Lightpath Re-optimization in Mesh Optical Networks

Eric Bouillet, Priyadarshi Mishra, Jean François Labourdette, Kathi Perlove, Steve French
Tellium Inc., 2 Crescent Place, Oceanport, NJ 07757

Abstract

In this paper we study two algorithms to re-optimize lightpaths in resilient mesh optical networks. A complete re-optimization algorithm that re-routes both primary and backup paths, and a partial re-optimization algorithm that re-routes the backup paths only. We show that on average, these algorithms allow bandwidth savings of 3 to 5% of the total capacity in scenarios where the backup path only is re-routed, and substantially larger bandwidth savings when both the working and backup paths are re-routed.

1. Introduction

Intelligent mesh optical networks, supported by dense wavelength division multiplexed (DWDM) equipment and optical switches, are firmly established as the core constituent of next-generation optical networks. A key requirement of these optical mesh networks is the ability to quickly provision and restore services via fast and capacity-efficient end-to-end restoration schemes.

During operations, requests for services are received and provisioned using an online routing algorithm that takes all of the information available at the time of the request to make the appropriate routing decision. With connection rates reaching tens of Gigabits per seconds (Gbps), the ability of the network management system to operate and maintain service continuation during failures has become a challenging requirement. In this work we consider end-to-end shared mesh restorations as supported by Tellium's family of Aurora optical switches[1]. In end-to-end dedicated (1+1) mesh protection, the ingress and egress OXCs of the failed connection attempt to restore the signal on a predefined backup path that is disjoint, or diverse, from the primary path. Path diversity guarantees that primary and backup paths will not simultaneously succumb to the same failure. This approach requires large amounts of capacity, that is more than the working capacity since backup paths are longer than working paths. However the backup path remains "live" in permanence, thus saving crucial path-setup latency when recovery takes place. In shared mesh restoration (Figure 1), backup paths can share capacity if the corresponding primary paths are mutually diverse. Compared to dedicated (1+1) mesh protection, this scheme allows considerable saving in terms of capacity required[1]. In addition, the backup resources can be utilized for lower priority pre-emptible traffic in normal network operating mode. However recovery is slower than dedicated (1+1) mesh protections, because it involves signaling and path-setup procedures to establish the backup path. In particular, we note that the restoration time will be proportional to the length of the backup path and the number of hops, and if

---

[1] Other categories include line protection and re-provisioning. These are not considered here.

recovery latency is an issue this length must be kept under acceptable limits. This latter constraint may increase the cost of the solution, as it is sometimes more cost-effective to use longer paths with available shareable capacity than shorter paths where new shareable capacity must be reserved.

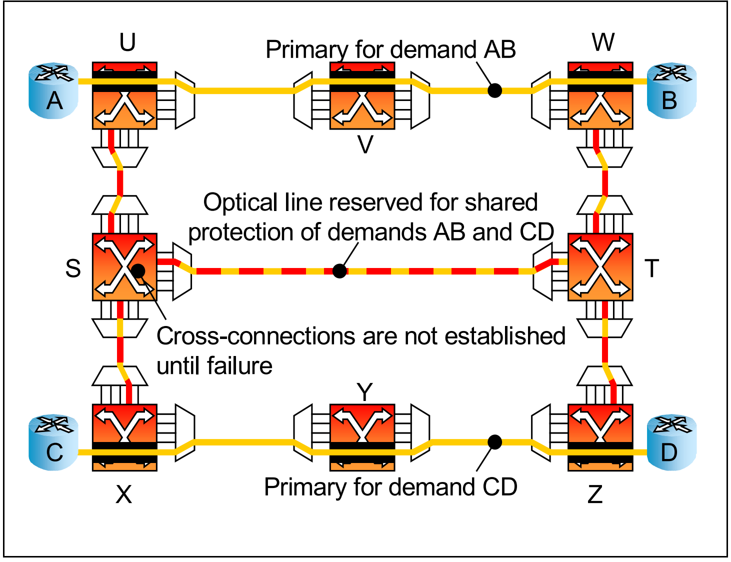

Figure 1. Shared Mesh Restoration

The primary and protection paths of each new demand are computed according to the current state of the network, which includes the routing of the existing demands. As the network and traffic evolve, the routing of the existing demands becomes sub-optimal. Demand churn and network changes such as the addition/deletion of new links and/or capacity, causes the routing to become sub-optimal, thereby creating opportunities for improvements in network bandwidth efficiency. Increasing customer churn and the continued demand for bandwidth services further exacerbates this problem.

Re-optimization seizes on these opportunities and offers the network operator the ability to better adapt to the dynamics of the network. This is achieved by regularly (or upon a particular event) re-routing the existing demands, temporarily eliminating the drift between the current solution and the best known solution that is achievable under the same conditions, as illustrated in Figure 2. Carriers can either re-route just the backup path so that existing services are not impacted, or re-route both the primary and backup paths, thus further improving network bandwidth optimization.

In this paper we study two re-optimization algorithms. A complete re-optimization algorithm that re-routes both primary and backup paths, and a partial re-optimization algorithm that re-routes the backup paths only. Re-routing backup paths only is a sub-optimal but attractive alternative that avoids any service interruption since the primary path is not affected (changed). In this paper we show that on average, these algorithms allow bandwidth savings of 3 to 5% of the total capacity in scenarios where the backup path only is re-routed. Substantially larger bandwidth savings can be achieved when both the working and backup paths are re-routed. These bandwidth savings are achieved through increased sharing of backup path capacity among several working paths, and substantial reductions in average path length, which also translates into shorter restoration times.

The paper is organized as follows. In section 2 we discuss the algorithm cost model and the main function used to compute the shared mesh restored paths that achieve the desired compromise between cost and restoration latency. In section 3, we describe the re-optimization algorithm. This algorithm uses the routing function discussed in section 2. The effectiveness of the re-optimization algorithm is measured for real customer networks and the results presented in section 4. We conclude this paper in section 5.

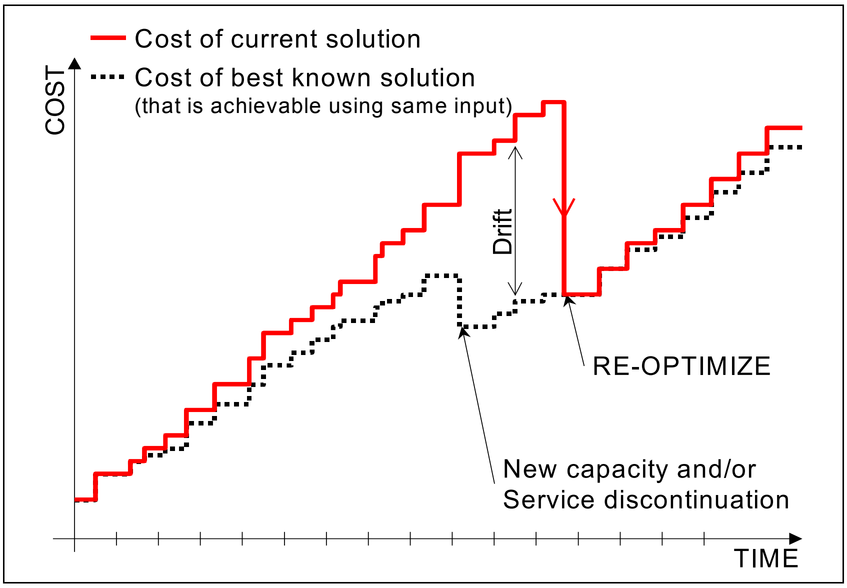

Figure 2. Current cost versus best possible cost with cost-benefit of re-optimization

2. Routing Algorithm

Cost model

We use the term Shared Risk Optical Group (SROG) to indicate a group of optical resources that share a common risk of failure. For the cost model we define the “length” of a path as the sum of the predefined weights of the links (or channels) that constitute it. The metric or policy used for weighting the links is different for primary paths and backup paths. For primary paths it is the real cost of using the links. For backup path it is a function of the primary path. A backup link *e* is assigned: (1) infinite weight if it intersects with an SROG of the primary path; (2) weight w_e if new capacity is required to provision the path; and (3) weight $s_e \leq w_e$ if the path can share existing capacity reserved for pre-established backup paths. The cost of a primary and its protection is then the sum of their respective lengths. Quite evidently, the underlying idea here is to encourage “sharing”, whereby existing capacity can be reused for provisioning multiple backup paths. The condition for sharing is that the backup paths must not be activated simultaneously, or in other words that their respective primaries must be pair-wise SROG-disjoint so that they do not fail simultaneously. The ratio s_e to w_e can be adjusted for the desired level of sharing. For smaller values of s_e, backup paths will be selected with the minimization of the number of non-shareable links (weights w_e) in view, eventually leading to arbitrary long paths (as expressed in number of hops) that consist uniquely of shareable links (weights s_e.) For larger values of s_e routing is performed regardless of sharing opportunities and backup paths will end-up requiring substantially more capacity.

Illustrative Routing Algorithm

Assume that provisioning of a lightpaths is performed in two steps: (1) computation of a primary and backup pair of routes, and (2) assignment of channels along the routes. Ideally the two steps are solved simultaneously and step (1) is optimized so that channel-assignment in step (2) reuses existing capacity for backup paths. For the only purpose of illustrating the cost model we present a K-shortest path based algorithm, keeping in mind that any other algorithm whose objective minimizes this cost model can also be used. The algorithm takes as input: (1) A network object N that encapsulates the state information of the switches, optical channels (busy and available), and existing demands with their routes; (2) the end nodes A and Z of the demand; and (3) a candidate primary path p_0 if partial re-optimization is desired. It operates as follows:

```
Compute_Pair_of_Paths (Network N, node A, node Z, candidate primary path $p_0$):

1. If $p_0$ is non-null, set P={$p_0$} and go to 5, otherwise compute a list of candidate primary
   paths:
2. For every link e in N set weight to cost $c_e$ of one channel in link (cost of transponders,
   regenerators and OAs.)
3. Compute set P of $k$ minimum-weight paths connecting node-pair A-Z, or all feasible paths
   if they are less than $k$ of them.
4. Set min_weight = infinity, and $\{p^*,q^*\}$=INFEASIBLE.
5. For each path $p$ in P, do
6.     Assign weight to every link e
         a. If e intersects SROG of primary path p, set weight to infinity.
         b. If e has at least one channel that is shareable with p, set weight to $s_e=\varepsilon c_e$.
         c. Otherwise, set weight to $w_e=c_e$.
7.     Using metric defined in 6, compute minimum-weight path q connecting node pair A-Z.
8.     If q does not exist, continue at step 5. with next path p in P.
9.     If min_weight < combined weight of p and q, then $\{p^*,q^*\}=\{p,q\}$ and min_weight =
       combined weight of $p$ and $q$.
10. Return $\{p^*,q^*\}$
```

If the minimum cost is sought (maximum sharing), the value of ε in step 6.b., determining the cost of “shareable” protection channels, is set to 0. Otherwise if shorter backup lengths, and faster restoration are desired, ε is set to a positive value. Extensive study has already been performed for ε=0 in [1]. In [7] we studied the effect of varying ε between 0 and 1. When ε tends toward 1, we expect the lengths of primary and backup paths, as expressed in number of hops, to resemble that of dedicated (1+1) mesh protection, though sharing is still implemented when available on the backup path and the capacity required remains lower than for dedicated (1+1) mesh protection. In the remainder of this paper we use ε = 0.3.

3. Re-optimization Algorithm

The re-optimization algorithm takes as input: (1) A network object N that encapsulates the state information of the switches, optical channels (busy and available), and existing demands with their routes; and (2) A list D of demands to be re-optimized with their respective re-optimization types (complete or partial). It operates as follows:

```
Reoptimize_Demands (Network N, list of demands D with respective re-optimization types)

1. Set REPEAT = 0
2. For each demand d in D, from A to Z
   a. Set $p_0$ = current primary path, and $q_0$ = current backup path of demand d.
   b. In network model N, free demand from paths $p_0$, and $q_0$.
   c. If partial re-optimization is desired do
            $\{p^*,q^*\}$ = Compute_Pair_of_Paths(N, A, Z, $p_0$)
                (note that $p^* = p_0$)
      else do
            $\{p^*,q^*\}$ = Compute_Pair_of_Paths(N, A, Z, null)
   d. If combined weights of $p^*$ and $q^*$ is less than combined weights of $p_0$ and $q_0$, then
      in network model N, provision demand d on paths $p^*$ and $q^*$, and set REPEAT = 1.
      Otherwise, in network model N, provision demand d back to paths $p_0$ and $q_0$.
3. If REPEAT > 0, repeat from step 1.
```

The key idea behind the re-optimization algorithm is not new[2]. Nevertheless, it is the first time to our knowledge that it has been applied to re-optimize shared mesh restored lightpaths. This algorithm is generic enough so that it is also applicable to re-optimize mixed protection types, i.e. combination of unprotected, dedicated mesh and shared mesh demands of various rates. It is also fast and easy to enhance with additional rules that improve the quality of the re-optimization. Finally, this algorithm provides the means to carry out the re-optimized solution by executing step 2.d. in the real network. The risks involved in step 2.d. are limited, since only one demand is re-routed at a time, and the operation does not impact the service if partial re-optimization is used.

Note that the optimum provisioning of shared mesh restored demands is a very difficult problem[1]. With the help of **Figure 3** we demonstrate the existence of at least one instance for which the algorithm fails to find the optimum solution. Part i) of **Figure 3** illustrates a 12 node network, with two demands, (**a**,**b**) and (**c**,**d**). We provision this demand using the re-optimization algorithm and all possible demand sequences $\mathbf{S_1}$={(**a**,**b**);(**c**,**d**)} and $\mathbf{S_2}$={(**c**,**d**);(**a**,**b**)}. Parts ii) and iii) of the figure depict two possible solutions. We find that these solutions require 2 primary channels, and reserve 8 channels for protection. The optimum solution, shown in part iv), requires 2 primary channels and reserves 6 channels for protection.

4. Experiments

We applied the algorithm to re-optimize the routes of 4 different networks, Net-A, Net-B, Net-C and Net-D. Net-A is a realistic network that consists of 45 nodes, 65 links, and 70 shared mesh restored demands with their routes provided by the operator. The available capacity for this scenario offers very little room for re-arranging the paths. Net-B consists of 25 nodes, 30 links and 290 demands. Net-C is a 45 node network with 75 links and 570 demands. Net-D is a 60 node network with 90 links and 195 demands. The demands of scenarios Net-B, Net-C and Net-D are provided without the routes. Henceforth, we created an initial configuration for these three scenarios by provisioning their demands sequentially following an arbitrary order, using the *Compute_Pair_of_Paths* procedure described in section 2. We added new channels as needed during that process. The demands of each scenario are then re-optimized, once partially and once completely, using the *Reoptimize_Demands* procedure of section 3.

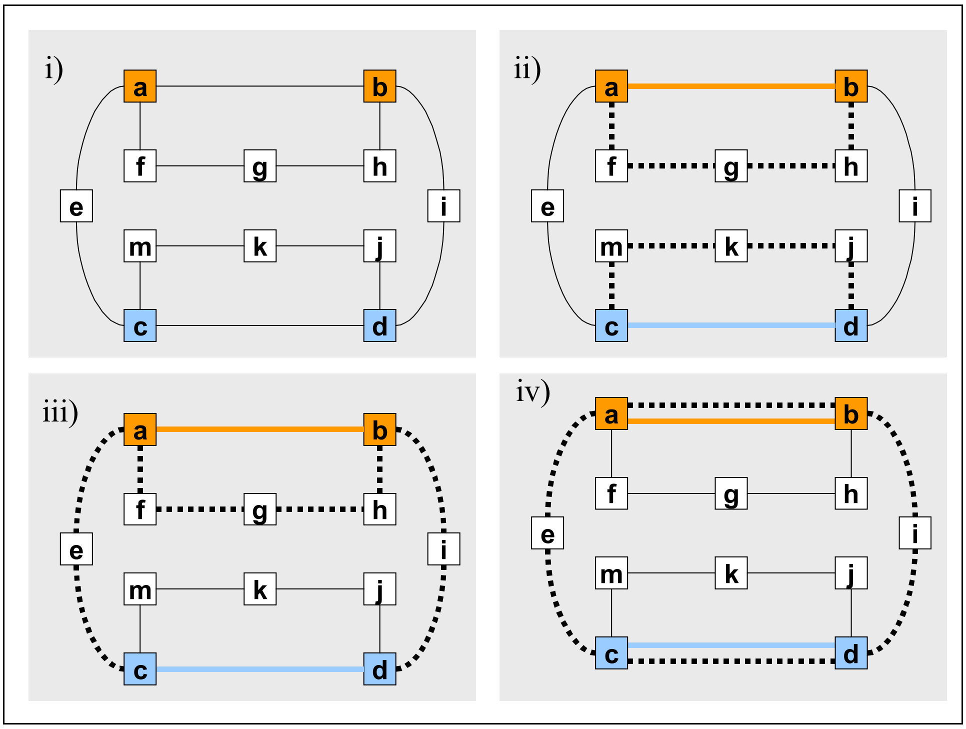

Figure 3. Shared mesh restoration architecture. i) Network with request for demands (a,b) and (c,d). ii) and iii) Two sub-optimum solutions, computed using the re-optimization algorithm. iv) An optimum solution

Tables 1 and 2 summarize the results for the partial and the complete re-optimization respectively. The tables show the quantities measured before and after re-optimization. For each scenario, the same network and routed demands are used for partial and for complete re-optimizations. The number of ports in Table 1 consists of ports used for the protection channels only, since the working channels remain the same. The number of ports in Table 2 consists of all the ports in the network, used for primary and protection channels. We observe that the partial re-optimization saves up to 4% of the total number of ports, and complete re-optimization up to 5%. The complete re-optimization offers the most cost efficient alternative, but most of the improvement is realizable using the partial re-optimization algorithm, without service interruption. Note that the savings for Net-A are substantial. It is possible that unlike the other scenarios where channel availability is not an issue, the demands of this network have been provisioned while new channels were being added, thus creating opportunities for optimization. The latter is the most realistic mode of operation, and the most likely to occur. Worth noticing for this scenario, is the reduction in protection latency measured as the average number of channels traversed by the protection paths, which decreases from 7.1 to 5.24 hops for the partial re-optimization.

Table 1. Partial Re-optimization

Scenario Name	Backup port count				Avg backup hops		Max backup hops	
	Before	After	% save	% save of total ports	Before	After	Before	After
Net-A	224	208	7%	4%	7.1	5.24	20	11
Net-B	2520	2452	3%	1%	5.83	5.76	11	10
Net-C	2290	2268	1%	0.5%	4.02	4.00	10	10
Net-D	556	508	9%	4%	3.71	3.74	9	9

Table 2. Complete Re-optimization

Scenario	Network port count			Avg backup hops		Max backup hops	
Name	Before	After	% save	Before	After	Before	After
Net-A	400	382	5%	7.1	5.43	20	10
Net-B	5088	4994	2%	5.83	5.72	11	10
Net-C	4864	4642	5%	4.02	4.20	10	11
Net-D	1138	1086	5%	3.71	3.73	9	9

5. Conclusion

In this paper we have presented a re-optimization algorithm to re-arrange shared mesh protected lightpaths. The proposed algorithm allows for two types of re-optimization. A complete re-optimization algorithm that re-routes both primary and backup paths, and a partial re-optimization algorithm that re-routes the backup paths only. Re-routing backup paths only is a sub-optimal but attractive alternative that avoids any service interruption. Our experiments indicate that the complete re-optimization achieves a 3% to 5% savings in the cost of the transport, and most of the improvement can be achieved by way of the partial re-optimization alone.

Contact author email: ebouillet@tellium.com, Tel: +1 (732) 483-2989, Fax: +1 (732) 728-9862

6. References

[1] G. Ellinas, E. Bouillet, R. Ramamurthy, J.F. Labourdette, S. Chaudhuri, K. Bala, “Routing and Restoration Architectures in Mesh Optical Networks”, to appear in October 2002 special issue of Optical Network Magazine (ONM).

[2] S. Baroni “Routing and Wavelength Allocation in WDM Optical Networks”. PhD Thesis, University of London, Dept of Electronic and Electrical Engineering, May 1998.

[3] K. Bala, T.E. Stern, “Multiwavelength Optical Networks, A Layered Approach”, Addision Wesley, 1999.

[4] J. 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*.

[5] S. Samieian, Z. Bogdanowicz, R. Ramamurthy, S. Chaudhuri, K. Bala, “Capacity Requirements of Ring, 1+1 and Mesh Configurations”, in *Lightwave Magazine*, August 2001.

[6] R. Ramamurthy, Z. Bogdanowicz, S. Samieian, D. Saha, B. Rajagopalan, S. Sengupta, S. Chaudhuri, and K. Bala, “Capacity Performance of Dynamic Provisioning in Optical Networks”, in *IEEE Journal of Lightwave Technology*, vol. 19, issue 1, pp. 40-48, January 2001.

[7] E. Bouillet, J.F. Labourdette, R. Ramamurthy, S. Chaudhuri, “Enhanced Algorithm Cost Model to Control tradeoffs in Provisioning Shared Mesh Restored Lightpaths”, OFC 2002, Anaheim, CA.