# Partial Path Protection for WDM Networks: End-to-End Recovery Using Local Failure Information

Hungjen Wang, Eytan Modiano, Muriel Médard

*Abstract*— **Path protection and link protection schemes are the main means of protecting wavelength-division multiplexed (WDM) networks from the losses caused by a link failure such as a fiber cut. We propose a new protection scheme, which we term partial path protection (PPP), to select end-to-end backup paths using local information about network failures. PPP designates a different restoration path for every link failure of every primary path. PPP allows the re-use of operational segments of the original primary path in the protection path. A novel approach used in this paper is that of a dynamic call-by-call model with blocking probability as the performance metric. This is in contrast with traditional approaches to restoration, which consider capacity-efficiency for batch call arrivals. Since optimizing the blocking probability is a large dynamic optimization problem, we present two heuristics for implementing PPP. We show that a simple method based on shortest path routing for which primary paths are selected first is more effective than a greedy approach that minimizes, for each call arrival, the number of wavelengths used by the primary and backup path jointly.**

## I. Introduction

A wide range of protection schemes for WDM networks have been investigated [1], [2], [3], [4], [5], [7], [9], [10], [12], [13], [14], [16], [17]. Among them, path protection and link protection have attracted the most attention [1], [10], [12], [13], [16]. Path protection requires the protection path of a request to be completely link-disjoint from the corresponding primary path, while the link protection scheme reroutes all affected requests over a set of replacement paths between the two nodes terminating the failed link. Primary capacity cannot be shared, but protection capacity can be shared as long as a single link failure does not activate more than one protection channel along any wavelength on any link. In general, path protection is more capacity efficient than link protection [12].

In this paper, we present a new protection scheme, the *partial path protection scheme* (PPP). In this scheme, the network identifies a specific protection path for each link along a considered primary path. Thus, similarly to the path protection scheme, the partial path protection scheme assigns "end-to-end" protection paths to primary paths. However, in PPP, one single protection path protects only one specific link failure on one primary path, instead of the whole primary path in path protection.

We consider a dynamic call-by-call system with random arrivals. Other research in the area of restoration efficiency has generally considered a batch model. That model is reasonable when call demands are known in advance. However, static batch models do not allow for dynamic provisioning of primary and protection paths in the network. Our call-by-call dynamic model is well suited to dynamic allocation of capacity for primary and protection paths. In our call-by-call model, every new call establishes its primary and protection paths according to the traffic already present in the network when the call arrives. Given the dynamic and probabilistic nature of our model, we take the call blocking probability to be the performance metric for our schemes, rather than traditional capacity efficiency metrics.

In order to optimize call blocking probability for selecting paths using PPP over some time horizon, we would have to solve a dynamic optimization problem. The extremely large state space of a dynamic program over a reasonable network and time horizon renders such an approach impractical. The complexity of a dynamic programming approach prompts us to consider two heuristics for implementing PPP.

The first heuristic is a *greedy* approach that, for each call arrival, the system uses the fewest previously unused wavelengths to establish the primary and protection paths jointly. Wavelengths already used for protection paths can be used for new protection paths as long as a single link failure does not entail the activation of more than one protection path on any wavelength on any link. The problem formulation is an integer linear program (ILP) [8], a common approach to network routing [4], [9], [10], [12], [14].

The second heuristic first selects the primary path, using a shortest path route. It then selects the protection paths using a shortest path algorithm in which wavelengths already assigned for protection can be used at no cost. We term the

This work was supported by DARPA under the Next Generation Internet Initiative.

whole of the second heuristic, involving the choice of primary and of protection paths, the *shortest path* approach (SP).

We show that the SP approach is not only significantly simpler computationally than the greedy approach, but also more effective in terms of blocking probability. This result may seem surprising at first. However, since protection paths can share bandwidth, while primary paths cannot, it is reasonable to select the most economical primary first, as done by SP, rather than consider primary and protection bandwidth jointly, as done by the greedy algorithm. The SP approach, by selecting the primary path first, in effect prioritizes the efficient use of primary path resources over protection resources. The greedy approach seeks to minimize the total use of new wavelengths by primary and backup paths jointly. However, in a dynamic system, the efficient use of protection bandwidth is not as important as the efficient use of primary bandwidth, since *in the future*, protection bandwidth has a high likelihood of being shared, whereas primary bandwidth cannot be shared. The fact that SP performs better than the greedy approach highlights the significant difference between a dynamic call-by-call model and a static batch system.

The main contributions of our paper are the introduction of the PPP method for establishing protection paths, the introduction of the greedy and SP approaches for implementing PPP and path protection and the use of a dynamic call-by-call model for protection. In the next section, we present PPP and related background. In Section III, we present the greedy and SP approaches to implementing PPP and path protection. In Section IV, we present simulation results over several backbone networks to compare the performance, in terms of call blocking probability, of path protection and PPP using SP and the greedy algorithm. We present our conclusions and directions for further research in Section V.

## II. Protection schemes

In this section, we introduce PPP and compare it to path and link protection. We also discuss the issue of protection resource sharing.

### A. Path protection and link protection schemes

There are two prevailing protection schemes to guard against link failure, path protection and link protection schemes. Path protection, as illustrated in Fig. 1, reserves network resources for a single protection path in addition to the primary path. Since it is impossible to foresee which link on the primary path will fail, the system allocates a protection path, which is completely link-disjoint from the primary path. The primary path therefore shares no common link with its associated protection path. When a link fails, the source and destination nodes of a call on the failed link are informed of the failure, and the communication is switched to the protection path.

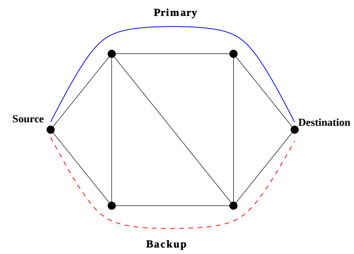


Fig. 1. Path protection scheme

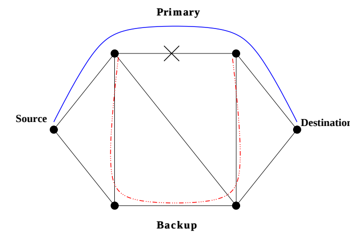


Fig. 2. Link protection scheme

Link protection, as shown in Fig. 2, reroutes all the connections on the failed link around it. When accepting a call request, the link protection scheme will reserve the network resource for the associated protection path. Note that the protection path connects the two nodes adjacent to the failed link. When a link failure occurs, the node adjacent to and upstream of the failed link immediately redirects the traffic along the predetermined protection path to the node on the other end of the failed link to restores transmission.

### B. Partial path protection scheme (PPP)

In PPP, the system reserves the protection resources while setting up a primary path. The major difference with path protection scheme is that the system now specifies a specific protection path for *each link* along the primary path. Thus, each protection path, rather than being associated with a single path as for path protection, or a single link as for link protection, is associated with a link/primary path pair. In the event of a link failure, the call is rerouted along the protection path corresponding to the failed link. For example, in Fig. 3, a call with source node $1$ and sink node $4$ has a primary path $1-2-3-5-4$. As illustrated in Table I, the system applying PPP takes $1-6-2-3-5-4$ as the protection path against the failure of link $(1,2)$. Similarly, the network assigns $1-2-5-4$ to protect against the failures of links $(2,3)$ and $(3,5)$, and finally, $1-2-3-4$ to protect against the failure of $(5,4)$. Each of these protection paths needs only to be link-disjoint only from the link it protects.

Comparing PPP with path protection, we see that the former is more flexible than the latter. Indeed, any path protection scheme is a valid PPP, whereas the reverse does not hold. We expect, therefore, that PPP will enhance our ability to provide protection over traditional end-to-end path protection. To illustrate this fact, consider Fig. 3. By applying traditional end-to-end path protection, the net-

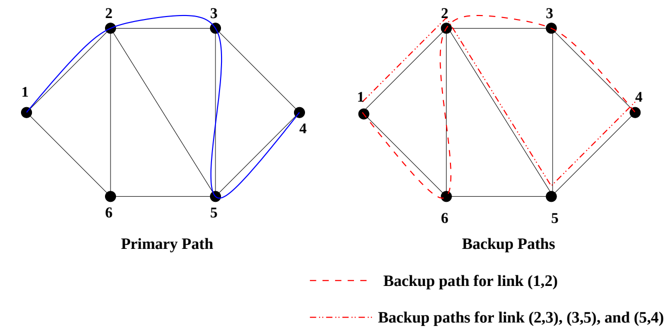


Fig. 3. An example for partial path protection scheme

| Link on Primary Path 1-2-3-5-4 | Corresponding Protection Path |
|---|---|
| (1,2) | $1-6-2-3-5-4$ |
| (2,3) | $1-2-5-4$ |
| (3,5) | $1-2-5-4$ |
| (5,4) | $1-2-3-4$ |

TABLE I

ILLUSTRATION OF PROTECTION PATHS FOR THE PRIMARY PATH IN FIG. 3

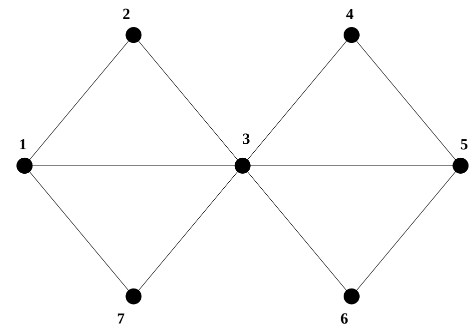


Fig. 4. An example network for illustrating the partial path protection and path protection schemes in protection sharing

| | SD Pair | Primary Path | Protection Path (protected link) | Total number of occupied $\lambda$'s |
|---|---|---|---|---|
| Path Protection Scheme | (1,5) | 1-3-5 | 1-2-3-4-5 (1-3) | 6 |
| | | | 1-2-3-4-5 (3-5) | |
| | (5,4) | 5-4 | 5-3-4 (5-4) | 8 (share (3,4)) |
| Partial Path Protection Scheme | (1,5) | 1-3-5 | 1-2-3-5 (1-3) | 6 |
| | | | 1-3-4-5 (3-5) | |
| | (5,4) | 5-4 | 5-3-4 (5-4) | 8 (share (3,4)) |

TABLE II

RESOURCE ALLOCATION FOR SOURCE DESTINATION PAIR $(1,5)$ AND $(5,4)$ OF THE NETWORK IN FIG. 4

work cannot find a protection path for the primary path shown. However, by applying PPP, we can provide protection service to the primary path. Since link protection schemes generally have a worse performance than path protection, we do not seek to compare PPP with link protection but only with traditional path protection.

### C. *Protection sharing*

For path protection, a system can allow primary paths with no link in common to share protection bandwidth against a link failure, because we assume a single link failure can occur at a time. In addition to this type of bandwidth sharing, PPP allows a protection path to share bandwidth with portions of the primary path that remain operational after link failure. The following example illustrates the different levels of protection sharing for path protection and PPP.

*Example 1:* Consider the network in Fig. 4 and assume the network is initially empty. The network now serves two call requests, $(1,5)$ and $(5,4)$, in sequence. Table II shows the resource assignments for primary and protection paths under the path protection and the PPP respectively. As shown in Table II, the two primary paths, $1-3-5$ and $5-4$, are completely link-disjoint from each other. By exercising protection sharing, the system reserves only one wavelength for protection on link $(3,4)$, thus improving the network resource utilization.

Example 1 illustrates the difference between path protection and PPP. Though the total number of occupied wavelengths to support the two requests is the same in both schemes, the protection wavelengths are used differently for path protection and for PPP. Consider, for example, link $(1,2)$. In the path protection scheme, a wavelength on this link is assigned to protect link $(1,3)$ and $(3,5)$, while in PPP, the wavelength protects only the link $(1,3)$. Hence, under PPP, this wavelength can be shared by a future call whose primary path includes link $(3,5)$, but cannot be shared by using path protection.

## III. PATH ASSIGNMENT APPROACHES

We consider two approaches to implement path protection and PPP. In the case of path protection, there is a single protection path per primary path. In the case of PPP, there is a protection path for every link in the primary path. The first approach we consider is the greedy approach, in which a system simultaneously allocates primary and protection paths to a new call by solving an ILP to minimize the use of previously unused wavelengths. The other approach is the SP approach. In SP, the system first assigns the shortest path between a source and destination as a request's primary path. After having assigned the primary path, the system assigns each protection path (a single one in the case of path protection and possibly several ones in the case of PPP) using wavelengths that are free or shared with other protection paths. The cost of using each previously free wavelength is 1 and the cost of using a wavelength shared with other protection paths is 0. In the case of path protection, the is no sharing of wavelengths with the primary path. In the case of PPP, a path protecting a primary path can re-use at no cost a wavelength over an unfailed link in that primary path.

### A. Greedy approach

To maximize network resource utilization, it is natural to seek to minimize the use of new resources for every call. We call this approach the greedy approach. We formulate the ILPs to realize path protection and PPP using the greedy approach. We first introduce the ILP formulation for path protection.

To begin with, we introduce the variables used in the formulation. Let

$L$ denote the set of all possible links,

$S$ denote the source node,

$D$ denote the destination node,

$$c_{ij} = \begin{cases} 1, & \text{if at least one wavelength is available on} \\ & \text{link } (i,j) \in L, \\ \infty, & \text{otherwise,} \end{cases}$$

$$d_{ij}^{lk} = \begin{cases} 0, & \text{if at least one wavelength on link } (l,k) \\ & \text{other than } (i,j) \text{ is already reserved to} \\ & \text{protect links other than } (i,j), \\ 1, & \text{else if at least one wavelength is available} \\ & \text{on link } (l,k) \in L, \\ \infty, & \text{otherwise,} \end{cases}$$

$$x_{ij} = \begin{cases} 1, & \text{if the primary path rests on an} \\ & \text{available wavelength in link } (i,j), \\ 0, & \text{otherwise,} \end{cases}$$

$$y_{ij} = \begin{cases} 1, & \text{if the system reserves a wavelength} \\ & \text{in link } (i,j) \text{ for protection,} \\ 0, & \text{otherwise,} \end{cases}$$

$$v_{ij}^{lk} = \begin{cases} 1, & \text{if a wavelength on } (l,k) \text{ is reserved to} \\ & \text{protect its associated primary path on } (i,j), \\ 0, & \text{otherwise.} \end{cases}$$

Note that, since we have no advance information about where the primary path will be placed, we need the variable $d$ to indicate which links have wavelengths available to protect some specific link on which the primary path may reside. Furthermore, we also need the variable $v$ to indicate the assignment of wavelengths to protection. The formulation of the ILP for a random call arrival is detailed below.

$$\text{Minimize} \qquad \sum_{(i,j)\in L} c_{ij}x_{ij} + \sum_{(i,j)\in L} y_{ij} \tag{1}$$

Eq.(1) represents the objective function, where $c$ indicates whether a link has a free wavelength, $x$ indicates the network resources for primary transmission and $y$ indicates the network resources reserved for protection. Notice that, in the ILP, the primary path and the protection path are considered concurrently. We next consider the constraint set.

$$\sum_{(S,j)\in L} x_{Sj} - \sum_{(j,S)\in L} x_{jS} = 1, \tag{2}$$

$$\sum_{(D,j)\in L} x_{Dj} - \sum_{(j,D)\in L} x_{jD} = -1, \tag{3}$$

$$\sum_{(i,j)\in L} x_{ij} - \sum_{(j,i)\in L} x_{ji} = 0, \quad \forall i \neq S, D, \tag{4}$$

$$\sum_{(S,l)\in L} v_{ij}^{Sl} - \sum_{(l,S)\in L} v_{ij}^{lS} \geq x_{ij}, \quad \forall (i,j) \in L, \tag{5}$$

$$\sum_{(l,D)\in L} v_{ij}^{lD} - \sum_{(D,l)\in L} v_{ij}^{Dl} \geq x_{ij}, \quad \forall (i,j) \in L, \tag{6}$$

$$\sum_{(l,k)\in L} v_{ij}^{lk} - \sum_{(k,l)\in L} v_{ij}^{kl} = 0, \quad \forall (i,j) \in L, \quad \forall k \neq S, k \neq D, \tag{7}$$

Eq.(2) to Eq.(4) provide the flow conservation for the primary path. Similarly, Eq.(5) to Eq.(7) give the flow conservation for the protection path. Note that Eq.(5) to Eq.(6) are only active when the primary path passes through link $(i,j)$, i.e., $x_{ij} = 1$.

$$v_{ij}^{ij} + v_{ji}^{ij} = 0, \quad \forall (i,j) \in L, \tag{8}$$

Eq.(8) enforces the path disjoint property.

$$y_{lk} \geq d_{ij}^{lk} v_{ij}^{lk}, \quad \forall (i,j), (l,k) \in L, \tag{9}$$

Eq.(9) indicates whether a unoccupied wavelength on link $(l,k)$ will be reserved for protection. Notice that $v_{ij}^{lk} = 1$ and $d_{ij}^{lk} = 0$ together mean that sharing protection bandwidth is possible.

$$x_{ij} \geq v_{ij}^{lk}, \quad \forall (i,j), (l,k) \in L, \tag{10}$$

Eq.(10) prevents the possibility of assigning a protection path for a link that is not used by the primary path.

$$v_{ij}^{lk} + x_{mn} \leq v_{mn}^{lk} + 1, \forall (i,j), (l,k), (m,n) \in L, \tag{11}$$

$$x_{ij}, y_{ij}, v_{ij}^{lk} \in \{0,1\}, \quad \forall (i,j), (l,k) \in L. \tag{12}$$

Eq.(11) ensures that each link reserved for protection must also protect the whole primary path. For example, if a wavelength on link $(l,k)$ is reserved to protect a primary path which passes through link $(i,j)$, then we have $v_{ij}^{lk} = 1$. Since link $(l,k)$ must also protect other links on the primary path, say link $(m,n)$ ($x_{mn} = 1$), we need to set $v_{mn}^{lk} = 1$. If the primary path does not pass through link $(m,n)$, i.e., $x_{mn} = 0$, then by constraint Eq.(10), $v_{ij}^{lk} = 0$ in this case. Hence, we assure the property that each link on a protection path protects every link of the associated primary path.

We next introduce the ILP formulation for PPP. Recall that, in this protection scheme, the system reserves a protection path for each link along the primary path and thereby the system reserves resources for one or multiple protection paths to protect the associated primary path.

The objective function for the path protection scheme remains the same for PPP. The constraint set of the formulation is as follows.

$$\text{Minimize} \quad \sum_{(i,j)\in L} c_{ij}x_{ij} + \sum_{(i,j)\in L} y_{ij}$$

$$\text{Subject to} \quad \sum_{(S,j)\in L} x_{Sj} - \sum_{(j,S)\in L} x_{jS} = 1, \qquad (13)$$

$$\sum_{(D,j)\in L} x_{Dj} - \sum_{(j,D)\in L} x_{jD} = -1, \qquad (14)$$

$$\sum_{(i,j)\in L} x_{ij} - \sum_{(j,i)\in L} x_{ji} = 0, \forall i \neq S, D, \qquad (15)$$

$$\sum_{(S,l)\in L} v_{ij}^{Sl} - \sum_{(l,S)\in L} v_{ij}^{lS} \geq x_{ij}, \quad \forall (S,l),(l,S),(i,j) \in L, \qquad (16)$$

$$\sum_{(l,D)\in L} v_{ij}^{lD} - \sum_{(D,l)\in L} v_{ij}^{Dl} \geq x_{ij}, \quad \forall (D,l),(l,D),(i,j) \in L, \qquad (17)$$

$$\sum_{(l,k)\in L} v_{ij}^{lk} - \sum_{(k,l)\in L} v_{ij}^{kl} = 0, \quad \forall (i,j) \in L, \forall k \neq S, k \neq D, \qquad (18)$$

$$v_{ij}^{ij} + v_{ji}^{ij} = 0, \quad \forall (i,j) \in L, \qquad (19)$$

$$y_{lk} \geq d_{ij}^{lk}(v_{ij}^{lk} - x_{lk}), \forall (i,j),(l,k) \in L, \qquad (20)$$

$$x_{ij} \geq v_{ij}^{lk}, \quad \forall (i,j),(l,k) \in L, \qquad (21)$$

$$x_{ij}, y_{ij}, v_{ij}^{lk} \in \{0,1\}, \forall (i,j),(l,k) \in L. \qquad (22)$$

Note that the difference between the two formulations is that we transform Eq.(9) into Eq.(20), and we also remove Eq.(11) from the previous formulation. Eq.(20) considers the situation where a protection path overlaps part of its links with the links on its associated primary path. The overlap incurs no cost. We eliminate Eq.(11) from the formulation for the path protection scheme, because there is no need to force a link on a protection path to protect the entire primary path in PPP.

### B. SP approach

Note that our system, with call arrivals and departures, is a discrete time system. The optimal solution can be obtained through dynamic programming, which would be prohibitively complex. The dynamic program takes into account the impact of present decisions on future system performance. The greedy algorithm only considers present resource usage, and thereby does not necessarily achieve optimality. The greedy approach can result in an inferior network resource utilization because it may choose paths with little opportunity for protection sharing (see Example 2). Therefore, we consider another implementation approach which encourages protection sharing as follows.

First, note that a request's primary path cannot be shared with other requests. Thus, it is natural to attempt to dedicate the fewest possible resources to a call's primary path. Therefore, we assign the shortest path for a call request as its primary path. After the call's primary path is identified, we then seek the protection paths for it. To encourage protection sharing, we construct a new graph. In the new graph, the network topology remains intact but the link costs are updated according to the resource usage status. Wavelengths that are in use by other protection paths have a cost of 0. In the case of path protection, links used by the primary path are not available in the new graph. In the case of PPP, unfailed links in the primary path are available for no cost in the protection path.

### C. Greedy versus SP

We may briefly compare our two implementation approaches. As mentioned, solving an ILP is computationally intensive. In contrast, since the algorithms for seeking the shortest paths, e.g. the Dijkstra's algorithm, are polynomial-time, the shortest primary path approach can place a new call rapidly. For static batch models, computational complexity is not very important, since decisions are not made in real time. For a dynamic call-by-call system, however, ease and speed of computation are more relevant.

Let us now consider resource efficiency. While the SP approach may at times require more resources for a given call, it is possible that over a number of calls, the SP approach may eventually result in more efficient bandwidth utilization. Example 2 illustrates this phenomenon.

*Example 2:* Consider the network in Fig. 5 and assume that the network employs PPP. The network is initially empty and serves three call requests, $(1,4)$, $(6,3)$, and $(3,5)$, in sequence. Table III shows the resource assignments for the greedy approach and the SP approach. In this

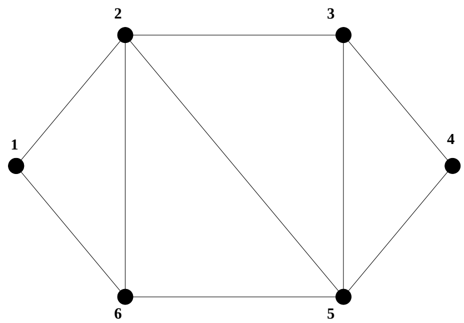


Fig. 5. An example network

| | SD Pair | Primary Path | Protection Path (protected link) | Total Number of Occupied $\lambda$'s |
|---|---|---|---|---|
| Greedy approach | (1,4) | 1-2-3-4 | 1-6-5-4 (1-2-3-4) | 6 (no sharing) |
| | (6,3) | 6-5-3 | 6-2-3 (6-5-3) | 10 (no sharing) |
| | (3,5) | 3-5 | 3-2-5 (3-5) | 13 (no sharing) |
| Shortest path approach | (1,4) | 1-2-3-4 | 1-6-2-3-4 (1-2) | 7 (share (2-3-4)) |
| | | | 1-2-5-4 (2-3) | (share (1,2)) |
| | | | 1-2-5-4 (3-4) | |
| | (6,3) | 6-5-3 | 6-2-3 (6-5) | 10 (share (6,2)) |
| | | | 6-2-3 (5-3) | |
| | (3,5) | 3-5 | 3-2-5 (3-5) | 12 (share (2,5)) |

TABLE III
RESOURCE USAGE FOR NETWORK EMPLOYING PARTIAL PATH PROTECTION SCHEME IMPLEMENTED BY DIFFERENT APPROACHES IN FIG. 5

example, the SP approach initially occupies more wavelengths to support the request $(1, 4)$ than does the greedy approach. However, as the calls accumulate, the SP approach uses fewer number of wavelengths to support the same requests than the greedy approach.

In this example, the greedy approach endeavors to serve each request using the minimum number of previously unused wavelengths. However, in doing so, the greedy approach happens to choose paths with no protection sharing, harming network resource utilization. In contrast, though the SP is not optimal at first, it performs better over the call arrivals, by encouraging protection sharing.

## IV. SIMULATION RESULTS

To investigate the protection schemes, we simulate path protection and PPP schemes implemented using both the greedy approach and the SP approach. We assume that the networks and the call requests have the following characteristics. First, all nodes in the network are equipped with wavelength converters. We therefore focus on the problem of whether an available wavelength exists on a link. Essentially, the network is regarded as a circuit-switched network. Second, in the simulation, the cost for placing a call refers to the aggregate link costs, as defined in Sec.III-A. Third, we assume full knowledge of the network resource status in our search for primary and protection paths. Fourth, the acceptance of a call request is completed only after the system reserves the available network resources for both primary and protection paths. Otherwise, we regard the incoming request as being blocked. Fifth, we assume that the arrival of call requests forms Poisson process and that calls have an exponentially distributed service time. The traffic load refers to the product of the arrival rate and the average service time. Finally, we assume uniform traffic, in which an arrival will choose one out of all possible source and destination pairs with equal probability.

In our simulations, we consider two nation-wide US networks, NSFNET (shown in Fig. 6) and Sprint's OC-48 network (shown in Fig. 7), and a regional network, the New Jersey LATA network (NJ LATA, shown in Fig. 8). Additionally, each link in the networks contains 16 bidirectional wavelengths. Note that the nodes in both national networks usually have a lower degree than those in the NJ LATA network, i.e., the regional network is denser.

Two measurements are investigated in the simulations to evaluate the performances of the protection schemes. The first measurement is the steady state blocking probability. Blocking probability is related to opportunity cost, referring to the additional revenue available if certain customers were not turned away. The second measurement is the aggregate number of occupied wavelengths on each link to support connections in the network. This measurement reflects the network resource utilization. For simplicity, we denote PPP implemented by the greedy and the SP approaches as Greedy-PPP and SP-PPP, respectively. We denote path protection using the greedy and the SP approaches as Greedy-PP and SP-PP, respectively.

Fig. 9 to Fig. 14 present our simulation results and Table IV summarizes the results. The results show that, with the same implementation approach, PPP is better than path protection. Still, for each of the protection schemes, the SP approach is better than the greedy approach as the calls

| | Path protection | PPP |
|---|---|---|
| Greedy approach | Inferior | Inferior |
| SP approach | Worst | Best |

TABLE IV
SUMMARY OF SIMULATION RESULTS

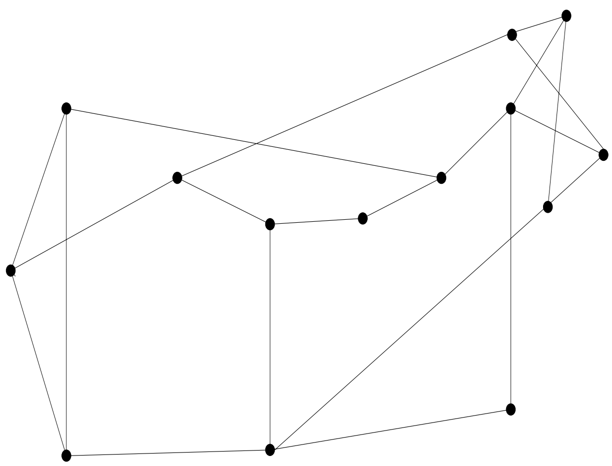

Fig. 6. NSFNET

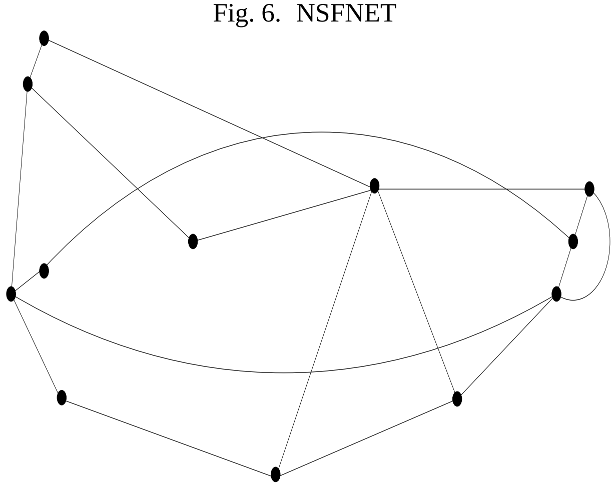

Fig. 7. Sprint OC-48 Network

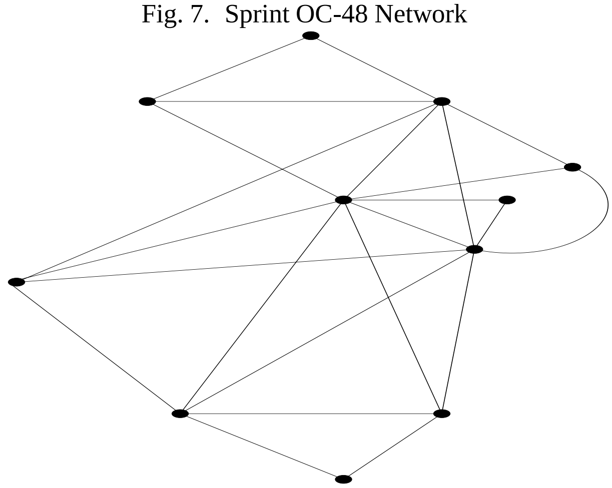

Fig. 8. New Jersey LATA Network

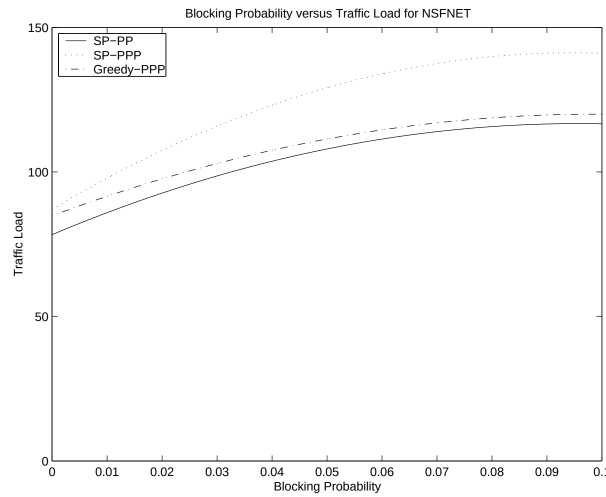


Fig. 9. Traffic Load vs. Blocking Probability in NSFNET

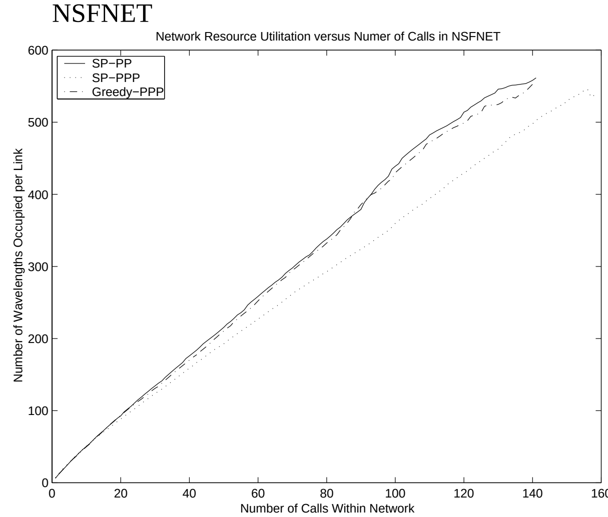


Fig. 10. Network Resource Utilization in NSFNET

accumulate. Our two major conclusions from our simulations are that, as shown in Table IV, the PPP scheme implemented using the SP approach has the best performance, and that the other combinations, Greedy-PP, Greedy-PPP and SP-PP, perform worse and are comparable among each other. We discuss these conclusions below.

Since SP-PPP is intrinsically more flexible than SP-PP in both the protection scheme and the implementation approach themselves, SP-PPP has lower blocking probabilities than other combinations in all networks simulated, shown in Fig. 9, 11, and 13. Note that, owing to the scale of the ILP for path protection, we provide the Greedy-PP result only for NJ LATA. In the figure, Greedy-PP performs better than SP-PP when the aggregate number of calls is smaller than 125 or so. But, as the aggregate number becomes larger, SP-PP becomes more efficient than Greedy-PP, a phenomenon which is very similar to that seen for PPP.

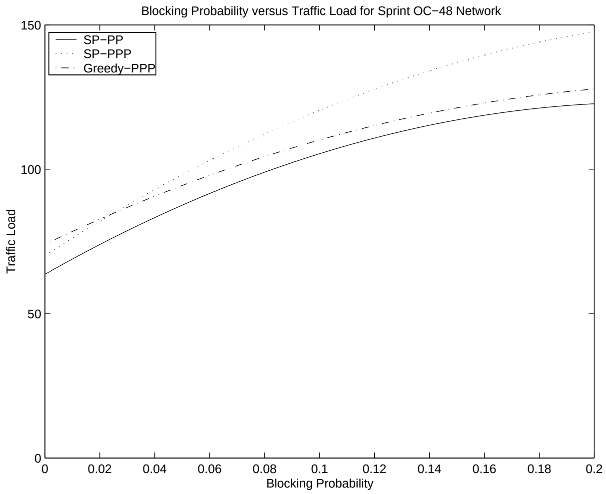


Fig. 11. Traffic Load vs. Blocking Probability in Sprint OC-48 Network

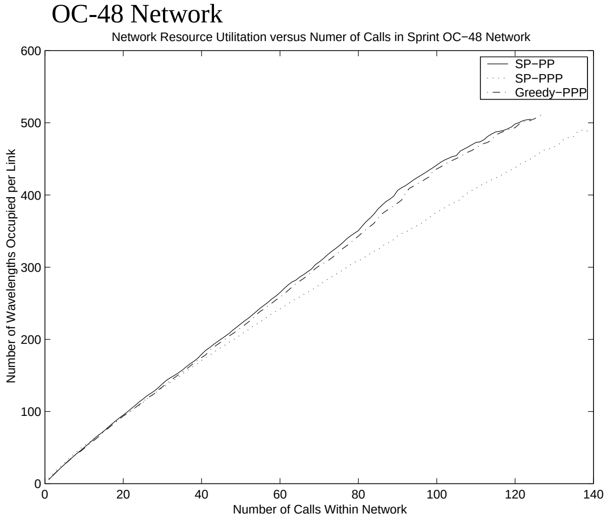


Fig. 12. Network Resource Utilization in Sprint OC-48 Network

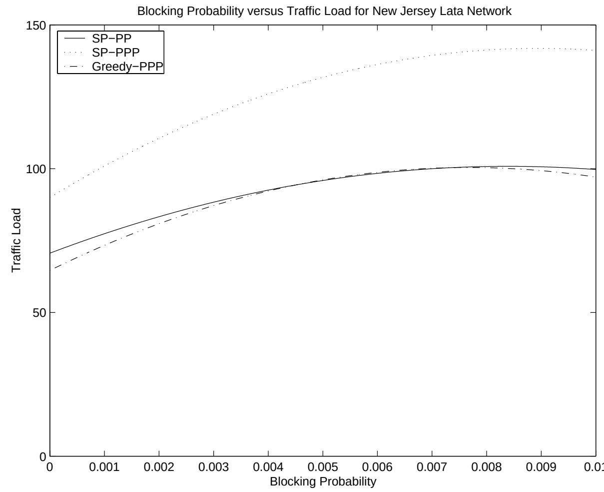


Fig. 13. Traffic Load vs. Blocking Probability in New Jersey Lata Network

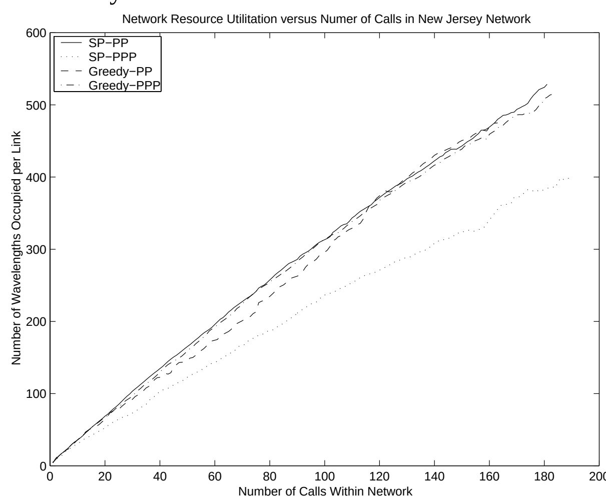


Fig. 14. Network Resource Utilization in New Jersey Lata Network

Example 3 illustrates why SP-PP, Greedy-PP and Greedy-PPP perform almost the same. Owing to the nature of the greedy algorithm, the Greedy-PPP approach attempts to occupy the minimum number of wavelengths to serve a call. To this end, Greedy-PPP will find the smallest possible number of wavelengths to protect the corresponding primary path. As a result, one single protection path for a primary path occurs in most cases in the simulation, even though the partial path protection scheme does not require all the protection paths to be the same. Hence the Greedy-PPP has an extremely similar performance to SP-PP and Greedy-PP, which are restricted to assign one single protection path per primary path.

*Example 3:* Consider Fig. 5 and a source destination pair $(1, 4)$. We have the resource allocation shown in Table V for SP-PP, SP-PPP, and Greedy-PPP. The table shows that the primary and protection paths for SP-PP are identical to those for Greedy-PPP. This is because Greedy-PPP attempts to fulfill the protection requirement with the minimum number of wavelengths. Note that SP-PPP has the worst performance in terms of network resource utilization in this case. This fact agrees with our simulation results showing that SP-PPP does not perform very well when the network is very lightly loaded. However, as calls accumulate, protection sharing becomes more important for resource utilization and thus SP-PPP is more efficient.

Another important observation is that the performance of the protection schemes are highly related to the network topology. For the highly connected regional network, i.e. the NJ LATA network, the blocking events are relatively rare. As shown in Fig. 13, when the blocking probability is set as 0.01, the achievable traffic load for NJ Lata network is far above 100. Conversely, from Fig. 9 and 11, the achievable traffic loads for the two relatively sparse

| | Primary Path | Protection Path (protected link) | Number of Occupied $\lambda$'s |
|---|---|---|---|
| SP-PP | 1-2-3-4 | 1-6-5-4 | 6 |
| SP-PPP | 1-2-3-4 | 1-6-2-3-4 (1-2)<br>1-2-5-4 (2-3)<br>1-2-5-4 (3-4) | 7 |
| Greedy-PPP | 1-2-3-4 | 1-6-5-4 (1-2)<br>1-6-5-4 (2-3)<br>1-6-5-4 (3-4) | 6 |

TABLE V

RESOURCE ALLOCATION FOR SOURCE DESTINATION PAIR $(1, 4)$ OF NETWORK IN FIG. 5

national-wide networks are both of 100 or so, as the blocking probability is fixed at 0.01 as well. The main reason for this phenomenon is that there exist many more choices in the regional network to serve a call request than in the nation-wide networks. Hence, a blocking event is relatively rare in the highly connected network.

Comparing together for the two nation-wide networks, we first note that NSFNET has a lower degree than Sprint OC-48 network and has a better performance than Sprint OC-48 network. From Fig. 9 and Fig. 11, we observe that NSFNET achieves higher traffic load than does the Sprint network for a fixed blocking probability. This indicates that, as the network becomes sparser, the benefits of SP-PPP over SP-PP and Greedy-PPP may increase. This argument also applies to the comparison of Fig. 10 and 12. One conjecture to buttress this observation is that, since NSFNET is relatively sparse, there are more occurrences of long primary paths, which enhance the usefulness of capacity sharing and thus make the improvements due to SP-PPP noticeable.

## V. CONCLUSIONS

We have introduced a novel protection scheme, PPP. Moreover, instead of considering traditional static capacity-efficiency measures for evaluating the efficiency of protection schemes, we considered a dynamic call-by-call model. To avoid the complexity of dynamic optimization, we presented two heuristics to implementing path protection and PPP. These approaches, which we termed greedy and SP, were compared to each other for both path protection and PPP. We have demonstrated that PPP is superior to path protection and that SP is superior to the greedy approach. As expected from the fact that PPP is more general and flexible than path protection, PPP outperforms path protection in terms of blocking probability. Moreover, the SP approach performs better than the greedy approach. It is the dynamic nature of our problem that renders SP superior to the greedy approach. Indeed, SP emphasizes reducing resource use among primary paths, since their bandwidth cannot be shared. The fact that SP may be less efficient than the greedy approach in its allocation of capacity for protection paths is mitigated by the fact that protection bandwidth can be shared.

The advantages of PPP over path protection has certain implications in the area of network management. Path protection only requires that the source and destination node be aware that a failure occurred somewhere along the primary path. Localization of the failure is unimportant, since protection takes place in the same way regardless of where the failure occurs. Thus, once the protection path has been set up, the network management does not need to have detailed knowledge of the nature of the failure to effect protection. Path protection can then be handled by higher layer mechanisms. For link protection, local information is needed by the nodes adjacent to the failure, but there is no need to manage protection on a path-by-path basis. Lower layers can therefore ensure link protection. PPP, on the other hand, requires on the part of the network management effecting protection both knowledge of the path and of the location of the failed link. Our results point to the fact that visibility by the network management system across layers may be useful for performing protection efficiently.

There are several further research directions for our work. One such direction is to consider the case of batch arrivals rather than dynamic call-by-call arrivals. We expect that the preferable approach in the static batch case is to solve some ILP similar to the one set up for our greedy approach. Comparing the results of the batch case with those of our dynamic system should yield insight into the effect of the dynamic assumption upon the effectiveness of protection schemes. Another area of further research is the generalization of our PPP algorithm to the case where failures are localized to segments, possible comprising several links. Such a generalization would allow us to study the effect upon blocking probability of different granularities of failure localization. Path protection can be viewed as the case where the whole segment is composed of a single segment and PPP as the case where every segment comprises a single link.

## REFERENCES


[1] T. Wu, *Fiber Network Service Survivability,* Norwood, MA: Artech House, 1992.

[2] O. Gerstel, *Opportunities for Optical Protection and Restorations,* Proc., OFC '98, San Jose, CA, vol. 2, pp. 269-270, February 1998.

[3] P. Bonenfant, *Optical Layer Survivability: A Comprehensive Approach*, Proc., OFC '98, San Jose, CA, vol. 2, pp. 270-271, February 1998.

[4] Rajiv Ramaswami and Kumar N. Sivarajan, *Routing and Wavelength Assignment in All-Optical Networks,* IEEE/ACM Transactions on Networking, vol. 3, no. 5, pp. 489-500, October 1995.

[5] Rajiv Ramaswami and Adrian Segall, *Distributed Network Control for Optical Networks,* IEEE/ACM Transactions on Networking, vol. 5, no. 6, pp. 936-943, Dec. 1997.

[6] Rajiv Ramaswami and Kumar N. Sivarajan, *Optical Networks: A Practical Perspective,* Morgan Kaufman, 1998.

[7] G. Ellinas and T.E. Stern, *Automatic Protection Switching for Link Failures in Optical Networks with Bi-Directional Links,* Proceedings of Globalcom, pp. 152-156, 1996

[8] Dimitris Bertsimas and John N. Tsitsiklis, *Introduction to Linear Optimization,* Athena, 1997.

[9] R.R. Iraschko, M.H. MacGregor, and W.D. Grover, *Optimal Capacity Placement for Path Restoration in Mesh Survivable Networks,* Proceedings of ICC, pp. 1568-1574, 1996.

[10] Murali Kodialam and T.V. Lakshman, *Dynamic Routing of Bandwidth Guaranteed Tunnels with Restoration,* Infocom 2000, April 2000.

[11] M. Kuznetsov, N. M. Froberg, Eytan Modiano, and et. al., *A Next-Generation Optical Regional Access Network*, IEEE Communications Magazine, vol. 38, Issue 1, pp. 66-72, Jan. 2000.

[12] S. Ramamurthy and B. Mukherjee, *Survivable WDM mesh networks, Part I - Protection,* INFOCOM '99, Eighteenth Annua Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE, vol. 2, pp. 744-751, 1999.

[13] S. Ramamurthy and B. Mukherjee, *Survivable WDM mesh networks. II. Restoration*, Communications, 1999. ICC '99. 1999 IEEE International Conference on, vol. 3, pp. 2023-2030, 1999.

[14] Eytan Modiano and Aradhana Narula, *Survivable Routing of Logical Topologies in WDM Networks,* Infocom 2001, Anchorage, Proceedings, IEEE, vol. 1, pp. 348-357, 2001.

[15] Bala Rajagopalan, Dimitrios Pendarakis, Debanjan Saha, Ramu S. Ramamoorthy, and Krishna Bala, *IP over Optical Networks: Architectural Aspects,* IEEE Communications Magazine, pp. 94-102, Sep. 2000.

[16] Muriel Médard, Steven G. Finn, Richard A. Barry, and Robert G. Gallager, *Redundant Trees for Preplanned Recovery in Arbitrary Vertex-Redundant Graphs,* IEEE/ACM Transactions on Networking, vol. 7, no. 5, PP. 641-652, Oct. 1999.

[17] Steven S. Lumetta, Muriel Médard, and Yung-Ching Tseng, *Capacity versus Robustness: A Tradeoff for Link Restoration in Mesh Networks,* Journal of Lightwave Technology, vol. 18, Issue 12, pp. 1765-1775, Mar. 2000.