similar and show that 70% of the time (77% in NetB) it was possible to determine almost certainly whether there would be a shareable reserved channel (probability 0.0 that an edge does not have shareable channel, 57% of the instances) or not (probability 1.0, 20% of the instances.)






Figure 6. Distribution of Sharing Probabilities (NetA and NetB)

## F. CONCLUSION

In this paper we describe a stochastic approach to identify shareable channels in a network when computing shared mesh-restored lightpaths. We show that a summarized information consisting of one fixed length array for every edge is sufficient to compute the paths efficiently while maximizing sharing opportunities. In contrast, the deterministic approach needs one such array for every protection channels, and thus does not scale when the demand grows.

Our results demonstrate that the stochastic approach completes the routing 6 to 20 times faster than the deterministic approach for networks ranging from 100 to 200 nodes. Although the stochastic approach uses several orders of magnitudes less information than what is necessary for a deterministic approach, their solutions are within 2% to 3% of each other in terms of capacity usage. In fact our experiments indicate that 70% of the time this little information is sufficient to determine with certainty whether a reserved channel could be shared or not.

One possible and natural application of the stochastic approach is for distributing the routing of shared mesh restored lightpaths to the optical switches. The local database of each switch may contain a summarized information that is necessary to compute the routes using the stochastic approach. Since this information is small, it can easily be disseminated by link-state protocols, such as

OSPF. Using this information each demand's ingress switch can compute a path equivalent to a path computed by a centralized deterministic algorithm with a complete view of the network's state.


## G. APPENDIX

In reference to question 2 of the problem presented in 1st, we show here how to compute the number of "non-blocking" combinations.

Let  $r(\max_i n_i - 1) = 0$ , and  $r(k) = \{C(k, M) \prod_j C(n_j, k)\} - r(k-1)$ .

Case 1: if  $M = \max_i n_i + 1$  then there is at most one bin empty, and the answer is the number of solutions in the remaining  $M-1$  bins. Thus  $D = C(M-1, M) \prod_j C(n_j, M-1) = r(M-1)$ , as expected.

Case 2: If  $M > \max_i n_i + 1$ , then there may be up to  $M - \max_i n_i$  empty bins. An incorrect answer would be to treat Case 2 in the same way as we treat Case 1, that is to remove 1 bin out of  $M$ , and compute all possible combinations in the  $M-1$  remaining bins. In order to understand why this is incorrect, take the case  $M = 2 + \max_i \{n_i\}$  and assume that we treat it as in Case 1. There can be up to 2 empty bins, and all combinations that have 2 empty bins will be counted twice, once for each of the two bins that is removed. Figure 7 illustrates this. The figure represents 3 bins, and one marble. If we remove one bin at the time and count the number of possible ways to place the marble in one of the two remaining bins, we observe that some combinations are equivalent. For instance, in combinations a) or b), c) or d) and e) or f) the marble respectively occupies the same position, but a different bin was removed. Therefore, although the computation in Case 1 would indicate 6 possible combinations, they are actually 3 of them. The argument presented in this simple example can be easily extended to cover the case of  $n$  marbles in  $M = n + 2$  bins. Observe now that if the number of combinations in  $(M-m)$  bins is known, then it is easy to derive from it the number of combinations in  $(M-m+1)$  bins. Let  $r(M-m)$  denote the number of combinations in  $(M-m)$  bins, then the number of possible combinations in  $(M-m+1)$  bins is  $C(M-m+1, M) \prod_j C(n_j, M-m+1)$  minus the number of combinations in  $(M-m)$  bins that would otherwise be counted twice, that is  $r(M-m+1) = \{C(M-m+1, M) \prod_j C(n_j, M-m+1)\} - r(M-m)$ . Replacing  $M-m+1$  by  $k$ , one recognizes the recursion presented in 1st part 2.



**Figure 7.** Combinations of 1 marble into 2 bins (out of 3 bins)

## H. REFERENCES

- [1] T.E. Stern and K. Bala, "Multiwavelength Optical Networks: A Layered Approach", Reading MA: Addison Wesley 1999.
- [2] B. Doshi, et al "Optical Network Design and Restoration" Bell-labs Technical Journal Jan-Mar 1999.
- [3] R.D. Doverspike et al, "Fast Restoration in a Mesh Network of Optical Cross-connects."
- [4] R.R. Iraschko, M.H MacGregor, and W.D. Grover, "Optimal Capacity Placement for Path Restoration in STM or ATM Mesh-Survivable Networks", IEEE/ACM Transactions on Networking, vol 6, Jun. 1998.
- [5] S. Chaudhuri, G. Hjalmtysson, and J. Yates, "Control of lightpaths in an optical network," in Optical Internetworking Forum. Jan 2000.
- [6] R. Ramamurthy et al, "Capacity Performance of Dynamic Provisioning in Optical Networks". IEEE JLT vol 19, Jan 2001.
- [7] R.D Doverspike and J. Yates "Challenges for MPLS in Optical Network Restoration", IEEE Communication Magazine, Aug. 1999.
- [8] Ramu Ramamurthy et al, "Comparison of Centralized and Distributed Provisioning of Lightpaths in Optical Networks" presented at OFC 2001, Anaheim CA.
- [9] J-F Labourdette, E. Bouillet, R. Ramamurthy, G. Ellinas, S. Chaudhuri, K. Bala, "Routing Strategies for Capacity-Efficient and Fast-Restorable Mesh Optical Networks," to appear in Photonic Network Communications, special issue on "Routing, Protection, and Restoration Strategies and Algorithms for WDM Optical Networks".