# Sub-path Protection: A New Framework for Optical Layer Survivability and its Quantitative Evaluation

**Vishal Anand *, Sunit Chauhan and Chunming Qiao**

Department of Computer Science and Engineering

State University of New York at Buffalo

Buffalo, New York 14260

Email:{vanand, chauhan, qiao}@cse.buffalo.edu

**Abstract**

*In wavelength division multiplexed networks (WDM) with mesh topolgies, past research has considered two typical survivability paradigms: (a) path protection/restoration and (b) link protection/restoration. In path based schemes a backup path (which is link-disjoint with the primary path) and wavelength is reserved during primary path setup, while in link based schemes, for every link along the primary path a link disjoint backup for that link is found and reserved. Each of these schemes have their advantages and disadvantages in terms of capacity utilization and restoration times.*

*In our study here, we introduce a new protection/restoration scheme called sub-path protection. We show that this scheme has advantages over the traditional link and path based schemes. Experimental results show that sub-path protection has interesting trade-offs in terms of capacity utilization and recovery times.*

* Corresponding Author

**Keywords:**

WDM networks, lightpath, wavelength routing, wavelength assignment, protection and restoration, path protection, sub-path protection, link protection, time to recover.

## 1 Introduction

*Optical wavelength division multiplexed (WDM)* [1, 2] networks are considered to be the future wide-area backbone networks. Wavelength division multiplexing divides the tremendous bandwidth of a fiber into many non-overlapping wavelengths (or wavelength channels). Bit rates of 2.5 to 10 Gigabit per second are now possible on each wavelength of a fiber, which can have 100 or more wavelengths. The fact that the adoption of the WDM technology has resulted in an ever increasing degree of concentration of traffic on a single fiber-link, whose failure in a WDM network may affect upwards of a Terabit of traffic per second, motivates the study of WDM network survivability and reliability.

Although higher protocol layers (such as IP) operating over a WDM network may have their own recovery procedures to deal with link failures, we consider a protection-based fault tolerant approach at the WDM layer in this paper for the following two main reasons [3, 4]; a) Some of the higher layers are not able to provide the required protection functionality, and optical layer protection can be used to provide additional resilience for all the higher layers. b) Given that a wavelength may carry a large number of higher layer connections, optical layer protection can provide high speed recovery and is more scalable. Note that since recovery at higher layers has its own advantages, multi-layer recovery schemes should also be pursued [3, 5, 6].

To survive *single-link* failures in an optical network, two basic approaches based on *path* and *link* protection and restoration have been used and studied in the past. Most of the earlier work also restricts itself to the case where it is assumed that at any time only a *single* fiber link can fail. In reality two or more links may fail simultaneously. Most of the algorithms and techniques developed for the single-link failures case would not work when multiple-link failures occur. The strategy we develop in this paper called *sub-path* protection, provides protection against any *single* link failure at the optical layer as in *1:1 protection* [7, 8], where, at the time of establishing a *working* (primary path), a link-disjoint *protection* (backup) path is also established (we assume that in the absence of any link failures, low-priority pre-emptable and unprotected traffic can be carried on a backup path which will not share bandwidth with other primary or backup paths.). In addition our strategy can survive certain multiple link failures (more such multiple failure combinations than path protection), as will be explained later. We compare our scheme with the path and link based schemes in terms of *capacity* (bandwidth) utilization and *recovery* times, which are the two most important criteria when evaluating a protection scheme. Extensive simulations have been performed and our results show that using sub-path protection has various advantages over the traditional *path* and *link* protection schemes.

## 2 Terminology and Related Work

We consider a wavelength-routed WDM network whose *physical topology* consists of (optical) wavelength-routers connected by point-to-point fiber links in an arbitrary mesh topology as shown in Figure 1. Between two connected wavelength-routers, there is a pair of unidirectional fibers (or equivalently, a bi-directional link). Each wavelength-routing node takes in a signal on a wavelength at one of its inputs, and routes it to a wavelength at a particular output, without undergoing opto-electronic conversion. An access station (e.g., an IP router) may be connected to each optical wavelength-router, which can transmit/receive signals through either a tunable transmitter/receiver or a transmitter/receiver array. A connection request is satisfied by establishing a *lightpath* from the source node (access station or wavelength-router) of the connection to the destination node. A lightpath uses one wavelength on each link it spans to provide a circuit-switched interconnection between the source and destination nodes. As shown in Figure 1, two lightpaths, one between nodes 3 and 5, and the other between nodes 3 and 7, must use different wavelengths, ($\lambda_1$ and $\lambda_2$) on a common fiber link ($3 \rightarrow 4$), in order to prevent interference of the optical signals.

In this paper, we assume that none of the wavelength-routers have wavelength conversion capabilities (since all-optical wavelength-conversion is an immature/expensive technology). Therefore, a lightpath has to occupy the *same* wavelength on all the links that it spans (this is called the *wavelength continuity constraint*). In such a wavelength-routed network, the routing and wavelength assignment (RWA) problem consists of two components. The first is a routing scheme which determines the route each of the lightpaths will take. The second component of the RWA problem is to assign a wavelength on each link along the chosen route.

Protection/restoration schemes to ensure fiber newtork survivability can be broadly classified based on whether the resources are *dedicated* for protection or *dynamically* allocated for restoration [7, 9, 10]. In dedicated resource protection, network service is restored utilizing the dedicated network resources. The network resources used in dedicated resource protection may be dedicated for each failure scenarios or alternatively, the dedicated network resources used for protection may be *shared* among different failure scenarios. In a dynamic scheme, the spare capacity available within the network *at the time* of failure is utilized for restoring services affected by the failure. Generally, dynamic restoration schemes are more efficient in utilizing capacity due to the multiplexing of the spare capacity and efficient utilization of resources. On the other hand dedicated schemes have a faster restoration time and provide guarantees on the restoration ability. We describe only the protection schemes here, which is the main focus of this paper.

To survive single-link failures, protection schemes can broadly be classified into: path and link protection.

- *Path protection*: In path protection, the source and destination nodes of each connection statically reserve backup paths and wavelengths during call setup. The primary and backup path are *fully link-disjoint*.

    1. *Dedicated-path protection*: In dedicated-path protection, at the time of call setup for each primary path, a link-disjoint[1] backup path and wavelength are reserved and dedicated to that call. The backup wavelength reserved on the links of the backup path are dedicated to that call, and are not shared with other backup paths.

    2. *Shared-path protection*: In shared-path protection at the time of call setup for a primary path, a link-disjoint backup path and wavelength are also reserved. However, the backup path may be shared with other backup paths. As a result, backup channels are multiplexed among different failure scenarios, which are not expected to happen simultaneously. Therefore shared-path protection is more capacity efficient when compared with dedicated-path protection.

- *Link protection*: In link protection, all the connections that traverse the failed link are re-routed around that link. The source and destination of the connections traversing that link are oblivious to the link failure. In link protection, during call setup, backup paths and wavelengths are reserved around each link.

[1] There is no common fiber-link between the backup and primary path for that connection.

1. *Dedicated-link protection*: In dedicated-link protection, at the time of call setup, for each link of the primary path, a backup path and wavelength are reserved around that link and are dedicated to that call.

2. *Shared-link protection*: In shared-link protection, at the time of call setup, for each link of the primary path, a backup path and wavelength are reserved around that link. However the backup path and wavelengths reserved on the links of the backup path may be shared with other backup paths. As a result, backup channels are multiplexed among different failure scenarios, which are not expected to happen simultaneously. Therefore shared-link protection is more capacity efficient when compared with dedicated-link protection.

In our work we consider only dedicated (or 1:1) type of protection (i.e. we assume no sharing occurs). During normal (failure-free) operation of the network, pre-emptable traffic is carried on the wavelengths reserved for the backup paths. This is to take advantage of the idle wavelength resources on the backup paths. When a link-failure occurs, the backup path and wavelength is chosen to route the traffic. The pre-emptable traffic which was using these wavelengths along the backup route is thrown out of the network.

Each of the above schemes (path and link) have been shown to posses certain desirable characteristics and outperform the other under different situations and circumstances. In [7, 8], the authors showed that path protection schemes provide efficient capacity utilization over link protection schemes, on the other hand link protection has a better service restoration time compared to path protection. In our work, we explore the benefits and shortcomings of each of these schemes and then design a scheme which is balanced in the sense that it incorporates the good features of each of these schemes. We call our scheme sub-path protection. The objective of such a scheme is to use as little resources as path protection and to provide fast restoration as in link protection.

## 3 Protection Architecture

Before we describe the sub-path protection scheme and its working, we describe the associated node model and protection architecture associated for such an optical network.

The optical layer consists of nodes interconnected by directed links which may consist of multiple fibers. However, we assume single fiber directed links, where each fiber can carry multiple wavelengths.

Figure 2 illustrates our protection model, where each of the wavelength channel is bi-directional. We assume that on each fiber the wavelengths can be distinguished on the basis of whether they carry data or signaling (or control, i.e. operation, administration and maintenance (OAM)) traffic. The primary and backup paths carrying data are directed, while the OAM/signaling channels are bi-directional. Let us assume that a primary and a backup path between source 1 and destination 5 is set up as shown. Now assume that link $4 \rightarrow 5$ breaks. Once the end nodes (4 and/or 5) of that link come to know of the failure (e.g. via loss-of-signal (LOS)), node 4 propagates an alarm signal via the signaling channel to node 1, the source of the connection. Once the source is informed of the failure, a signal is sent to all the nodes on the (pre-determined) backup path. This signal is used to configure the nodes along the backup path. The destination node is also informed to start receiving data along the backup path. It is important to note that with such a scheme, the *time to recover* (TTR) from a failure depends mainly on the lengths of the paths the signaling information has to travel (which often depends on the lengths of the primary and backup paths). Thus, one can immediately see why link protection schemes can be faster than path based protection schemes.

## 4 Sub-path Protection

As explained earlier, in our work we explore the benefits of the sub-path protection scheme to achieve a balance between short service restoration times and high capacity utilization achieved by the link and path protection schemes respectively.

Figure 3 illustrates the link, path and and sub-path protection schemes in detail. Let us assume a network as shown, and the source and destination be as indicated and a primary path spanning four links from the source to destination. In path protection (see Figure 3 $(a)$), we may find the backup path after removing the links along the primary path and recompute the second shortest path, which is now guaranteed to be link-disjoint with the primary path.

In link protection as depicted in $(b)$ of the Figure, we may remove one link along the primary path at a time. starting with the first link. We find a shortest path between the end nodes of the link, which is guaranteed to be link disjoint with the removed link. Note also that while finding such backup paths, one for each link, all the links along the primary path may be removed from the network, first to guarantee that *each* backup path is link-disjoint with *all* other backup paths and the primary path. It is also possible to find backup paths such that they are link disjoint with each other. We will explain the significance of this later.

Figure 3($c$) illustrates the concept of sub-path protection. A sub-path is defined as a subset of the links along a primary path. Let us assume that we opt for a sub-path of size 2 links (this size can be varied by the network administrator etc. depending on what needs to be to optimized – bandwidth, restoration time etc.). So now the primary path has been divided into two sub-paths $P_1$ and $P_2$, of size 2 links each. Note that though the primary path is already determined in this example, one may also find the primary paths by first running Dijkstra's algorithm to find the shortest path from source node 1 to node 3 and then repeat the process for the second sub-path from node 3 to node 5. In either case, to find the backup paths for these sub-paths of the primary path, we remove the links along each of the primary sub-paths and then run the Dijkstra's algorithm to find the link disjoint shortest backup paths. We can easily see that link and path protection are special cases of sub-path protection when the size of the sub-path is equal to $n$, ($n$=length of primary path and $n = 1$ respectively).

Another point to note here is the survivability of each of these schemes in the case of *multiple-link* failures. It is obvious that path protection can survive as many link failures as there are links in the primary path. In link protection the primary path and backup paths for each link typically overlap (or share wavelengths). So link protection can survive single-link failures only. Note, as long as the backup paths of each of these links is disjoint, $n$ (n=number of links in the primary path) simultaneous failures can be tolerated. A similar comment can be made for the sub-path protection. So each of our scheme can survive $n$ multiple-link failures along the primary path. What will differ from scheme to scheme is the time to recover. Let us assume link $4 \rightarrow 5$ (Figure 3) breaks, then in path protection scheme the signal has to travel all the way to the source node 1 and then the data has to be sent on the backup path. In link protection, the alarm signal has to travel the length of one link *only*, to node 4 from where the data is switched onto the backup path for this link. The rest of the primary path up to node 4 is still used to send data and the switch to the backup happens only at node 4. In sub-path protection, the signal has to travel up to node 3 (a shorter fiber-distance compared to path protection but longer than link protection), before the switch to the backup occurs. In terms of device cost and complexity, note that in both path and link protection scheme *each* node along the primary path should have the intelligence to able to detect a link failure, and in path protection *only* the source node needs to have the ability of switching to the backup path, while in link protection *each* node needs to be able to do such a switching. In sub-path protection, on the other hand,*only* nodes 1 and 3 need to have this switching capability.

Depending on how we assign wavelengths to the primary and backup paths (sub-paths), we can classify the sub-path protection schemes into 3 types as shown in Figure 4.

In *Case 1*, the primary path and the protection sub-paths both occupy the same wavelength ($\lambda_1$). The advantage of this is that node 3 need not have any *wavelength conversion* capabilities. Say there is a failure along $P_2$, any failure along $P_2$ will be handled by the respective backup path found. For the no-failure path, $P_1$, the original primary path can still be used and a switch to the backup path can be made only at node 3. Since the protection and primary path use the same wavelength no conversion is necessary.

By restricting the primary and backup paths to the same wavelength, we are restricting our choices of backup paths. In *Case 2*, the primary and backup paths can be assigned different wavelengths ($\lambda_1, \lambda_2$), the backup paths of each of the sub-paths use the same wavelength, again a similar analysis can be done and it can be shown that there is a need for wavelength conversion capability at node 3. The trade-off being the additional flexibility and choices of finding backup paths and hence a potential of savings in bandwidth. In *Case 3*, the primary and backup paths can be assigned different wavelength ($\lambda_1, \lambda_2, \lambda_3$), in addition the backup paths of each of the sub-paths can be assigned different wavelengths ($\lambda_2$ and $\lambda_3$). This is a more complicated case in terms of signaling and switching, it is obvious that the nodes need wavelength conversion capabilities, and there may be a need for more than one wavelength conversion at each of the intermediate nodes.

## 5 Path, Link and Sub-path protection - Capacity utilization and Time to recover

As discussed earlier, the efficiency in capacity utilization (throughput achieved) is high, but the time to recover is long in the path protection scheme. Link protection on the other hand, is just the opposite with a low capacity utilization but short recovery time. In sub-path protection we try and strike a balance in capacity utilization and recovery time. This is explained in Figure 5. The balance point (or the ideal sub-path length) can be chosen by the network operator or administrator, depending on what his optimization criteria, this could be to maximize the connections carried, minimize recovery time or to strike a balance between the two. Furthermore, providing different recovery times can be a good way in which network providers can provide *service-differentiation* and *quality of service guarantees* (QoS) to the carried traffic. Sub-path protection could be used to provide this varying recovery times.

## 6 Numerical Results

### 6.1 Simulation Model

Before we present the numerical results, we describe the details of the simulation model used to get these results. The results presented here are for the network topologies shown in Figure 6 and Figure 9. To compare the different protection schemes, we define 3 parameters, $M$, the mean of the number of demands generated per node pair, $W$, number of wavelengths per link in the network, $S$, the percentage of the total number of demands satisfied by each of the protection schemes. In our simulations we vary $M$ from 0.3 to 1.0, and $W$ from 1 to 80. For each value of $M$ and $W$, 1000 random traffic matrices are generated such that the mean is $M$, the percentage of demands satisfied in each case is calculated and averaged over the 1000 runs, and $S$ is calculated. The value $M$ takes needs some more explanation. It is the average of all the demands in the traffic matrix. For example, when $M = 0.3$, it means that on an average there are 0.3 demands per node pair. This explains the results in the Tables 4 to 6, when $M$ is small (say 0.3) more demands (percentage wise) are satisfied (see column 2 of tables), than when $M$ takes a larger value (say 1).

The following parameters are shown in Tables 1 to 6, one in each column. To obtain results for different sub-path sizes, we try to divide the primary path in (up to) $1, 2, ...5$ sub-paths. The first column in each table thus indicates the average number of sub-paths per primary path (i.e. the number we *actually* get during our simulations). The number in the parenthesis indicates the maximum number of sub-paths we *try* to divide the primary path into. For example, if a primary path has only 3 hops and we intend to divide it into a maximum of 5 sub-paths, we will divide it into only 3 sub-paths (of size 1 hop each). Note that, when the average number of sub-paths per primary path is 1 (undivided), it is path protection. When this number is greater or equal to the average length of the primary path, link protection is usually implied.

Column 2 of each table (Number of demands satisfied (%)), is the percentage of the total demands that we were able to satisfy in the different instances (no. of sub-path). Column 3 (Avg. length of primary path in # of links), is self-explanatory. Column 4 (Avg. length of backup paths in # of links), is the average length of a backup path. For example, suppose we have a primary path of length 2 (links), and we divide this into 2 sub-paths each of length 1 (link). Now assume that the protection paths of each sub-path is 7 links, then the length of the backup paths would be 14. We then average this over all the sub-paths of all the demands and present it in Column 4. Column 5 is the average length of a backup path of each of the sub-paths. Note also that, value in column 4 = value in column 1 * value in column 5. Futhermore, column 4 gives the actual number of links (bandwidth) used for protection. This value increases with the number of sub-paths. The values in column 5, on the other hand is directly proportional to the restoration time, which should decrease (as in the case of the 54 node network) with an increase in the number of sub-paths.

### 6.2 Results for a 14 node network

The results for the network shown in Figure 6 are shown in Tables 1 to 3 and Figures 7 and 8. The primary paths and the backup paths for each of the schemes are calculated using Dijkstra's algorithm as explained in Section 4. The results in the table shows that as the number of sub-paths is increased from 1 to 5 (or 6) the percentage of demands satisfied falls. Path protection satisfying the maximum number of demands and link protection the least, with sub-path falling in between (this is in accordance with Figure 5a). We assume that the time to recover depends on the length (average number of links) of the backup paths. We notice from Tables 1 to 3 and Figure 8 that for a network topology shown in Figure 6, the time to recover (average length of backup paths of sub-paths) does not vary much and tends to remain more or less a constant (which is not what one would expect from Figure 5b & 5c).

### 6.3 Results for a 54 node network

Tables 4 to 6 show similar results for the network topology considered in Figure 9. Notice that on the average it is now possible to divide the primary path into more sub-paths. This is due to the fact that we have a larger network and primary paths tend to be longer in such a network compared to a smaller network. The capacity utilization of the various protection schemes is in accordance with Figure 5a, but the time to recover (as shown in

Figure 10) is not constant as in the case of the 14 node network. Further notice that variation (in terms of % of demands satisfied) does not vary much in this case. The reason for this is as follows, link/sub-path protection can make use of extra unused bandwidth in the network which path protection cannot, because path protection always requires a *totally* link disjoint path, which is more difficult to find. In other words, there could be scenarios where link (or sub-path) protection works but path protection does not. Furthermore, the chances that link (sub-path) protection can find (make use of) unused bandwidth on certain links, due to fragmentation of the newtork is higher than path protection for the same reasons as discussed above. This is more pronounced in larger networks as in smaller networks the large amount of bandwidth consumed by link and sub-path protection dominates, and thus results in a lower throughput.

For similar reasons, there is a larger variation in the length of the backup paths of the sub-paths and hence a larger variation in the time to recover in the 54 node network as shown in Figure 10 than in the 14 node network. This is in accordance with Figure 5b & Figure 5c, i.e, with an increase in the number of sub-paths (and hence a decrease in their lengths), there is a decrease in lengths of their respective backup paths, and accordingly, a decrease in the time to recovery time. Hence a natural trade-off between time to recover and efficiency of bandwidth usage (throughput) occurs. Note that the proportion with which the bandwidth consumption increases is much more than the proportion with which the restoration times (length of sub-paths) decrease. This implies that fast recovery using sub-path protection is possible, but is not directly related to protection bandwidth allocated. We note that the average length of the sub-paths and hence the bandwidth utilization efficiency and speed of recovery depend very much on the network topology (e.g. average connectivity, nodal degree etc.) and also the protection mechanisms used, as these factors affect the length and number of backup paths.

## 7 Conclusion and Future Work

Optical networks using WDM technology enable the transfer of several Gbps of data on each fiber link in the network. However, this large link capacity has the drawback that a single link failure can lead to the loss of huge amounts of data. Hence, network survivability in the event of link failures is a major concern in optical WDM networks.

Traditionally path and link based schemes have been used to survive link failures in optical networks. In this paper, we have provided a new framework for survivability in the event of link failures, called sub-path protection. Time to recover in the event of a failure and capacity utilization are two parameters that can be used to evaluate protection schemes. We have compared our scheme with the path and link based schemes using these parameters. We have shown through extensive simulations that our scheme can be used to combine the advantages of path and link schemes. A network provider or designer can make the trade-offs that our scheme provides to optimize his network design.

In our future work, we intend to study the effect of having *limited*, *full* and *no* wavelength-conversion capability. The affect of allowing sharing of bandwidth among the backup paths of different sub-paths will also be studied. Heuristic algorithms for efficient computation and wavelength assignment of these sub-paths and their respective backups will also be studied.

## References

[1] B. Mukherjee, *Optical Communication Networks.* New York: McGraw-Hill, 1997.

[2] R. Ramaswami and K. Sivarajan, *Optical Networks: A practical perspective.* San Francisco: Morgan Kaufmann, 1998.

[3] O. Gerstel and R. Ramaswami, "Optical Layer Survivability: A Services Perspective," *IEEE Comm. Mag.*, vol. 38, pp. 104–113, March 2000.

[4] O. Gerstel, "Opportunities for Optical Protection and Restoration," in *Proc., OFC '98*, vol. 2, (San Jose, CA), pp. 269–270, February 1998.

[5] P. Demeester et al., "Resilience in Multilayer Networks," *IEEE Comm. Mag.*, vol. 37, pp. 70–75, Aug. 1999.

[6] R.D. Doverspike, S. Phillips and J.R. Westbrook, "Future Transport Network Architectures," *IEEE Comm. Mag.*, vol. 37, pp. 96–101, August 1999.

[7] S. Ramamurthy and B. Mukherjee, "Survivable WDM Mesh Networks, Part I-Protection," in *Proc., IEEE INFOCOM '99*, (New York, NY), pp. 744–751, March 1999.

[8] S. Ramamurthy and B. Mukherjee, “Survivable WDM Mesh Networks, Part II-Restoration,” in *Proc., ICC '99*, (Vancouver, Canada), pp. 2023–2030, June 1999.

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

[10] T. Wu, “Emerging Technologies for Fiber Network Survivability,” *IEEE Comm. Mag.*, vol. 33, pp. 58–74, February 1995.

| **Avg. Number of sub-paths per primary path** | **Number of demands satisfied** (%) | **Avg. length of Primary path (# of links)** | **Avg. length of Backup path (# of links)** | **Avg. length of Backup path of Sub-paths (# of links)** |
|---|---|---|---|---|
| 1.00(1) | 58.47 | 2.52 | 3.93 | 3.93 |
| 1.66(2) | 40.33 | 3.07 | 6.75 | 4.07 |
| 1.99(3) | 34.71 | 1.88 | 8.81 | 4.43 |
| 2.03(4) | 34.76 | 1.88 | 8.88 | 4.35 |
| 2.04(5) | 35.02 | 1.97 | 8.91 | 4.37 |

Table 1: Performance of sub-path protection when $M = 1.0$ and $W = 16$, for the 14 node network.

| **Avg. Number of sub-paths per primary path** | **Number of demands satisfied** (%) | **Avg. length of Primary path (# of links)** | **Avg. length of Backup path (# of links)** | **Avg. length of Backup path of Sub-paths (# of links)** |
|---|---|---|---|---|
| 1.00(1) | 81.07 | 2.59 | 3.99 | 3.99 |
| 1.70(2) | 57.05 | 3.03 | 7.15 | 4.21 |
| 2.08(3) | 47.86 | 2.03 | 9.25 | 4.45 |
| 2.13(4) | 48.08 | 1.98 | 9.33 | 4.38 |
| 2.14(5) | 48.29 | 2.06 | 9.39 | 4.39 |

Table 2: Performance of sub-path protection when $M = 1.0$ and $W = 24$, for the 14 node network.

| **Avg. Number of sub-paths per primary path** | **Number of demands satisfied (%)** | **Avg. length of Primary path (# of links)** | **Avg. length of Backup path (# of links)** | **Avg. length of Backup path of Sub-paths (# of links)** |
|---|---|---|---|---|
| 1.00(1) | 91.91 | 2.65 | 4.01 | 4.01 |
| 1.73(2) | 71.87 | 3.01 | 7.31 | 4.22 |
| 2.15(3) | 59.23 | 2.16 | 9.61 | 4.47 |
| 2.20(4) | 59.55 | 2.13 | 9.68 | 4.40 |
| 2.21(5) | 59.75 | 2.10 | 9.75 | 4.41 |

Table 3: Performance of sub-path protection when $M = 1.0$ and $W = 40$, for the 14 node network.

| **Avg. Number of sub-paths per primary path** | **Number of demands satisfied (%)** | **Avg. length of Primary path (# of links)** | **Avg. length of Backup path (# of links)** | **Avg. length of Backup path of Sub-paths (# of links)** |
|---|---|---|---|---|
| 1.00(1) | 34.86 | 3.97 | 5.15 | 5.15 |
| 1.78(2) | 32.38 | 3.91 | 6.99 | 3.88 |
| 2.45(3) | 30.70 | 3.72 | 8.45 | 3.45 |
| 2.76(4) | 29.84 | 3.60 | 9.11 | 3.30 |
| 3.15(5) | 29.07 | 3.59 | 9.83 | 3.12 |

Table 4: Performance of sub-path protection when $M = 1.0$ and $W = 40$, for the 54 node network.

| **Avg. Number of sub-paths per primary path** | **Number of demands satisfied (%)** | **Avg. length of Primary path (# of links)** | **Avg. length of Backup path (# of links)** | **Avg. length of Backup path of Sub-paths (# of links)** |
|---|---|---|---|---|
| 1.00(1) | 90.05 | 5.02 | 6.09 | 6.09 |
| 1.90(2) | 83.32 | 4.76 | 8.21 | 4.32 |
| 2.80(3) | 78.12 | 4.73 | 10.36 | 3.70 |
| 3.25(4) | 76.26 | 4.72 | 11.37 | 3.50 |
| 3.85(5) | 74.29 | 4.68 | 12.47 | 3.24 |
| 4.08(6) | 73.28 | 4.47 | 12.97 | 3.18 |

Table 5: Performance of sub-path protection when $M = 0.6$ and $W = 60$, for the 54 node network.

| **Avg. Number of sub-paths per primary path** | **Number of demands satisfied (%)** | **Avg. length of Primary path (# of links)** | **Avg. length of Backup path (# of links)** | **Avg. length of Backup path of Sub-paths (# of links)** |
|---|---|---|---|---|
| 1.00(1) | 98.11 | 5.02 | 6.20 | 6.20 |
| 1.92(2) | 98.12 | 5.23 | 8.53 | 4.44 |
| 2.85(3) | 98.05 | 5.17 | 10.83 | 3.80 |
| 3.37(4) | 97.80 | 4.98 | 12.01 | 3.58 |
| 4.08(5) | 97.22 | 4.88 | 13.38 | 3.28 |
| 4.73(6) | 96.41 | 4.80 | 14.61 | 3.09 |

Table 6: Performance of sub-path protection when $M = 0.4$ and $W = 80$, for the 54 node network.

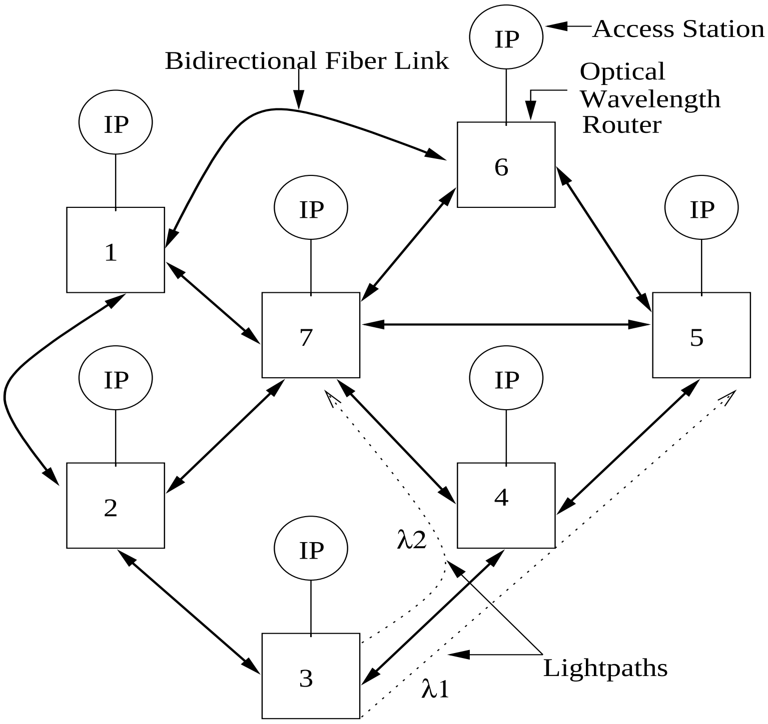


Figure 1: Architecture of a WDM Network.

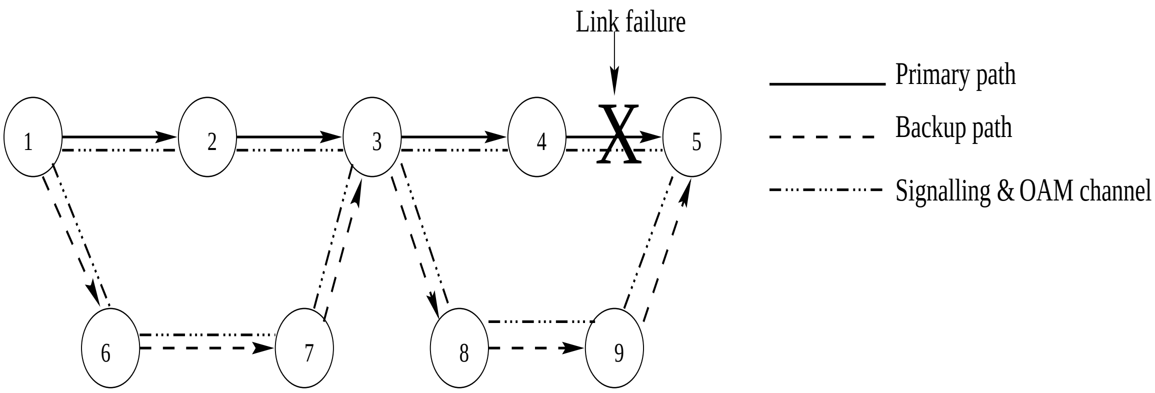


Figure 2: Protection model.

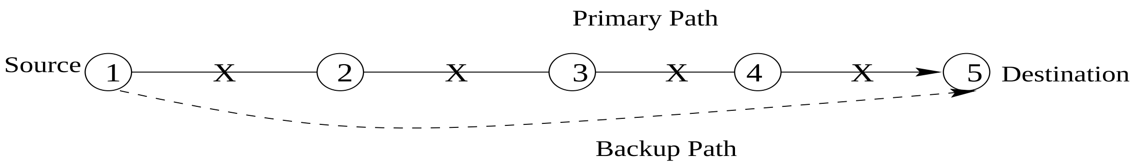


(a) Link Disjoint Path Protection

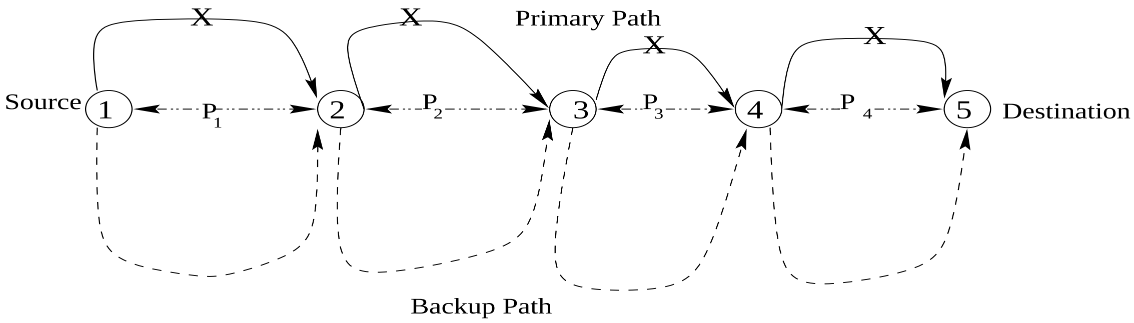


(b) Link Disjoint Link Protection (Special Case of Sub-path protection, Sub-path length = 1)

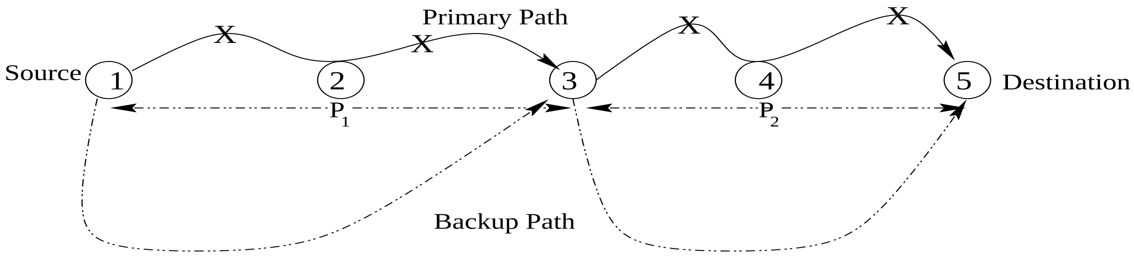


(c) Link Disjoint Sub-Path Protection (Sub-Path length = 2)

Figure 3: Illustration of the different protection schemes.

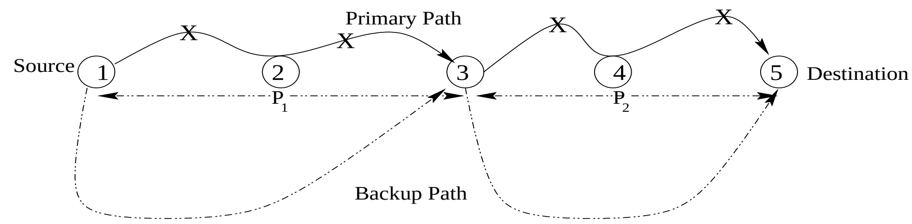


(a) Link Disjoint Sub-Path Protection (Sub-Path length = 2)

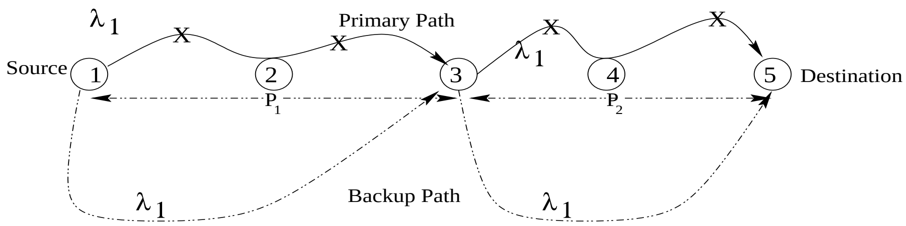


(b) Primary and backup paths use the same wavelength while using 1:1 protection

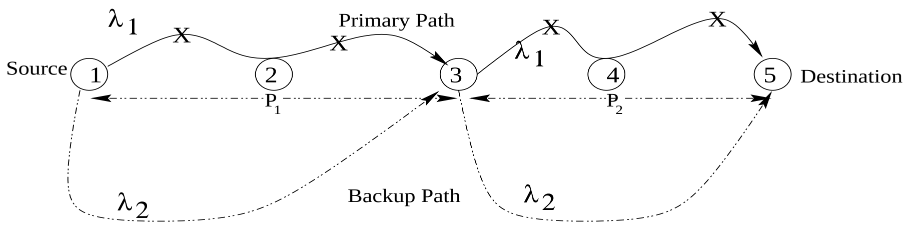


(c) Primary and backup paths use different wavelengths while using 1:1 protection but wavelengths used for the protection of each sub-path is the same

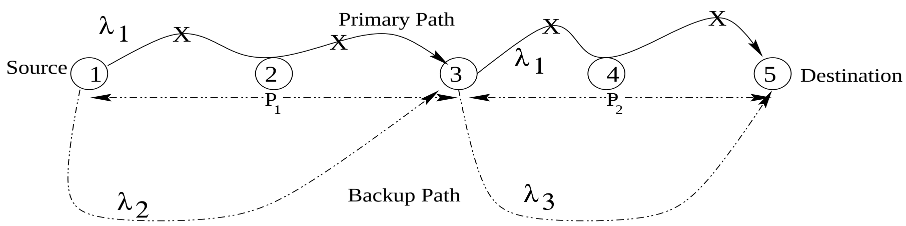


(d) Primary and backup paths use different wavelengths while using 1:1 protection but the wavelengths used for the protection of each sub-path is different.

Figure 4: Wavelength assignment dependent variations of the sub-path protection scheme.

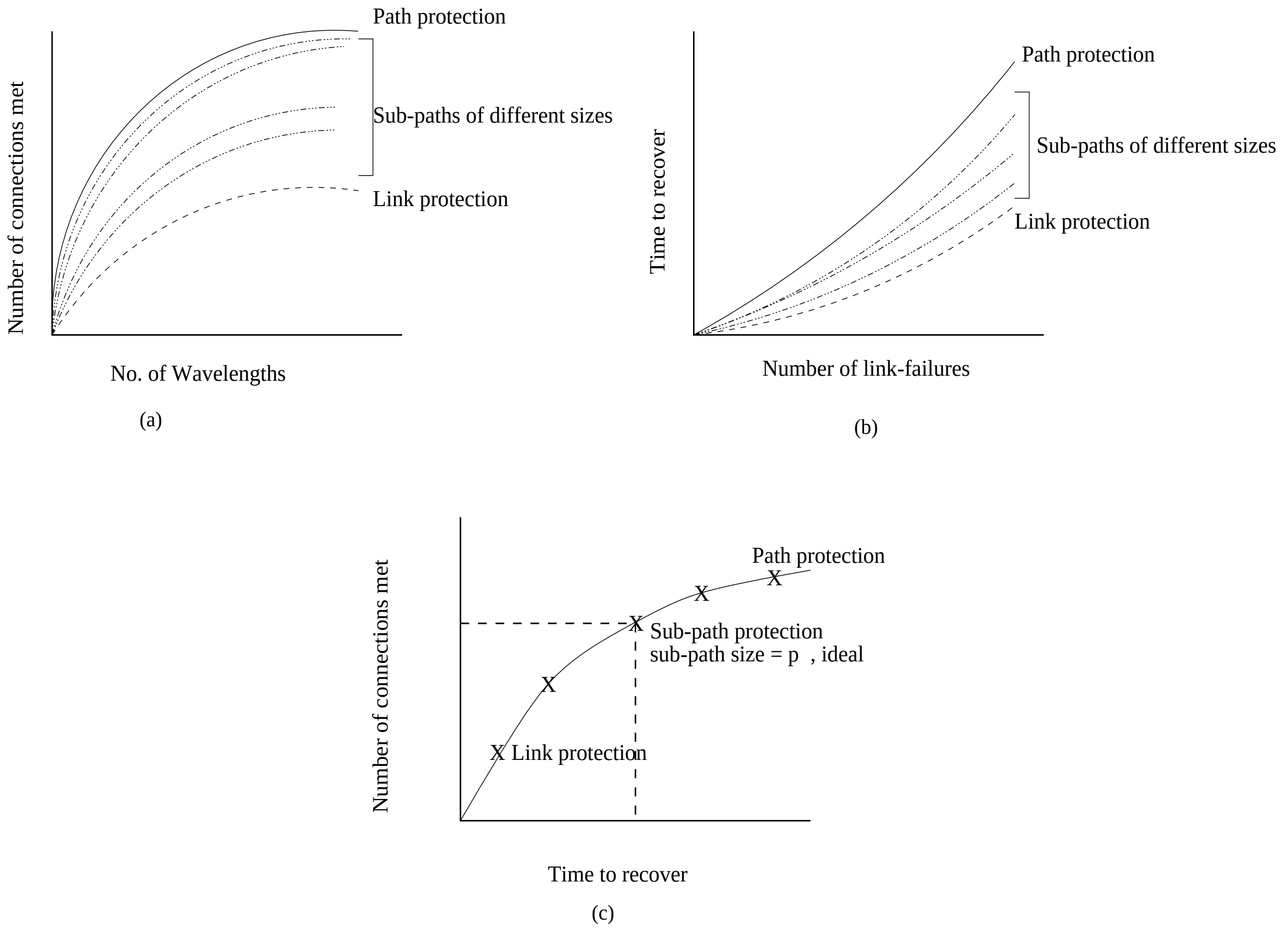


Figure 5: Illustration of the typical capacity utilization and recovery times of each of path, sub-path and link protection schemes.

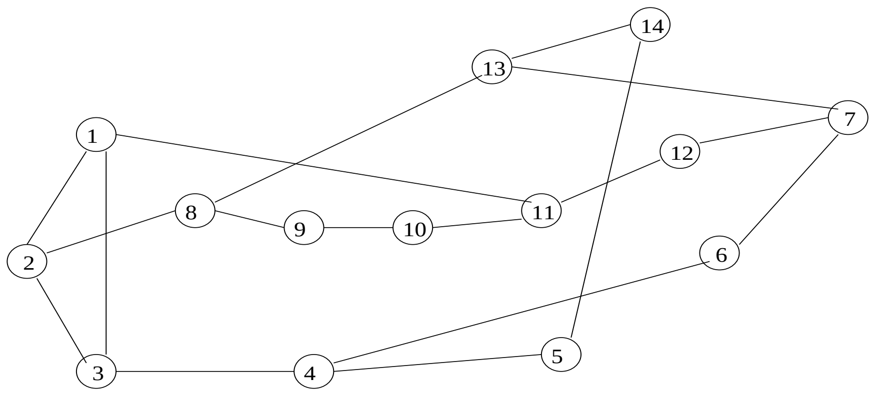


Figure 6: 14 node NSFNET backbone.

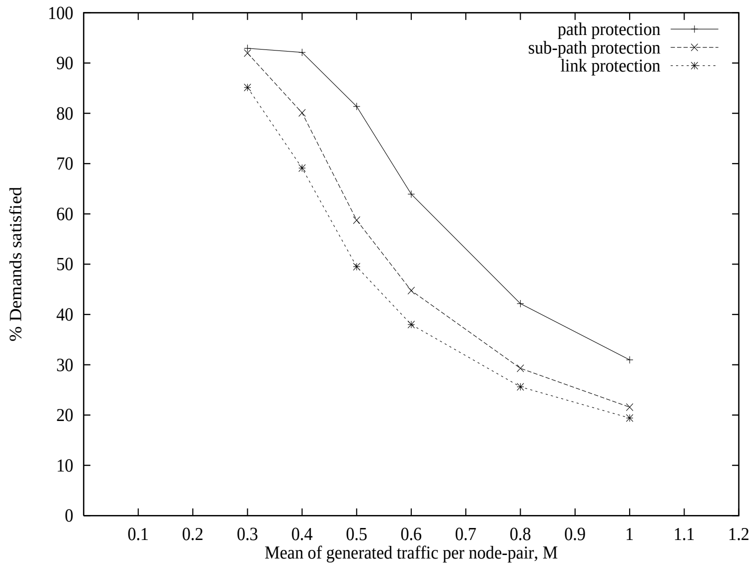


Figure 7: Capacity utilization of protection schemes, 14 node network.

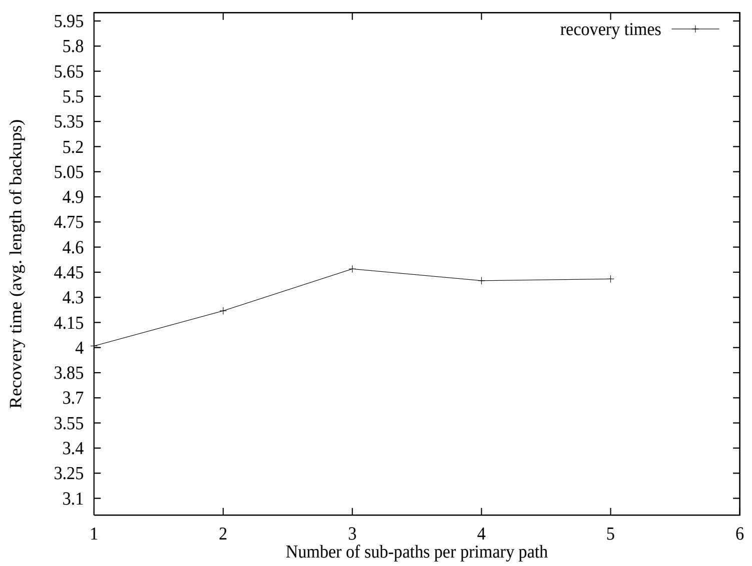


Figure 8: Recovery times of protection schemes, 14 node network.

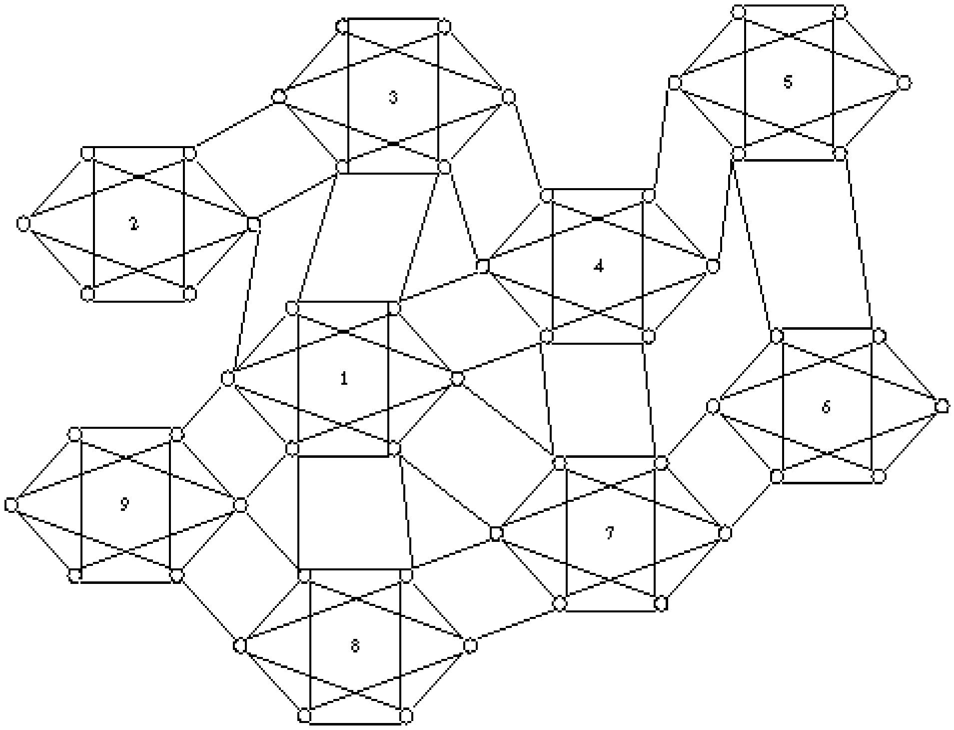


Figure 9: 54 node network.

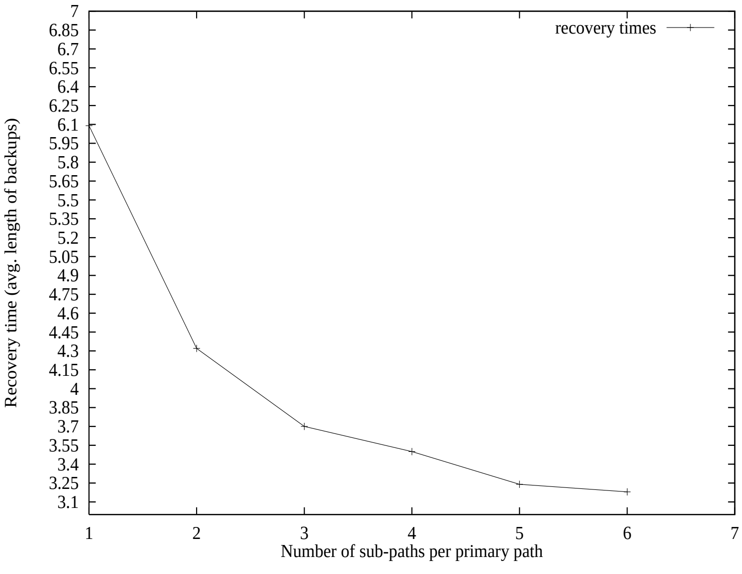


Figure 10: Recovery times of protection schemes, 54 node network.