

Revenue Maximization in Survivable WDM Networks

Murari Sridharan and Arun K. Somani

Dependable Computing & Networking Laboratory
Department of Electrical & Computer Engineering
Iowa State University, Ames, IA 50011

ABSTRACT

Service availability is an indispensable requirement for many current and future applications over the Internet and hence has to be addressed as part of the optical QoS service model.¹⁻³ Network service providers can offer varying classes of services based on the choice of protection employed which can vary from full protection to no protection. Based on the service classes, traffic in the network falls into one of the three classes viz., full protection, no protection and best-effort. The network typically relies on the best-effort traffic for maximizing revenue. We consider two variations on the best-effort class, (i) all connections are accepted and network tries to protect as many as possible and (ii) a mix of protected and unprotected connections and the goal is to maximize revenue. In this paper, we present a mathematical formulation, that captures service differentiation based on lightpath protection, for revenue maximization in wavelength routed backbone networks. Our approach also captures the service disruption aspect into the problem formulation, as there may be a penalty for disrupting currently working connections. Since the combined problem for solving all demands together can be quite complex, we propose a solution methodology for extracting a feasible solution. We also compare the increase in revenue got by the two variations of the best-effort class with the case of accepting demands without any protection. If the cost of a backup path is α times the cost of a primary path, then, for particular instances of demands, the best-effort variation 1 results in a 67% gain in revenue and variation 2 achieves an additional 6% gain in revenue, for $\alpha = 1$

Keywords: WDM, Optical Layer Protection and Restoration, Optimization, Revenue, Survivability, Service Differentiation

1. INTRODUCTION

An explosion in the growth of web-related services offered over the Internet is creating a growing demand for bandwidth. All-optical networks with wavelength-division multiplexing (WDM) are considered to be a promising technology for next generation transport networks, as they can effectively satisfy this growing demand for bandwidth. In WDM networks, the huge bandwidth available on an optical fiber is divided into multiple channels. Each channel can carry bandwidth upto several gigabits per second. A minimum unit of resource allocation is an optical channel, which consists of a route and a wavelength assigned on each link along the route. If wavelength translation is performed in optical switching, then each channel may be assigned different wavelengths on each link along the route; otherwise the wavelength continuity constraint must be satisfied on all links along the route.

Today's Internet is dominated by applications and services based on the ubiquitous Internet Protocol (IP). The trend is likely to continue as IP continues to provide a form of protection and restoration by enabling packets to be dynamically rerouted around link or node failures. With TCP providing a reliable transport service, it is very likely that IP based applications will continue to dominate the Internet traffic for years to come. It is therefore evident that the WDM backbone networks be optimized for IP services. Many factors make it attractive to carry fast growing IP traffic directly over an optical network without the intervening SONET/SDH layer. In such cases, the entire network needs a new restoration strategy. SONET has its own protection schemes providing fast recovery (of the order of milliseconds). The relative benefits of providing restoration at either the service or the optical layer will continue to be debated. Restoration at the optical layer has several advantages a) recovery mechanisms will be

Further author information: (Send correspondence to Murari Sridharan)

Murari Sridharan: E-mail: murari@iastate.edu

Arun K. Somani: E-mail: arun@iastate.edu

much faster b) optical layer can better optimize resources such as wavelengths c) provides protection to higher layer protocols which do not have their own recovery mechanisms.

The challenge is to react quickly to these increasing bandwidth requirements while maintaining reliable service, and to design and operate networks to provide adequate capacity in the geographical areas where demand is growing fastest without over provisioning to the point of reducing the network revenue.

1.1. Related Work

To date, design problems in mesh-survivable WDM networks have been studied in.⁴⁻⁹ The study in⁴ proposes an optimal design scheme for survivable WDM transport networks in which fast restoration can be achieved by using predetermined restoration paths. Integer programming (ILP) based design problems were formulated to optimally determine working and their corresponding restoration paths, the number of fibers in each span, and the optical cross connects in each node. In,⁵ ILP and simulated annealing (SA) were used to solve optimization problems for routing, planning of working capacity, rerouting, and planning of spare capacity in WDM networks. The purpose of the study was the design of a fiber topology and optical path layer for WDM Networks, with a fixed channel plan, minimizing the total cost for a given traffic demand. The work in,⁶ aims at providing design protection that is well adapted to WDM networks, where many channels share the same fiber. The design protection, however, does not guarantee carrying all the traffic that was carried prior to the failure. Instead, it aims at maintaining connectivity between all pairs of network ports following a single failure and lets the higher level network layers reconfigure itself so as to carry only the high priority traffic. Joint optimization of primary and restoration routes to minimize the network capacity was studied in.⁷ Given a network, a set of point-to-point demands, the optimization problem was formulated to find primary and restoration route for each demand so that the network capacity is minimized. The study also tried to determine the best restoration route for each wavelength demand, given the network topology and the capacities and primary routes of all demands. The work in⁸ mainly concerns connection provisioning for optical networks. An heuristic algorithm was developed for routing and wavelength assignment for a set of static connections and an adaptation of the algorithm was proposed to handle a set of failures. The study in⁹ examines different approaches to protect mesh-based WDM optical networks from single-link failures. ILPs were formulated to determine the capacity requirements for a static traffic demand based on path/link protection/restoration survivability paradigms.

To date, the design problems in optical network design and restoration have considered a static traffic demand and tried to optimize the network cost assuming various cost models and survivability paradigms. Fast restoration has been a key feature addressed on all of the designs. In our earlier work,¹⁰ we formulated an ILP to address the following problem. Given the current demands and the new set of demands to be provisioned, optimize the network capacity, while avoiding service disruption to the current working connections. We use the terms demands and connections interchangeably in the paper. This optimization problem can be treated as a static formulation by removing all current connections and optimizing the network capacity for the complete demand set which includes the current working and the new demands. On the other hand, we could avoid disrupting any of the current working demands (by removing the links used by the current working demands) and optimizing the network capacity for the new demands. The former treatment provides the best capacity optimization, but all the current paths are disrupted which is not acceptable. The latter case avoids disruption to the current working paths and tries to optimize on the remaining capacity which may not yield the best solution. To the best of our knowledge, none of the existing formulations capture the service disruption aspect into the problem formulation.

1.2. Proposed Extension

In this paper, we extend our work to include a service differentiation model based on lightpath protection and formulate an ILP problem for maximizing the total revenue generated by the network. Network service providers can offer varying classes of services, based on the choice of protection employed, which can vary from full protection to no protection. Based on the service classes, traffic in the network falls into one of the three classes viz., full protection, no protection and best-effort. The first class comprises of high priority traffic which require full protection in the optical layer. Many carriers may have already invested hugely in their networks and their equipment may not support protection and such networks have to rely on the optical layer for protection. The second class comprises of high priority traffic with no protection, as they may already be protected by higher layers such as SONET. The best-effort class tries to provide protection for the connections based on the resources available. These may include IP traffic which have their own protection mechanisms, which tend to be slower. Optical layer protection is also beneficial for traffic which does not have any stringent protection requirements, but can pay for protection if the network

has enough resources available. The network typically relies on the best-effort traffic for maximizing revenue. In this paper, we present a mathematical formulation for revenue maximization in wavelength routed metropolitan networks that considers a service differentiation model based on lightpath protection.¹⁻³ There may be a penalty for disrupting currently working connections and we capture this service disruption aspect into our problem formulation. We consider two variations on the best-effort class, (i) all connections are accepted and network tries to protect as many as possible and (ii) a mix of protected and unprotected connections and the goal is to maximize revenue. Since the network typically relies on best-effort traffic for revenue, we compare the increase in revenue got by the two variations of the best-effort class with the case of accepting demands without any protection.

The paper is organized as follows. Section 2 introduces the network model and explores the choices for a restoration architecture, Section 3 formulates the optimization problem, the solution methodology for solving the combined problem for all classes of demands is presented in Section 4, Section 5 discusses the results and Section 6 concludes the paper.

2. RESTORATION ARCHITECTURE

In this section, we discuss the network model and motivate the restoration architecture adapted for our formulation. This discussion is based on a restoration model proposed in.¹⁰

The optical layer model (shown in Figure 1) consists of nodes interconnected by links which can accommodate multiple fibers. In our formulation, we assume a single fiber model. Each fiber can carry multiple wavelengths. The number of wavelengths which can be carried on a fiber is a technological constraint, which is expected to increase from a few tens to a few hundreds in the coming years. A connection request between nodes is satisfied by establishing a lightpath from the source node to the destination node. A lightpath is an all optical channel which is assigned the same wavelength on all links along the route, to provide a circuit switched connection between the nodes. Each node

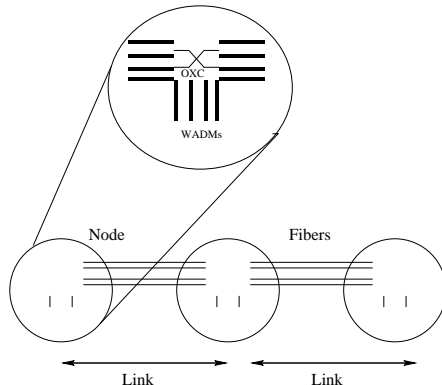


Figure 1. Optical Layer Model

consists of an optical cross-connect (OXC) and optical terminating equipment. This may not always be the case as some nodes may act as through nodes where optical channels are in transit. An optical channel passing through the optical cross-connect may be routed from an input fiber to an output fiber without undergoing O-E-O conversions. In our model we assume that the same wavelength is assigned on all links along the route. So no wavelength translation function is performed in the OXC, all cross-connects are wavelength-selective. An optical channel is terminated by optical terminating equipment such as Wavelength Add/Drop Multiplexers (WADMs). WADMs are used to add or drop selected wavelengths to and from the fiber. So any node can be a source or destination to a connection.

A connection request between a s-d pair is provided a primary route and a backup route. We assume that each path, primary or backup, always accommodates an OAM (operation, administration, and maintenance) channel terminated by the same s-d pair as the path. The restoration model is shown in Figure 2. When a primary path fails, an alarm indication signal is generated by the node which detects the link failure and is transferred over this OAM channel. When the source receives the alarm signal in its OAM channel, it prepares to setup the pre-computed




Figure 2. Restoration Model

backup path and sends messages to the controllers along the backup path to configure the ports accordingly. Since the backup is dedicated, the capacity is assumed to be reserved, so no run time link capacity search needs to be performed. Once the backup path is setup, the destination prepares to receive on that path. There is no restriction in our model for the choice of wavelength on the backup path. It may or may not be the same as the primary path. The tuning time and the associated cost is assumed to be negligible.

Several survivability paradigms have been explored for surviving single link failures in mesh-based networks.^{3,4,7,9,11} They can be classified based on their route computation and execution mechanisms as centralized/distributed, by their re-routing as path/link based, by their computation timing as pre-computed/real time, and their capacity sharing as dedicated/shared. Link based restoration methods re-route disrupted traffic around the failed link, while path based re-routing replaces the whole path between the source and destination of a demand. Link based approach requires the ability to identify a failed link at both ends and makes restoration more difficult when node failures happen. The choice of restoration paths is limited, and thus may use more capacity. The pre-computed approach calculates restoration paths before a failure happens and real time approach does so after the failure occurs. The former approach allows fast restoration as the routes are pre-computed, while the latter approach is slow, as the alternate route is computed after the failure is detected. Centralized restoration methods compute primary and restoration paths for all demands at a central controller where current information is assumed to be available. The routes are then downloaded into each node's route tables. These algorithms are usually path based. They may use pre-computed routes or detect routes at real time. As explained above, since this step needs to identify failure, ascertain the remaining topology and capacity and then find the best alternate route for the affected demands, the procedure is very slow. Given the importance of restoration speed and potential difficulty in fast failure isolation in optical networks, this approach is not very attractive. Centralized schemes which involve pre-computed routes are more conducive for practical implementations. However, maintaining up-to-date information requires frequent communications between the nodes and the central controller. This overhead becomes a potential problems as the network size grows. Distributed methods may involve pre-computed tables of routes, and discovers capacity in real time. Real time capacity discovery is slow and the capacity utilization may be inefficient. Distributed pre-computation of restoration route is an attractive approach. Capacity sharing among the primary and restoration paths can be dedicated or shared. The dedicated technique uses 1:1 protection where each primary path has a corresponding restoration path. In the shared case several primaries can share the same backup path as long as the primaries are node and link disjoint. This scheme is sometimes called backup multiplexing technique.³ These paradigms serve as a good framework for analyzing the different design methodologies, as each design methodology uses a restoration model which is a combination of the different paradigms just described.

2.1. Restoration Model

We consider a centralized, pre-computed, shared restoration model, with 100% restoration guarantee against any single node or link failure. This means that primary and restoration paths are allocated the same capacity, and are

node and link disjoint. We employ backup multiplexing to increase the wavelength utilization. We have the following constraints in our restoration model.

- Number of connections (lightpath) on each link is bounded
- Levels of protection
 - Full protection: Every demand is assigned a primary and a backup path
 - No protection: Every demand is assigned only a primary path
 - Best-effort protection: (i) Every demand is assigned a primary path. A backup if resources are available
(ii) Accept as many demands as possible with or without backup.
- No backups are admitted without a primary i.e., for every node pair, the number of primaries accepted is equal to or greater than the backups.
- Primary path wavelength restrictions: Only one primary path can use a wavelength λ on link l , no restoration path can use the same λ on link l
- Backup multiplexing constraint: Many restoration paths can share a wavelength λ on link l iff their corresponding primary paths are link and node disjoint
- Restoration path wavelength usage indicator to identify if a wavelength λ is used by some restoration route (i, r) traversing link l
- Primary and backup paths for a given demand should be node and link disjoint.

3. FORMULATION OF THE OPTIMIZATION PROBLEM

In this section, we develop an ILP formulation for a service differentiation model based on lightpath protection and maximize the total revenue generated by the network. There may be a penalty for disrupting currently working connections and we capture this service disruption aspect into our problem formulation. Based on the service classes, traffic in the network falls into one of the three classes viz., full protection, no protection and best-effort. In the full protection class, every demand is assigned a primary and a backup path. We provide 100% restoration guarantee for protected connections against any single node or link failure. This implies that primary and restoration paths are allocated the same capacity, and are node and link disjoint. In the no protection class, every demand is assigned only a primary path and in the best-effort class, (i) every demand is assigned a primary path, a backup path is also provided, if resources are available (ii) accept as many demands as possible with or without backup.

The following information is assumed to be given: the network topology, a demand matrix consisting of the new connections to be established for each class, and the set of current working connections. We also assume that a set of alternate routes between each node-pair is pre-computed and given. Each route between every s-d pair is viewed as W wavelength continuous paths, one path corresponding to every wavelength and therefore, we do not have an explicit constraint for wavelength continuity. Information regarding whether any two given paths are link and node disjoint (except the source and destination nodes) are also assumed to be given. The ILP solution determines the primary and backup paths for the demand set and hence the routing and wavelength assignment.

3.1. Notation

The network topology is represented as a directed graph $G(N, L)$ with N nodes and L links with W wavelengths on each link. We also assume that two alternate paths, which are node and link disjoint, for each s-d pair, are used to provide survivability. It has been shown in¹² that two alternate paths are generally sufficient to achieve good performance. The following notations are used.

- $n = 1, 2, \dots, N$: Number assigned to each node in the network
- $l = 1, 2, \dots, L$: Number assigned to each link in the network
- $\lambda = 1, 2, \dots, W$: Number assigned to each wavelength in a fiber

- $i, j = 1, 2, \dots, N(N-1)$: Number assigned each s-d pair
- $p, r = 1, 2, \dots, KW$: Number assigned to a path for each s-d pair
- (i, p) : Refers to the p th path for s-d pair i

The following cost parameters are employed.

- C_l : Cost of using a link l (data)
- C_w : Cost of disrupting a currently working path (data)
- C_{ND} : Cost of a primary path (data)
- C_D : Cost of a backup path (data)

Information regarding whether two given paths are link and node disjoint

- $I_{(i,p),(j,r)}$ takes a one if (i, p) and (j, r) have at least one link in common, zero if the paths are disjoint (data).
if $i = j$ then $p \neq q$.

The following notations are used for path related information

- $\delta^{i,p}$: Path indicator which takes one if (i, p) is chosen as a primary path, zero otherwise (variable)
- $\nu^{i,r}$: Path indicator which takes one if (i, r) is chosen as a restoration path, zero otherwise (variable)
- $\delta_l^{i,p}$: Link indicator which takes one if link l is used in path (i, p) , zero otherwise (data)
- $\psi_\lambda^{i,p}$: Wavelength indicator which takes one if wavelength λ is used by the path (i, p) , zero otherwise (data)
- $g_{l,\lambda}$: Restoration wavelength indicator takes a one if wavelength λ is used by some restoration route (i, r) that traverses link l (variable)

The following notations are used for currently working (primary) path related information. We are only interested in the primary route of the currently working connection as the restoration paths can be reassigned.

- $\chi^{i,p}$: Path indicator which takes one if (i, p) is a currently working primary path, zero otherwise (data)

The following notation is used to denote the demand in terms of lightpath requests for every node pair

- d_i : Demand for node pair i , in terms of number of lightpath request.

3.2. Problem Formulation

Objective: The objective is to maximize the revenue, i.e., to accept as many demands as possible. Each demand translates into a primary and a backup path for full protection class, or only primary for no protection class, and either only primary or both primary and backup for best-effort class depending on the bandwidth available. The first term in Equation 1 denotes the revenue generated from primary paths, and the second term denotes the revenue from backup paths. The last term indicates that if a currently working connection ($\chi^{i,p} = 1$) is not picked in the final solution ($\delta^{i,p} = 0$), then the objective value is penalized by subtracting a cost C_w to it.

Maximize

$$\sum_{i=1}^{N(N-1)} \sum_{p=1}^{KW} \delta^{i,p} C_{ND} + \sum_{i=1}^{N(N-1)} \sum_{p=1}^{KW} \nu^{i,p} C_D - \sum_{i=1}^{N(N-1)} \sum_{p=1}^{KW} \chi^{i,p} (1 - \delta^{i,p}) C_w \quad (1)$$

Restoration path wavelength usage indicator constraint:

$$X_{l,\lambda} = \sum_{i=1}^{N(N-1)} \sum_{r=1}^{KW} \nu^{i,r} \delta_l^{i,r} \psi_\lambda^{i,r} \quad (2)$$

$$g_{l,\lambda} \leq X_{l,\lambda} \quad (3)$$

$$N(N-1)WK g_{l,\lambda} \geq X_{l,\lambda} \quad (4)$$

$$1 \leq l \leq L, 1 \leq \lambda \leq W, X_{l,\lambda} \geq 0$$

Link capacity constraint:

$$\sum_{i=1}^{N(N-1)} \sum_{p=1}^{KW} \delta_l^{i,p} \delta_l^{i,p} + \sum_{\lambda=1}^W g_{l,\lambda} \leq W \quad 1 \leq l \leq L \quad (5)$$

Demand constraints for each node pair

- Full protection: Every demand is assigned a primary and a backup path

$$\sum_{p=1}^{KW} \delta^{i,p} = d_i \quad 1 \leq i \leq N(N-1) \quad (6)$$

$$\sum_{r=1}^{KW} \nu^{i,r} = d_i \quad 1 \leq i \leq N(N-1) \quad (7)$$

- No protection: Every demand is assigned only a primary path

$$\sum_{p=1}^{KW} \delta^{i,p} = d_i \quad 1 \leq i \leq N(N-1) \quad (8)$$

- Best-effort protection: (i) Every demand is assigned a primary path. A backup if resources are available

$$\sum_{p=1}^{KW} \delta^{i,p} = d_i \quad 1 \leq i \leq N(N-1) \quad (9)$$

$$\sum_{r=1}^{KW} \nu^{i,r} \leq d_i \quad 1 \leq i \leq N(N-1) \quad (10)$$

- Best-effort protection: (ii) Accept as many demands as possible with or without backup

$$\sum_{p=1}^{KW} \delta^{i,p} \leq d_i \quad 1 \leq i \leq N(N-1) \quad (11)$$

$$\sum_{r=1}^{KW} \nu^{i,r} \leq d_i \quad 1 \leq i \leq N(N-1) \quad (12)$$

Constraint to avoid backups without a primary: No backups are admitted without a primary i.e., for every node pair, the number of primaries accepted is equal to or greater than the backups. This constraint is required to ensure that when best-effort class demands are admitted, the ILP does not admit a more backups than primaries.

$$\sum_{p=1}^{KW} \delta^{i,p} - \sum_{r=1}^{KW} \nu^{i,r} \geq 0 \quad 1 \leq i \leq N(N-1) \quad (13)$$

Primary path wavelength usage constraint:

$$\sum_{i=1}^{N(N-1)} \sum_{p=1}^{KW} \delta^{i,p} \delta_l^{i,p} \psi_\lambda^{i,p} + g_{l,\lambda} \leq 1 \quad (14)$$

$$1 \leq l \leq L, 1 \leq \lambda \leq W$$

Backup multiplexing constraint:

$$(\nu^{i,p} \delta_l^{i,p} \psi_\lambda^{i,p} + \nu^{j,r} \delta_l^{j,r} \psi_\lambda^{j,r}) I_{(i,\bar{p}), (j,\bar{r})} \leq 1 \quad (15)$$

$$1 \leq i, j \leq N(N-1), 1 \leq p, \bar{p}, r, \bar{r} \leq KW$$

Constraint for topological diversity of primary and backup paths:

$$\sum_{p=1}^W \delta^{i,p} = \sum_{r=W+1}^{KW} \nu^{i,r} \quad 1 \leq i \leq N(N-1) \quad (16)$$

$$\sum_{p=W+1}^{KW} \delta^{i,p} = \sum_{r=1}^W \nu^{i,r} \quad 1 \leq i \leq N(N-1) \quad (17)$$

In the objective function, the last term indicates that if a currently working connection ($\chi^{i,p} = 1$) is not picked in the final solution ($\delta^{i,p} = 0$), then the cost C_w is subtracted from the objective and since the objective is to maximize, it ensures that service is not disrupted unless otherwise to increase revenue. The choice of C_w offers a lot of flexibility to the network provider. Although the network would like to avoid service disruption to all connections, there may be some customers who are willing to pay more and do not wish to be disturbed. This can be accommodated by modifying C_w to be path specific ($c_w^{i,p}$) and setting a higher cost for disrupting such connections.

In the next section, we describe a solution methodology to extract a feasible solution for the combined problem.

4. SOLUTION METHODOLOGY

In this section, we describe the solution methodology for solving the combined problem for all classes of demands.

The number of variables $\delta^{i,p}$ and $\nu^{i,p}$ grow rapidly with network size. This effect is more pronounced with an increase in the number of wavelengths. For a network of size $N = 14$, $W = 32$ and $K = 2$, considering one class of demands, there are $K * W = 2 * 32$ instances of each variable for every node pair. Since there are $N * (N - 1) = 182$ node pairs, we have 11,648 $\delta^{i,p}$ variables and 11,648 $\nu^{i,p}$ variables. The number of equations will be roughly 125 million ($11,648^2$). For the combined problem to be solved, with multiple classes of demands, the following changes are made to the demand constraints in the formulation. Let $\delta_1^{i,p}$, $\delta_2^{i,p}$, $\delta_3^{i,p}$ denote variables for classes full protection, no protection, and best-effort variation 2 respectively. We consider best-effort variation 2 class for demonstrating our solution methodology and results. The same arguments hold for best-effort variation 1. Replace $\delta^{i,p}$ by $\delta_1^{i,p}$ in equations (6),(7), $\delta^{i,p}$ by $\delta_2^{i,p}$ in equations (8), and $\delta^{i,p}$ by $\delta_3^{i,p}$ in equations (11),(12). The number of variables increases making the problem harder to solve.

Multistage Approach We propose a three stage approach to solving the problem. At each stage, one instance of the problem is solved, for one of the classes, and the result is used in successive stages. If the problems are solved independently, the resulting solution may be infeasible, as the same primary path might be chosen for different demands from different classes. In order to avoid infeasibility, we feed the information about one stage to the next

through the $\chi^{i,p}$ variable. Typically this variable is used to feed information about existing paths to avoid service disruption. We exploit this aspect of our formulation by feeding the solution of one stage to the next. In our case, we first solve the problem for the demands which require full protection. The result of this stage, which is a set of primary paths, is fed to the next stage. We then solve the problem for demands from full protection and no protection classes. The result of the second stage, which is a set of primary paths for all demands from the two classes, is fed to the third stage. The third stage solves the problem for all classes.

Effect of C_w The effect of the solution depends on the value of C_w , the higher the value, more the guarantee that the path will remain unaffected. This value is set to be some β times the cost of primary paths. Typically the value of β is set to 2 or 3. This implies that the increase in the objective value for picking β primary paths is lost for disrupting one existing path.

Complexity To understand the reduction in complexity at each stage, let us first examine the first stage of the solution. Since we are interested only in the primary paths for the full protection class in the first stage (backups will be picked in the last stage of the solution), all the $\nu^{i,p}$ variables can be assumed to be zero, which results in a reduction of $N * (N - 1) * K * W$ variables. The same explanation can be applied for stage two. The last stage complexity depends on the value of C_w . If the value is high compared to the cost of a primary, then the solution from the previous stage will remain unaffected. This reduces the number of combinations to be explored thus resulting in a decrease in the computational complexity of the ILP solution. The penalty paid is a sub-optimal arrangement of paths. Optimal arrangement of paths is possible only if all the demands from all three classes are solved simultaneously without any restrictions. The deviation from optimality is not discussed here as it is beyond the scope of the paper. Our goal is to extract a feasible solution.

Finally, in order to ensure that the backup paths for full protection class demands are picked at the last stage, equation (12) is modified as follows.

$$\sum_{r=1}^{KW} \nu^{i,r} \geq dl_i \quad 1 \leq i \leq N(N-1) \quad (18)$$

where dl_i is the number of full protection class demands for node pair i . Equation (13) ensures that the number of backups is less than or equal to the number of primaries, thus enforcing a bound on the modified equation above.

5. RESULTS

We use CPLEX Linear Optimizer 5.0.1¹³ to solve the ILPs. The combined routing and wavelength assignment problem is known to be NP-Complete¹⁴ and the problem addressed in this paper is expected to be NP-Complete as well. The major difficulty in using the above formulation for larger and more practical networks arises due to the combinatorial nature of the ILP. Although several algorithmic approaches exist for solving ILP problems, it should be observed that as the number of 0-1 variables increases, the computational complexity grows exponentially. The ILP problem has been shown to belong to the class of NP-complete problems.¹⁵ Various decomposition techniques, based on lagrangean relaxation^{7,16} and LP relaxation techniques,^{15,17} can be employed to reduce the computational complexity of the original problem.

We demonstrate the effectiveness of our formulation on the 14 node 21 link NSFNET topology (shown in Figure 3) with one fiber per link and 10 wavelengths per fiber. Consider the following cost relationship between the primary and backup paths. $C_D = \alpha * C_{ND}$, where $0 \leq \alpha \leq 1$. The total revenue is calculated as $(\#totalprimaries * C_{ND} + \#totalbackups * \alpha * C_{ND})$ costunits (cu). The network relies on the best-effort class to increase revenue. We compare the increase in revenue got by the two variations of the best-effort class with a base case of accepting all connections without any protection. We show results for $C_{ND} = 500cu$ and for two values of $\alpha = \{1, 0.5\}$. The results for various demand sets are shown in Table 1. For particular instances of demands, we see that the best-effort variation 1 results in a 67% gain in revenue and variation 2 achieves an additional 6% gain, for $\alpha = 1$. The cases are compared to the revenue generated by accepting all demands without protection $(\#primaries * C_{ND})$. For example, consider the case of 48 demands for $\alpha = 1$. The base case accepting all demands without any protection results in $48 * C_{ND} = 24,000cu$. The total revenue for variation 1 is $48 * C_{ND} + 32 * C_D = 40,000cu$, which is a 66.7% gain. The revenue for variation 2 is $44 * C_{ND} + 39 * C_D = 41,500cu$, which is a 72.9% gain. Although, both schemes employ




Figure 3. The 14 node 21 link NSFNET

Demand	$\alpha = 1$					$\alpha = 0.5$				
	Best-effort 1		Best-effort 2			Best-effort 1		Best-effort 2		
	Primary	Backups	Primary	Backups	Rejected	Primary	Backups	Primary	Backups	Rejected
12	12	8	12	8	0	12	8	12	8	0
20	20	16	20	16	0	20	16	20	16	0
24	24	12	21	18	3	24	12	21	18	3
32	32	20	28	27	4	32	20	29	26	3
36	36	22	33	28	3	36	22	33	28	3
44	44	30	41	36	3	44	30	41	36	3
48	48	32	44	39	4	48	32	46	36	2

Table 1. Increase in revenue for the two variations of best-effort class

backup multiplexing, the best-effort variation 1 has no choice but to choose all the primary paths and then tries to accommodate backups and so is restricted. There is no such restriction on best-effort variation 2 as it accommodates as many demands as possible with or without backups. Best-effort variation 2 better exploits the backup resource consumption by effectively multiplexing more connections on the same wavelength, thus accepting more connections and generating a slight increase in revenue.

We now demonstrate our multistage solution methodology on the NSFNET topology. We consider a demand set comprising of 48 demands with 12 demands in full protection class, 12 demands in no protection class and 24 demands in the best-effort (variation 2) class, distributed uniformly across four node pairs. The cost values used are $C_{ND} = 500$, $C_D = 500(\alpha = 1)$, $C_w = 500(\beta = 1)$. In the first stage, the problem is solved for full protection demands. We assume that there are no currently working connections. Thus, the value of $\chi^{i,p}$ for all the node pairs is zero. The ILP determined a feasible solution, which is a set of paths, with a route and wavelength associated with each of them, for all the 12 demands in the full protection class. This set of paths, is fed into the second stage through $\chi^{i,p}$ variables. The problem is solved for full protection and no protection classes. The 12 paths chosen for full protection class are assumed to be working paths in this stage. The ILP determined a feasible solution for all 24 demands with an objective value of 11,500. Although, the objective value is of no relevance as long as we know the number of primary and backups selected, it is interesting to see how the ILP handles service disruption. Since the ILP determined a feasible solution for all 24 demands, the objective value is expected to be 12,000, but the value got is 11,500 ($24 * C_{ND} - 1 * C_w$). This was due to the fact that one of the full protection demand's primary path was reassigned. The objective value incurred a penalty for disturbing the connection. Thus, by appropriately choosing C_w , as explained in Section 4, this aspect of the formulation can be used to avoid service disruptions to existing connections in the network. This set of primary paths is then fed to the third stage. The value of C_w is set to 1500 ($\beta = 3$) for this stage to ensure that the paths chosen for the full protection and no protection demands are retained in the final solution. Equation (18) ensures that backups for all the full protection class demands are chosen in the last stage. The final solution at the end of the third stage is shown in Table 2. The demands rejected are those belonging to the best-effort class. The total revenue generated for provisioning the complete demand set

Node pair	Class 1	Class 2	Class 3	Primary paths	Backup paths
1	3	3	6	10	10
2	3	3	6	11	10
3	3	3	6	12	8
4	3	3	6	12	8
				45	36

Table 2. Solution at the end of the third stage

for all classes is $45 * C_{ND} + 36 * C_D = 58,500\text{cu}$.

6. CONCLUSIONS

Service availability is an indispensable requirement for many current and future applications over the Internet and hence has to be addressed as part of the optical QoS service model. We considered a service differentiation model based on lightpath protection. Based on the service classes, traffic in the network falls into one of the three classes viz., full protection, no protection and best-effort. We captured service differentiation based on lightpath protection, into a integer linear formulation for revenue maximization in wavelength routed backbone networks. Our approach also captures the service disruption aspect into the problem formulation, as there may be a penalty for disrupting currently working connections. We exploited this aspect to propose a solution methodology for extracting a feasible solution, as the combined problem for solving all demands together can be quite complex. We also showed that for particular instances of demands, best-effort variation 1 results in a 67% gain and variation 2 achieves an additional 6% gain in revenue, for $\alpha = 1$. To adapt the formulation for larger networks, we propose to use heuristics and decomposition techniques, based on relaxation techniques, to significantly reduce the computational complexity of the problem.

ACKNOWLEDGMENTS

This work is funded in part by NSF grant ANI 9973102 and by Defense Advance Research Projects Agency under Contract N66001-00-1-8949

REFERENCES

1. T. D. Ndousse and N. Golmie, "Differentiated Optical Services: Quality of Optical Service Model for WDM Networks," *SPIE All Optical Networking: Architecture, Control and Management Issues*, September 1999, pp 79-88.
2. O. Gerstel and R. Ramaswami, "Optical Layer Survivability: A Services Perspective," *IEEE Communications Magazine*, March 2000, Vol.38 No. 3, pp 104-113.
3. G. Mohan and A. K. Somani, "Routing Dependable Connections With Specified Failure Restoration Guarantees in WDM Networks," *IEEE INFOCOM*, 2000.
4. Y. Miyao and H. Saito, "Optimal Design and Evaluation of Survivable WDM Transport Networks," *IEEE JSAC*, Vol. 16, No. 7, September 1998, pp 1190-1198.
5. B.V. Caenegem, W.V. Parys, f. De Turck, and P.M. Demeester, "Dimensioning of Survivable WDM Networks," *IEEE JSAC*, Vol. 16, No. 7, September 1998, pp 1146-1157.
6. O. Crochat and J.Y. Boudec, "Design Protection for WDM Optical Networks," *IEEE JSAC*, Vol. 16, No. 7, September 1998, pp 1158-1165.
7. B.T. Doshi, S. Dravida, P. Harshavardhana, O. Hauser, and Y. Wang, "Optical Network Design and Restoration," *Bell Labs Technical Journal*, January-March 1999, pp 58-83.
8. M. Alanyali and E. Ayanoglu, "Provisioning Algorithms for WDM Optical Networks," *IEEE INFOCOM*, 1998.
9. S. Ramamurthy and B. Mukherjee, "Survivable WDM Mesh Networks, Part I - Protection," *IEEE INFOCOM*, 1999, pp 744-751.
10. M. Sridharan, M.V. Salapaka, and A.K. Somani, "Operating Mesh-Survivable WDM Transport Networks," to appear in *Intl. Symposium on SPIE, Terabit Optical Networking: Architecture, Control, and Management Issues* 2000

11. T.H. Wu, "*Fiber Network Service Survivability*," Norwood, MA: Artech House, 1992.
12. L. Li and A.K. Somani, "Dynamic Wavelength Routing Using Congestion and Neighbourhood Information," *IEEE/ACM Transactions on Networking*, October 1999.
13. <http://www.cplex.com>
14. I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath Communications: An approach to high-bandwidth optical WAN's," *IEEE Transactions on Communications*, Vol. 40, No. 7, July 1992, pp 1171-1182.
15. C.A. Floudas, "*Nonlinear and Mixed-Integer Optimization: Fundamentals and Applications*," Oxford University Press, 1995
16. M.L. Fischer, "The Lagrangean Relaxation Method for Solving Integer Programming Problems," *Management Science*, Vol. 27, No. 1, 1981, pp 1-18.
17. D. Banerjee and B. Mukherjee, "A Practical Approach for Routing and Wavelength Assignment in Large Wavelength-Routed Optical Networks," *IEEE JSAC*, Vol. 14, No. 5, June 1996, pp 903-908.
18. M. Sridharan, M.V. Salapaka, and A.K. Somani, "A Practical Approach to Operating Large Mesh-Survivable WDM Transport Networks," , *under preparation*.