Shared Recovery in Transparent Optical Networks

Krishnan Kumaran, Iraj Saniee

Bell Labs, Lucent Technologies
600 Mountain Ave., Murray Hill, NJ 07974
{kumaran,iis}@research.bell-labs.com

***Abstract* --We consider the problem of transparent routing of lightpaths in a shared-protection optical network. Although optical routing problems have been extensively studied, there is no previous work that combines transparent routing *and* shared protection in a mesh network. This problem is important, however, because substantial network cost and capacity reduction can be achieved through sharing of protection lightpaths on a mesh. After a brief review of the "opaque" routing problem covering the commonly used *1:1,1:N* and shared protection, we present novel formulations of the transparent routing problem for each of these schemes. In particular, we present a model of shared mesh protection that combines *graph coloring* with packing of wavelengths into fibers on a *path intersection graph* to maximize wavelength sharing. An important feature that is incorporated into the model is elimination of fault isolation, which is typically infeasible in optical networks due to the stringent recovery time requirements. We propose and study heuristic solution techniques and bounds, via LP relaxation and rounding, on large example networks consisting of up to thousands of paths.**

***Index terms*-- transparent optical networks, shared protection in meshes, routing and wavelength assignment.**

A. Introduction

Optical networks are poised to successfully meet the explosive growth in bandwidth requirements of the future Internet. To make high bandwidths available to the end users, sophisticated switching, monitoring and control capabilities for optical signals are necessary. This is especially true of the core of the network, where these functions have been available for electronic transmission, but currently lacking for optical transmission. As a result, the backbone of most current networks operate in ''layers", with optics providing purely the transport layer while the intelligence resides at the electronic layer.

However, repeated conversion of signals between the optical and electronic layers is expensive, and ultimately infeasible, with concomitant delays and change in control that offset the high-speed advantages of optics. It is, therefore, advantageous to make networks ''transparent" [3,6], i.e. eliminate all Optical-to-Electronic-to-Optical (O-E-O) conversions. In a transparent backbone network, switching, monitoring and control are also provided at the optical layer.

The key enabling technologies to meet this goal include 1. the Ultra Long Haul (ULH) dense wavelength division multiplexing systems, 2. Optical Cross-Connect (OXC)s , 3. optical Wavelength Converters (WLC) and 4. optical monitors. ULHs can carry *wavelengths*, the basic units of resource in optical networks, across several thousand (~6000) kms without O-E-O regeneration. These systems can currently carry 80-160 wavelengths concurrently over a single fiber. OXCs switch wavelengths between multiple fibers carrying DWDM signals, and WLC devices convert between wavelengths optically. These devices provide the necessary switching, control, multiplexing and transport capabilities at the optical layer, thus enabling transparency in the next generation of networks, with substantial reduction in costs.

These advantages of transparency, however, come with higher complexity for the network design and management. For example, since a fiber cannot concurrently carry wavelengths with the same frequency, routing needs to be more carefully performed. This problem, known as routing with wavelength assignment (RWA) [3], leads to generalizations of the well known NP-hard problem of *graph coloring* even in very special cases such as when routing is fixed and all demands are for a single wavelength. It is expected that in operational networks simpler approximations with good performance will be used. The analysis and development of such schemes for transparent shared meshes is the goal of this report.

A critical consideration in the design and routing in the backbone network is that of reliability against failures of physical components. The most common forms of failure in optical networks include link (multiple fiber) cuts and OXC malfunctions (*node failures*). Under the assumption that these failures happen rarely, and with statistical independence, our analysis in this paper only addresses *single* component failures, as these would then be the most likely failure modes for the network. Further, due to the very high data rates supported by the network, recovery from failures needs to be extremely rapid [4], within 50-250 msec. This recovery speed is achievable through pre-provisioned *backup* paths in the network. Furthermore, depending upon the recovery scheme chosen, the backups may be *dedicated* for each demand, or *shared* across demands. However, there is strong evidence that mesh recovery is substantially more efficient and less costly than dedicated protection [1]. Although shared mesh recovery has been studied in the context of non-transparent optical networks [2,5], there is no work, to our knowledge, on transparent optical networks with shared restoration. It is natural to ask if good routing approximations exist or indeed if the efficiency and cost optimization observed in "opaque," or non-transparent, shared meshes carry over to transparent networks. These are also addressed in the following sections.

In Section B, we provide an overview of opaque routing in mesh networks, derive analytical results for the 1:N protection and tabulate computational results for the shared protection in meshes. In Section C, we introduce a novel formulation of the problem of transparent routing with protection. In particular, this formulation aims to maximize the sharing of protection for the transparent mesh. We propose both heuristic solutions to the model and linear programming (LP)-based solutions. In Section D we present computational results on large example networks.

B. SHARED PROTECTION IN OPAQUE MESHES

An optical network is *opaque*, as the name suggests, when there is O-E-O conversion at every node. In opaque networks there are no wavelength assignment considerations and only capacity needs to be provisioned. The simplest way to provide protection against single physical failures is to duplicate, via a diverse route, the capacity needed for each node pair. This scheme, known as 1:1 or 1+1, is also the most costly and the least efficient. To improve efficiency, multiple demands between the same node pair could share protection as long as there are more than two disjoint paths between the node pair. This scheme gives rise to 1:N protection for sharing between the *same* node pairs. Further improvements can be obtained by sharing *across* node pairs. The ring is an example of such sharing. Rings have been studied in much detail and we will not consider them explicitly. However, routing in 1:N protection and optimized meshes will be reviewed.

In the formulation that follow, we *decouple* the path selection problem from both routing and protection assignment. This is done for both the opaque and transparent routing problems. This constraint on the set of admissible paths is necessary to meet quality of service requirements, such as delay and jitter. This qualification typically reduces the problem to the best (shortest) set of *K*-disjoint paths. In real networks, the set of feasible candidate routes is sufficiently small, on account of the low degree of physical layer connectivity. Such set of *K*-disjoint paths can be efficiently generated, see [7] and its generalizations.

1st. 1:N Sharing

A computationally simple method for the shared-protection routing in meshes can be obtained by considering routing for single point-to-point demand, see **Figure 1**.

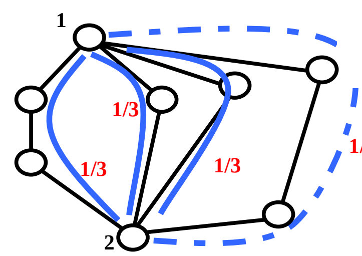

Figure 1. Short primary (solid) paths and longer protection (dotted) path for demand between nodes 1 &2. The fractions in red denote demand carried by respective path.

The procedure is clearly more efficient than 1:1 protection and can be used to provide a fast solution to a special case of a more comprehensive mesh design to be discussed in the next subsection. We refer to this algorithm as the *1:N* shared protection.

Consider the restoration-adapted version of the standard uncapacitated *multi-commodity flow problem:*

$$\text{Minimize} \quad \sum_{p \in P} x_p * c_p$$

$$\text{Subject to} \quad \sum_{p \in P_k,\, p \neq q} x_p \geq d_k, \quad \forall k \in K, \forall q \in P_k$$

where

N	Set of all network nodes
E	Set of all network edges
K	Set of all node-pairs with demand
P	Set of all paths for all node-pairs
D	Set of all node-pair demands
d_k	Demand for node pair k
x_p	Flow (in units of λ) assigned to path p

P_k Set of all paths for node pair k – assumed (edge) disjoint

The goal is to minimize the total wavelength distance traversed in the network while protecting against single link/node failures. The source-destination pair *s-t* is assumed multiply connected by N candidate paths, each path i with a given cost c_i of carrying unit flow. The costs may be ordered as $c_1 \le c_2 \le .... \le c_N$. For single-path-failure reliability, it is easy to see that $c_N \le c_1 + c_2$, i.e. the set of candidate paths may be pruned in this manner without loss of optimality. We seek to route, at minimum cost, a total demand D between s and t along the paths such that service is not interrupted by single path failure. This problem has the following simple LP formulation directly derived from (1) above:

Letting $\sum_{k=1}^{N} x_k = D + D'$, the above LP is written as a classical *knapsack problem* for fixed D'.

$$\begin{aligned} &\text{Minimize} \quad \sum_{k=1}^{N} c_k x_k \\ &\text{subject to} \quad \sum_{k=1}^{N} x_k = D + D', \qquad (3) \\ &0 \le x_k \le D' \quad \forall\ j \in \{1,2,...,N\}. \end{aligned}$$

It can be shown that (3) has the following simple solution: Let $K = \lceil 1 + D/D' \rceil$.

$$\begin{aligned} x_k &= D'\ \forall k < K; \\ x_k &= 0\ \forall k > K; \\ x_K &= D - (K-2)D'. \end{aligned}$$

Further, a simple search over D': $D/(N-1) \le D' \le D$ produces the globally optimal x_k independent of whether or not they are restricted to take integer values. The end result of this procedure is that the optimal set of chosen paths, for large demands, will satisfy $\frac{1}{K-1}\sum_{k=1}^{K} c_k < c_{K+1}$. This relation has the interesting interpretation that the optimal solution, independent of the demand, chooses to add new paths to route the demand only until the incremental additional cost of unit restoration capacity continues to decrease, beyond which it load balances among the previously chosen paths. This implies (except for integrality of flows) equal allocation of flows to the set of chosen paths. It is also possible to show that complete load balancing over all paths, i.e. the choice $x_k = \lceil D/(K-1) \rceil$, is never more than a factor of 2 in total cost from the optimal solution. Finally, all of the above conclusions can be generalized for the case of reliability against multiple failures, individual capacity limits for each path, variable degrees of protection from full backup of affected demand to smaller percentage, as well as differing fixed installation costs for each path. Based on the solution to this single demand routing problem, a heuristic solution to the more general shared mesh network design that accounts for sharing across demands can also be constructed. This algorithm provides a limited form of sharing between demands, as well as between routes for a given demand, and produces a sub-optimal approximation to the fully shared mesh network design that we discuss next.

2nd. Optimal Sharing in Opaque Meshes

Most current optical networks employ ring architectures [4] (see **Figure 2**) to perform efficient shared restoration. This architecture allows sharing of restoration capacities due to the fact that the backup ring, whose capacity is equal to the maximum link capacity of the primary ring, is never simultaneously used by demands that don't share a link on the primary ring. The concept of mesh restoration generalizes this notion to arbitrary network topologies to gain even further efficiency from sharing, which we shall attempt to quantify through example networks later.

We first define the problem of shared restoration for opaque networks in its general mathematical form.

N	Set of all network nodes
E	Set of all network edges
K	Set of all node-pairs
P	Set of all paths for all node-pairs
D	Set of all node-pair demands
d_k	Demand for node pair k
x_p	Flow (in units of λ) assigned to path p
P_k	Set of all paths for node pair k – assumed to be Disjoint
Q_e	Set of all paths that include edge e
C_e	Capacity (in number of λ s) on edge e
l_e	Distance of edge e
y_p^f	Flow (in units of λ) assigned to path p when edge f is in failure mode

To ensure that bi-directional demands are routed the same way, we restrict $(i, j) \in K$, to $i < j$. Furthermore, to reduce problem dimensionality, we include a path in P only if its end nodes have demand. To ensure that restoration from failures does not require fault *isolation*, we assume paths in each P_k are disjoint for each commodity k. Therefore, when end node pairs recognize a fault in (one of) their primary path(s), the node pair initiate restoration over the remaining paths. This assumes that due to their being disjoint from the failed path, the backup paths must be operational in cases of single (link) failures.

To avoid solutions with meandering paths–*e.g.*, those that are either unusually long or visit the same node more than once–the set of paths P needs to constrained. Examples include paths that do not exceed pre-specified hop counts or distance limits between their end nodes. Further, more careful path generation reduces the problem complexity. In realistic networks K is rarely above 10 and typically no more than 6.

The *mesh network routing problem* can be formulated according to a variety of criteria using the foregoing notation. For example, it may be important to route as much of the demand on shortest possible routes. In this case the natural criterion may be the sum of wavelengths times the distance each wavelength traverses. Or it may be desired to spread the load as evenly as possible so that all the edges in the network carry more or less the same number of wavelengths. For the latter criterion, the following formulation results.

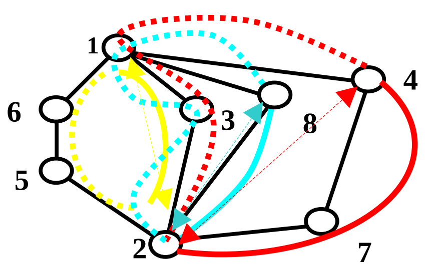

Figure 2. Active and restoration paths and wavelengths for demands between nodes 1-2 (1-3-2 and 1-6-5-2), 2-4 (2-7-4 and 2-3-1-4) and 2-8 (2-8 and 2-3-1-8)

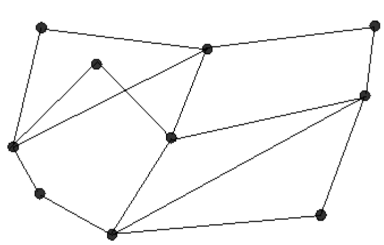

Figure 3. A prototypical backbone network based on NSFnet.

$B4(\lambda * dist)$

$$Minimize \quad \sum_{e \in E} \lambda_e * l_e$$

$Subject \ \ to:$

$$(1) \qquad \sum_{p \in P_k} x_p \geq d_k, \qquad \forall k \in K$$

$$(2) \qquad \sum_{p \in Q_e} x_p \leq \lambda_e, \qquad \forall e \in E$$

$$(3) \qquad \lambda_e \leq C_e, \qquad \forall f \in E$$

$$(4) \qquad \sum_{\substack{p \in P_k - Q_f \\ f \in q \in P_k}} y_p^q + \sum_{p \in P_k - Q_f} x_p \geq d_k, \qquad \forall f \in E, \forall k \in K$$

$$(5) \qquad \sum_{\substack{p \in Q_e - Q_f \\ f \in q \in P_k}} y_p^q + \sum_{p \in P_k - Q_f} x_p \leq \lambda_e, \qquad \forall e \neq f \in E$$

Constraint set (1) ensures that demands (in number of λs) are met for each node-pair and (2) the link/DWDM capacities are not exceeded. Constraint set (3) counts the number of wavelengths used on each edge, which is used in the objective function. Constraint set (4) ensures that for each failure, the restored demand is at least the required capacity and finally (5) ensures that the capacity used in an edge under all failure conditions is no more than the capacity provisioned.

We note that in the above formulation, the recovery procedure does not depend on the failure condition *f*, it only depends on the paths that are affected by the failure. The restoration mechanism thus does not requires *fault isolation*. Note that in this formulation for each failure only a small number of reconfigurations will be necessary.

3rd. Computational Results for Opaque Design

A prototypical backbone network consisting of 10 nodes and 16 links with demands ranging between 5-20 λ-circuits between all node-pairs was examined, see **Fig 3**. The design involves optimizing for the 1+1 protection and the shared mesh as described in Section B.2. The results are tabulated in Table 1.

As this table shows, the shared mesh protection scheme is subtantially more resource efficient than the 1+1 and subsequently less costly. The ~50-75% increase in efficiency due to sharing of protection wavelengths is typical across several of other realistic examples ranging in size from 10 to 100 nodes we examined.

	No Protection		1+1 Protection		Shared Protection Mesh	
	# Units	Cost	# Units	Cost	# Units	Cost
Bi-dir λs	1215	NA	3K	NA	1967	NA
Bi-dir λ*km	688k	NA	1.8M	NA	1.2M	NA
λ-efficiency	***100%***		***~40%***		***~60%***	
Bi-dir fiber km	14K	21M	26M	40M	16K	25M
Normalized Cost	***1***		***1.9***		***1.2***	

Table 1. Computational results for a typical network.

C. SHARED PROTECTION IN TRANSPARENT MESHES

Design of *transparent* meshes involves the additional step of assigning a single wavelength for each unit of demand end-to-end.[1] It is well-known that the wavelength assignment problem, for a given set of routes between node pairs, can be formulated as a graph coloring problem [8]. We generalize this formulation to include both the use of multiple fibers and also allow for shared protection as described in Section B. Further, the model admits computationally simple approximations and heuristics that we also describe. As before, we decouple the *selection* of candidate routes from the flow and wavelength assignment problem. However, feasible paths have even more stringent requirements in transparent routing since regeneration is not performed.

With the route selection and wavelength assignment decoupled, we first describe the problem without protection and follow this up with the maximal shared protection formulation. We subsequently present computational results and approximations.

4th. Transparent 1:N Shared Protection

Sharing protection lightpaths between the same node pairs, or the 1:N scheme introduced at the beginning of Section B, admits a simple extension to transparent lightpaths as follows. Assigning a fixed set of *D'* wavelengths to all paths, that were selected in the opaque design, provides the required transparency on these paths. Repeating this over different node-pairs completes the assignment.

[1] Technically, transparency is defined in this way only in the absence of all-optical wavelength conversion. However, the fewer convertors used the better and hence the problem as defined is still relevant even when optical wavelength convertors become available.

5th. Unprotected/Dedicated Wavelengths with Multiple Fibers

Consider a network of N nodes and E links operating with F wavelength channels. Each link l has f_l fibers, which allows f_l copies of each wavelength channel $\nu \in \{1,..,F\}$ on the fiber. The fiber costs are c_l per fiber on link l. Each source destination pair (i,j) has a set P_{ij} of candidate paths over which a total demand of d_{ij} must be satisfied. Define the variables n_{ν_p} that denote the number of copies of wavelength ν to path p. Using the paths P_{ij}, which are specified as a sequence of links traversed in the network connecting nodes i and j, we may construct the so-called *Path Intersection Graph* (PIG)[8] as follows: Each node in the PIG corresponds to a path in the network, and two nodes in the PIG are joined by an edge iff the corresponding paths share at least one link in the network.

Given the PIG, with specified demands for each path, the wavelength assignment may now be performed using any good heuristics for graph coloring, which is hard to do optimally, in general NP-hard. In the transparanet routing problem, in addition to wavelength assignment, it is also critically important to "pack" the wavelengths onto exisiting/minimal additional fibers. Also, note that the multiple fibers on a link provide "duplicate" copies of each wavelength, and the number of these copies is itself another decision variable in a greenfield network design version of the wavelength assignment problem.

To take the above-described features into account adequately, we depart from the conventional graph coloring approach and formulate the following *Integer Linear Program* (ILP) for the Optimal Wavelength Assignment (WA) problem. A useful notion provided by the PIG view of the problem is that of *cliques*, a set of mutually connected nodes. It is easy to see that the set of flows that traverse a common link in the network from a clique on the PIG, and the identification of these cliques is particularly useful not only for graph coloring[REFS] but also in the formulation of constraint sets in the ILP. For a given number of fibers f_e on edge e,the ILP is given by

$$(C1) \quad Minimize \quad \sum_{e=1}^{E} \sum_{p:e\in p} \sum_{\nu=1}^{F} n_{\nu p} l_e$$

subject to

$$\sum_{p\in P_k} \sum_{\nu=1}^{F} n_{\nu p} \geq d_k \qquad \forall\, k \in K$$

$$\sum_{p:e\in p} n_{\nu p} \leq f_e \qquad \forall\, \nu, p \in P$$

It is easy to see that, even in the special case of $d_k=1\ \forall k$, F very large and $l_e=1\ \forall e$, and $f_e=1$, the WA problem gives an instance of the well-known NP-hard *graph coloring* problem. If we seek to solve for a greenfield network design, the f_e may be treated as variables, and the objective may include the total weighted fiber length $\sum_e l_e f_e$.

The ILP could be approximately solved using the LP relaxation followed by a rounding procedure that preserves feasibility. Indeed, when the f_e and d_k are large enough, this method could prove very effective. Even when LP-rounding while preserving feasibility does not provide good results, as may occur for small demands between node pairs, the procedure still provides useful bounds for the problem.

6th. Shared Protection Wavelength Assignment with Multiple Fibers

We generalize the PIG construct for the shared protection mesh routing problem with multiple fibers. We note that different PIGs can be constructed depending on (1) If the opaque routing problem is first solved and flows on primary and backup paths are given (separate routing and wavelength assigment), (2) If the primary and back up paths are identified without flow values, (3) No information about flow values or primary and backup paths is given (simultaneous solution of the routing and wavelength assignment).

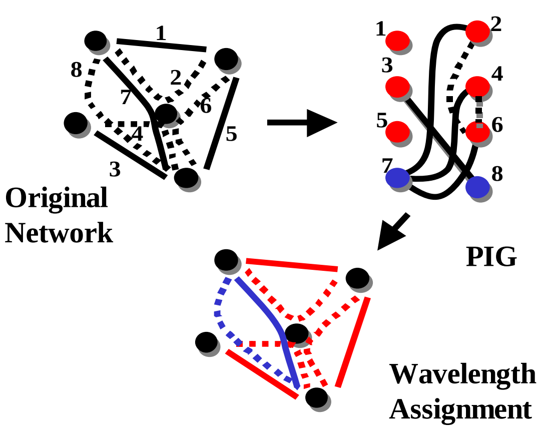

Figure 4. The Path Intersection Graph construction for shared mesh. Dotted lines are protection paths for the active solid lines. Observe that even though paths 2 & 6 share a link, their corresponding nodes in PIG are not connected (shown by dotted line) due to non-concurrency of any failure that could invoke them.

To start with, we first consider cases (1) & (2). The part of the PIG involving any "primary" node is constructed as in Section C.4, with an edge between two nodes (either both primaries or one primary and one backup) if the corresponding paths share a link in the original network. For each pair of "backup" nodes *p* and *p'*, we perform the following:

1. If the backup node *p'* shares links with *p* (in the original graph) ascertain whether *p* and *p'* may be simultaneously invoked upon failure of some network component.
2. If yes, connect *p* and *p'* by an edge on the PIG.

See Fig. 3 for an example. Observe that a valid wavelength assignment on the PIG constructed as above provides a feasible solution with shared protection, as intersecting backup paths that are not simultaneously invoked are permitted to share a wavelength. It is however not clear that the optimal assignment on this PIG corresponds to the optimal wavelength assignment in the network without failure isolation. To see that this is true, note that, in the absence of fault isolation (a common requirement, see Section B) we must account for the worst case possibility that a given backup path may be invoked together with any of the other backups that any failure in the network may necessitate. That is, restoration should only depends on the paths that are affected by each failure, and not by the specific network element that fails. This implies that the wavelength assigned to the backup paths is independent of the failed network element.

In case (1), the ILP formulation of the problem is similar to (C1) presented above, except that more constraints are necessary to capture the correlations between backups. We give this formulation below for illustration. Let P_k denote the set of primary paths for node pair k and Q_k the set of backup paths. Each path $p\in P_i$ is paired with a corresponding $q\in Q_i$, with the q possibly being shared by many ps. When a given p needs to be protected by multiple qs, we break up p into multiple instances with the flow on p apportioned appropriately. We finally obtain a set of paths, each carrying a given flow $d_{p/q}$ for each path p/q. We then obtain the ILP formulation

$$(C2) \quad \textit{Minimize} \quad \sum_{e=1}^{E} \sum_{p:e\in p} \sum_{\nu=1}^{F} n_{\nu p}\, l_e$$

subject to

$$\sum_{\nu=1}^{F} n_{\nu p} \geq d_p \qquad \forall\, p \in P_k, Q_k$$

$$\sum_{\substack{p:e\in p \\ p\in P_k}} n_{\nu p} + \sum_{\substack{q=q(p) \\ p:e'\in p, e\in q}} n_{\nu q} \leq f_e \qquad \forall\, \nu, e \neq e'$$

In case (2), where the flows on each path are not specified, but only the primary demands d_k and the sets P_k ,Q_k are given, it is possible to combine formulation (C2) with that of the 1:N protection described in Section B to model WA problem accurately. For brevity, we omit the explicit formulation here, but mention that, along the lines of the 1:N formulation, the RHSs of demand satisfaction constraints are modified as D_k= d_k+d_k' where d_k' are now treated as variables, and the flow on each path for the demand *k* is restricted to be atmost d_k'.

7th. A Heuristic for Transparent Shared Mesh Routing

Based on the above PIG formulation, we develop the following simple greedy heuristic to assign end to end wavelengths while simultaneously packing wavelengths into fibers. The steps of the method are as follows:

1. Sort the demands in decreasing order of bandwidth*hop product, and index the wavelengths as 1,…,F.
2. Starting at the top of the sorted list, assign each demand the lowest indexed end-to-end wavelength that does not conflict with previous assignments.

➢ Note that the test of conflict, while straightforward for primary demands, is more involved for backups and uses the PIG structure outlined above to compute maximum concurrent backup flow on each link.

3. When a conflict cannot be resolved with existing number of fibers on each link, add the minimal number of additional fibers to find an end-to-end wavelength.

The above simple heuristic is generally of low complexity for most networks. However, note that its complexity grows with the total demand in the network, as also the lengths of the paths, while the LP relaxation and rounding approach is insensitive to these parameters. The method is hence most suited to moderate sized demands and path lengths, but can accommodate large numbers of nodes, edges and paths in the network.

D. Computational Results

The heuristic solution and LP bounding methods were applied to a set of real networks of varying sizes and connectivity as described in Table 3. The results of the experiments are summarized in Tables 4 and 5, with the number of fiber hops as the performance metric. Observe that the heuristic method comes close to the LP lower bound (atmost 10-15% from optimum for the larger examples not counting N3) in cases when the lower bound is computable using standard LP solvers, and is also comparable to the shared protection opaque designs (10-25% difference). For the larger networks, the LP lower bound could not be computed due to the inability of existing LP solvers in dealing with very large numbers of variables and dense constraint matrices. Table 5 compares 3 transparent designs – unprotected, shared mesh and dedicated – for the same input data. While the difference from the unprotected design is significant, sometimes a factor of 2, note that the sharing still provides significant gains compared to dedicated protection, sometimes as high as 30%.

	Problem Size			
	Nodes	**Links**	**Paths**	**Total Demand (λ)**
N1	18	25	90	2988
N2	10	16	106	1814
N3	15	31	108	473
N4	225	301	824	2576
N5	30	38	902	6998
N6	38	48	1406	3403

Table 3. (Real) Example networks used in numerical experiments.

	Number of Fiber Hops		
	Shared Mesh Heuristic	**LP Lower bound**	**Opaque Shared Mesh**
N1	140	125	98
N2	50	41	46
N3	30	18	22
N4	433	-	346
N5	437	-	389
N6	226	-	178

Table 4. Heuristic solution compared with Shared mesh opaque design and the LP lower bound.

	Unprotected Transparent Design	Shared Mesh Heuristic	Dedicated (1+1/1:1) Design
N1	70	140	163
N2	31	50	71
N3	22	30	37
N4	311	433	546
N5	253	437	535
N6	123	226	280

Table 5. Comparison of transparent designs: Shared Mesh vs. Unprotected and Dedicated protection designs.

E. CONCLUSIONS

We have presented novel formulations and solution methods for transparent routing and protection that combines graph coloring and packing of wavelengths into fibers. In particular, we have shown how the formulation can be adapted to take sharing into account effectively in mesh networks. Numerical results suggest that our methods are reasonably efficient, both in the design produced as well as in computational requirements, over a broad range of network sizes. In particular, the results demonstrate that the heuristic presented is both efficient and scalable.

REFERENCES

[1] R. Davis, K. Kumaran, G. Liu, and I. Saniee, "SPIDER: A simple and flexible tool for design and provisioning of protected lightpaths in optical networks", *Bell Labs Technical Journal*, Vol. 6, No. 1, January 2001.

[2] R. Iraschko, M. MacGregor, W. Grover, "Optimal capacity placement for path restoration in STM or ATM mesh survivable networks", *IEEE/ACM Trans. on Networking*, vol.6, No.3, June 1998, 325-336.

[3] R. Ramaswami, K. Sivarajan, *Optical Networks: A Practical Perspective,* Kaufmann Morgan Publishers. 1998.

[4] F. Ritchie, "SONET lays the roadbed for broadband networks", *Networking Management,* 1990.

[5] I. Saniee, "Optimal routing in self-healing communications networks", *Int. Trans. in Operations Research*, 3, pp. 187-195, 1996.

[6] T. Stern, K. Bala, *Multiwavelength Optical Networks, A Layered Approach*, Addison Wesley Longman, Reading MA, 1999.

[7] J. W. Suraballe, “Disjoint paths in a network”, *Networks*, Vol 4, pp 125-145, 1974.

[8] A. C. Tucker, "Coloring a family of circular arc graphs", *SIMA Journal on Applied Math*, Vol 29, pp. 493-502, 1975.