

MIT Open Access Articles

Survivable paths in multilayer networks

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Parandehgheibi, Marzieh, Hyang-Won Lee, and Eytan Modiano. "Survivable paths in multilayer networks." In 2012 46th Annual Conference on Information Sciences and Systems (CISS), 1-6. Institute of Electrical and Electronics Engineers, 2012.

As Published: <http://dx.doi.org/10.1109/CISS.2012.6310813>

Publisher: Institute of Electrical and Electronics Engineers

Persistent URL: <http://hdl.handle.net/1721.1/81419>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Survivable Paths in Multilayer Networks

Marzieh Parandehgheibi, Hyang-Won Lee and Eytan Modiano
 Laboratory for Information and Decision Systems
 Massachusetts Institute of Technology
 Cambridge, MA 02139

Abstract—We consider protection problems in multilayer networks. In single-layer networks, a pair of disjoint paths can be used to provide protection for a source-destination pair. However, this approach cannot be directly applied to layered networks where disjoint paths may not always exist. In this paper, we take a new approach which is based on finding a *set of paths* that may not be disjoint but together will survive any single physical link failure. First, we consider the problem of finding the minimum number of survivable paths. In particular, we focus on two versions of this problem: one where the length of a path is restricted, and the other where the number of paths sharing a fiber is restricted. We prove that in general, finding the minimum survivable path set is NP-hard, whereas both of the restricted versions of the problem can be solved in polynomial time. We formulate the problem as Integer Linear Programs (ILPs), and use these formulations to develop heuristics and approximation algorithms. Next, we consider the problem of finding a set of survivable paths that uses the minimum number of fibers. We show that this problem is NP-hard in general, and develop heuristics and approximation algorithms with provable approximation bounds. Finally, we present simulation results comparing the different algorithms.

I. INTRODUCTION

Multilayer network architecture such as IP-over-WDM has played an important role in advancing modern communication networks. Typically, a layered network is constructed by embedding a logical topology onto a physical topology such that each logical link is routed using a path in the physical topology. While such a layering approach enables to take advantage of the flexibility of upper layer technology (e.g., IP) and the high data rates of lower layer technology (e.g., WDM), it raises a number of challenges for efficient and reliable operations. In this paper, we focus on the issue of providing protection in layered networks.

The protection problem in single-layer networks is rather straightforward; namely, providing a pair of disjoint paths (one for primary and one for backup) guarantees a route between two nodes against any single link failure. This approach, however, cannot be directly applied to layered networks, because a pair of seemingly disjoint paths at the logical layer may share a physical link and thus simultaneously fail in the event of a physical link failure. To address this issue, [1] introduced the notion of *physically disjoint* logical paths.

In [2], this notion was generalized as Shared Risk Link Group (SRLG) disjoint paths, i.e., two paths between the source and destination nodes that do not share any risk (e.g.,

fiber and conduit). Nearly all the previous works in the context of layered network protection have focused on finding SRLG-disjoint paths [3], [4], [5], [6].

Although the SRLG-disjoint paths problem has been well studied, there are several challenges to this approach. First, SRLG-disjoint paths may not always exist. Second, such a pair of paths could be very long and thus vulnerable. Third, by associating appropriate cost to a path, the SRLG-disjoint paths problem can be modified to find a path set avoiding long paths. However, the modified problem is only known to be NP-hard [2] and there is no known algorithm with provable approximation guarantee.

In order to address these challenges, we take an alternative approach that is based on finding a set of paths that together will survive any single physical link failure. Thus, in the case that SRLG-disjoint paths do not exist, we may find three or more paths such that in the event of a fiber failure, at least one of the paths remain connected. This notion of *survivable path set* generalizes the traditional notion of SRLG-disjoint paths, and enables to provide protection for a broader range of scenarios. Our contributions can be summarized as follows:

- We introduce a new notion of survivable path set so as to provide protection even for the case where SRLG-disjoint paths do not exist.
- We prove the NP-hardness of various protection problems that seek to find a survivable path set with different objectives, and identify the conditions for the protection problems to become polynomial time solvable.
- We develop heuristics and approximation algorithms for the survivable path set problems.

In Section II, we present the network model. In Section III, we study the problem of finding a minimum set of paths that will survive any single fiber failure and develop several approximation algorithms. In Section IV, we design approximation algorithms for finding a survivable path set that uses the minimum number of fibers. Finally, we provide simulation results in Section V and conclusions in Section VI.

II. NETWORK MODEL

We consider a layered network that consists of a logical topology $G_L = (V_L, E_L)$ built on top of a physical topology $G_P = (V_P, E_P)$ where V and E are the sets of nodes and links respectively. Each logical link (i, j) in E_L is mapped onto an $i - j$ path in the physical topology. This is called lightpath routing. Different lightpaths may use the same fiber (physical link), therefore when a fiber fails, all the lightpaths

using that fiber will fail. Hence, a logical path survives the failure of any fiber that it does not use.

As mentioned above, we generalize the traditional notion of SRLG-disjoint paths to account for the case where there does not exist a pair of SRLG-disjoint paths. In a layered network, a set of logical paths is said to be *survivable* if at least one of the paths remain connected after any single physical link failure. Hence, a survivable set consisting of two paths is a pair of SRLG-disjoint paths. Note that there may exist a survivable path set while SRLG-disjoint paths do not exist. For example, consider the physical and logical topologies in Fig. 1. Each dashed line in Fig. 1(c) shows the lightpath routing of each logical link over the physical topology. Under this lightpath routing, each pair of logical paths between nodes 1 and 4 shares some fiber.

Suppose that we want to find a set of logical paths between nodes 1 and 4 in Fig. 1 that can survive any single physical link failure. Clearly, there does not exist a pair of SRLG-disjoint paths as each pair of logical paths shares a fiber. However, it is straightforward to check that the set of 3 paths can survive any single fiber cut, although they are not SRLG-disjoint. This example shows that the traditional protection schemes based on SRLG-disjoint paths (such as the ones in [2]) may fail to provide protection against single physical link failures, while there exists a set of paths that can together provide protection. Our goal in this paper is to address the problem of finding a set of survivable paths that together will survive any single fiber failure.




Fig. 1. Topologies in Multilayer Networks

III. MINIMUM SURVIVABLE PATHS SET (MSP)

We start with the problem of finding a minimum survivable path set, i.e., the minimum cardinality set of paths between a pair of nodes s and t that survive any single physical link (fiber) failure. We first present a path-based Integer Linear Program (ILP) formulation for this problem, assuming that the entire set of $s - t$ paths with their routings over fibers is given. For each path j , let P_j be a binary variable which takes the value 1 if path j is selected, and 0 otherwise. The matrix $A \in R^{m \times n}$ refers to the mapping of all n paths over the m fibers such that $a_{ij} = 0$ if path j uses fiber i and $a_{ij} = 1$

otherwise. Let e be a $m \times 1$ vector of ones.

$$\text{minimize} \quad \sum_{j=1}^n P_j \quad (1)$$

$$\text{subject to} \quad A \times P \geq e \quad (2)$$

$$P_j \in \{0, 1\}, \quad j = 1, \dots, n \quad (3)$$

In the above, the objective function is the number of selected paths. Each row $i \in \{1, \dots, m\}$ in constraint (2) requires that at least one selected path survives the failure of fiber i , i.e., the selected path set should be survivable. Hence, the optimal solution to the above optimization problem gives a minimum survivable path set. Although this formulation requires the knowledge of every path (which is possibly exponential in the number of fibers), the compact and clean expression of the path-based formulation enables us to analyze the useful properties of survivable path sets. Later, we will use this formulation to develop heuristics and approximation algorithms for finding a minimum survivable path set.

The MSP problem can also be formulated using a polynomial number of constraints and variables without enumerating all of the paths. Let P_{tot} denote the number of selected $s - t$ logical paths, E_L denote the set of logical links and E_L^k denote the set of remaining logical links after the failure of fiber k . Note that for survivability, each E_L^k should contain at least one of the selected paths. Let x_{ijk} be 1 if link (i, j) in E_L^k is selected to form an $s - t$ path over the remaining graph $G_L^k = (V_L, E_L^k)$, and 0 otherwise. Let y_{ij} be 1 if the selected path set uses logical link (i, j) , and 0 otherwise. The following link-based formulation describes the MSP problem.

$$\text{minimize} \quad P_{tot} \quad (4)$$

$$\text{subject to} \quad \left. \begin{aligned} \sum_{(s,j) \in E_L} y_{sj} &= P_{tot} \\ \sum_{(i,t) \in E_L} y_{it} &= P_{tot} \\ \sum_{(i,j) \in E_L} y_{ij} - \sum_{(j,i) \in E_L} y_{ji} &= 0, \quad \forall i \neq s, t \end{aligned} \right\} \quad (5)$$

$$\left. \begin{aligned} \sum_{(s,j) \in E_L^k} x_{sjk} &= 1, & \forall k \\ \sum_{(i,t) \in E_L^k} x_{itk} &= 1, & \forall k \\ \sum_{(i,j) \in E_L^k} x_{ijk} - \sum_{(j,i) \in E_L^k} x_{jik} &= 0 & \forall k, \forall i \neq s, t \end{aligned} \right\} \quad (6)$$

$$y_{ij} \geq x_{ijk} \quad \forall k, i, j \quad (7)$$

$$x_{ijk} \in \{0, 1\} \quad \forall k, i, j \quad (8)$$

The constraints in (6) require that each remaining logical graph E_L^k should contain an $s - t$ path, which guarantees the survivability against any single physical link failure. By the constraints in (7), logical link (i, j) is selected if it has been used in some remaining logical graph $G_L^k = (V_L, E_L^k)$. Hence,

the constraints in (5) require that there should be total P_{tot} flows between nodes s and t over the selected logical links specified by y_{ij} 's. Consequently, the variable P_{tot} counts the total number of paths selected for survivability. In Section V, we will use this formulation to verify the performance bound of our approximation algorithms.

A. MSP in general setting

In this section, we show that the MSP problem is NP-hard in general and discuss some algorithms that can be used to solve the problem. In Sections III-B and III-C, we will study the MSP problem under practical constraints. Our first result pertains to the complexity of the MSP problem as stated in Theorem 1 below.

Theorem 1: Computing the minimum number of survivable paths in multilayer networks is NP-hard. In addition, this minimum value cannot be approximated within any constant factor, unless $P = NP$.

The proof of Theorem 1 relies on a mapping between the survivable path set problem and the minimum set cover problem. Suppose that each path corresponds to a set of fibers that are not used by that path, i.e., survived. Then, finding a minimum survivable path set is equivalent to finding a minimum path set that survives (covers) all of the fibers. The complete proof can be found in Appendix A.

Since the problem is computationally hard to solve, we consider heuristics and approximation algorithms that give a set of survivable paths in polynomial time. Owing to the similarity to the set cover problem, the heuristics that have been developed for set cover problems can be used here. In particular, a common approach to solve the set cover problem is the greedy algorithm. In order to apply the greedy algorithm to our setting, one needs to enumerate all of the paths with their routings on the fibers. In general, the number of paths in a multilayer network is exponential in the total number of fibers. Moreover, in each iteration, the greedy algorithm tries to find a path that survives the maximum number of fibers. This is equivalent to the Minimum Color Path problem, which is known to be NP-hard. [7]

Another approach which can be used to approximate the set cover problem is randomized rounding. Randomized rounding gives an $O(\log m)$ approximation, where m is the number of fibers [8]. This is the best possible approximation for the MSP problem, which is due to the fact that the minimum set cover problem cannot be approximated within better than a $\log m$ factor [9].

Fortunately, practical systems impose certain physical constraints that make the survivable path-set problem easier to solve. For example, due to physical impairments and delay constraints, paths are typically limited in length. Furthermore, in WDM networks, the sharing of a fiber by the logical links is limited by the number of available wavelengths. In the following, we show that these physical limitations make the MSP problem tractable.

B. The Path Length Restricted Version

In this section, we assume that each logical path is restricted to use at most K fibers. Restricting the length of paths (i.e., number of fibers on each path) is a realistic assumption, because each logical link is typically constrained in the number of fibers that it may use, and each logical path is constrained in the number of logical links.

Lemma 1: Under the path length restriction, the optimal number of survivable paths is at most $K + 1$.

Proof: By the assumption, each path uses at most K fibers, and thus at least $m - K$ fibers are survived by a path. Suppose that we have selected an arbitrary path, and want to add other paths to form a survivable path set. In the worst case, each of the newly selected paths can survive only a single fiber which is not survived by the previously selected paths. Since there are at most K fibers that are not survived by the first path, we need at most K additional paths to survive the rest of the fibers. Therefore, the total number of paths will not exceed $K + 1$. ■

Lemma 2: In the path length restricted version of MSP, the total number of paths is polynomial in the number of fibers m , and can be enumerated in polynomial time.

Proof: Under the assumption, a path can consist of up to K fibers, and thus at most K logical links. In a graph with n nodes there can be $O(n^K)$ paths of length up to K . Since the number of nodes is at most $2m$, the total number of logical paths of length up to K is $O(m^K)$. A simple exhaustive search can be used to enumerate the paths. ■

Theorem 2: The path length restricted version of the MSP problem can be solved in polynomial time.

Proof: By Lemma 1, MSP needs at most $K + 1$ paths to survive any single failure. Therefore, one can find the exact solution by searching through all subsets of paths with sizes $2, 3, \dots, K + 1$. This will take $O(P^{K+1})$ iterations where P is the total number of paths. On the other hand, by Lemma 2, the total number of paths is $O(m^K)$. Therefore, the total running time of exhaustive search is $O(m^{K(K+1)})$ which is polynomial in the total number of fibers. ■

Although this exhaustive search returns an optimal solution, its running time can be prohibitive for large values of m and K . This motivates us to study heuristics and approximation algorithms with better running time. First, we consider a greedy algorithm, followed by a randomized algorithm based on ϵ -net which is a well-known technique in the area of computational geometry.

1) *Greedy Algorithm:* The first heuristic we consider is a greedy algorithm which is similar to the greedy algorithm for the minimum set cover problem. The input to the greedy algorithm is the set of paths with the set of fibers used by each path and the set of all fibers. The greedy algorithm is an iterative algorithm that works as follows. In the first iteration, it selects a path using the minimum number of fibers, and updates the set of fibers not survived by the selected path. This greedy path selection is repeated until the selected path set survives all of the fibers. Following the proof of Lemma 1, it

can be shown that the greedy algorithm also finds a survivable path set with size at most $K + 1$.

As discussed in Section III-A, the greedy algorithm generally gives an $O(\log m)$ approximation to the minimum survivable path set. However, under the assumption of restricted path length, it provides a better approximation as stated in Theorem 3.

Theorem 3: The greedy algorithm provides an $O(\log K)$ approximation in polynomial time for the path length restricted version of MSP.

Proof: Let ξ be the size of minimum survivable path set. Let n_i be the number of fibers that are not survived after the i^{th} iteration of the greedy algorithm. Clearly, we have $n_1 \leq K$. Now, note that there is a path that survives at least $\frac{n_1}{\xi}$ of the remaining n_1 fibers, because otherwise the size of the optimal path set would be larger than ξ . Hence, in the second iteration, the greedy algorithm would select a path that survives at least $\frac{n_1}{\xi}$ of fibers. Thus,

$$n_2 \leq n_1 - \frac{n_1}{\xi} \leq K(1 - \frac{1}{\xi}). \quad (9)$$

Similarly,

$$n_3 \leq n_2 - \frac{n_2}{\xi} \leq K(1 - \frac{1}{\xi})^2, \quad (10)$$

and in general,

$$n_i \leq K(1 - \frac{1}{\xi})^i. \quad (11)$$

The greedy algorithm will terminate when $n_t < 1$, and this condition is satisfied when

$$K(1 - \frac{1}{\xi})^t < 1, \quad (12)$$

where t is the total number of iterations. Since $1 - x < e^{-x}$ for $x > 0$, inequality (12) is satisfied when

$$Ke^{-\frac{t}{\xi}} \leq 1 \Leftrightarrow t \leq \xi \times \log K. \quad (13)$$

Therefore, the greedy algorithm provides an $O(\log K)$ approximation.

To prove the polynomial time complexity, note that in each iteration of the greedy algorithm, the best path can be found in $O(m^K)$ by searching through all the paths (see the proof of Theorem 2). Furthermore, as mentioned above, the greedy algorithm terminates in at most $K + 1$ iterations. Therefore, the computational complexity of the greedy algorithm is $O(Km^K)$. ■

Although the greedy algorithm runs significantly faster than the exhaustive search algorithm, its running time can still be prohibitive for large K and m . Hence, we develop a novel randomized algorithm which has a considerably better running time. This algorithm builds upon solutions to the closely related Set Cover and Hitting Set problems. In particular, the algorithm is based on ε -net, a concept in computational geometry, which provides an approximation algorithm for the Hitting Set problem.

2) *ε -net Algorithm:* Our ε -net algorithm is an iterative algorithm which selects each path with some probability. If all the fibers are survived by the selected path set in the first iteration, the algorithm terminates. Otherwise, it changes the probability of selecting each path and selects a new set of paths using the new probabilities, until all fibers are survived.

Let W_j be the weight of path j , initialized as $W_j = 1$. Define the weight of each fiber i to be the sum of the weights of paths surviving fiber i , i.e.,

$$W(f_i) = \sum_{j:a_{ij}=1} W_j. \quad (14)$$

Definition 1: A fiber is said to be ε -Survivable if

$$W(f_i) \geq \varepsilon \sum_{j=1}^n W_j \text{ for some } \varepsilon \in (0, 1), \quad (15)$$

where n is the total number of paths.

Note that when all the paths have the same weight of 1, a fiber is ε -Survivable if it is survived by at least $\varepsilon \times n$ paths. Hence, if a fiber is ε -Survivable with large ε , then it is likely to be survived by randomly selected paths. This observation is exploited in our ε -net algorithm as discussed below.

By applying the randomized algorithm for the hitting set problem from [10] and [11], we can obtain a path-selection algorithm for selecting a random subset of paths that will survive all of the ε -Survivable fibers, with high probability. In particular, the algorithm finds a set of paths via s independent random draws (with replacement), such that in each draw, a path is selected from the entire path set according to the probability distribution $\mu(P_j) = \frac{W_j}{\sum_{j=1}^n W_j}, \forall j$.

Our ε -net algorithm iteratively applies this random path selection as follows. After each iteration, it checks the survivability of the selected path set. If not all fiber failures are survived, the algorithm doubles the weight of all paths that survive the failure of fibers in \tilde{S} , where \tilde{S} is the set all the fibers that are not survived yet (so that such fibers are more likely to be survived by the new path set). The random path selection is repeated with the new probability distribution.

Let ξ be the optimal solution to the MSP problem. By applying the results in [12], [13], the following theorem can be proved.

Theorem 4: Assume $s = c \frac{\log K}{\varepsilon} \log \frac{\log K}{\varepsilon}$, where c is a constant. The ε -net algorithm finds a set of survivable paths of size $O(\log K \log \xi) \xi$, with high probability.

This theorem together with Lemma 1 implies that the ε -net algorithm finds a survivable path set of size $O(\log^2 K) \xi$. Moreover, it can be shown that the algorithm requires $O(K \log(\frac{m}{K}))$ iterations to achieve this performance bound. On the other hand, the path-selection algorithm needs to select $O(\frac{\log K}{\varepsilon} \log \frac{\log K}{\varepsilon})$ paths in each iteration. Therefore, the computational complexity of the ε -net algorithm is $O(K \log(K) \log(m) \log(\log(K)))$. Table I summarizes the performance of each algorithm under the path length restriction.

Method	Approximation	Running Time	T
ExS	Exact Solution	$O(m^{K(K+1)})$	D
Greedy	$O(\log K)$	$O(Km^K)$	D
ε -net	$O(\log K \log \xi)$	$O(K \log(K) \log(m) \log(\log(K)))$	P

TABLE I

PERFORMANCE BOUNDS UNDER PATH LENGTH RESTRICTED VERSION:
EXS-EXHAUSTIVE SEARCH, T-TYPE, D-DETERMINISTIC,
P-PROBABILISTIC

C. Wavelength Restricted version

Another important practical constraint is that in WDM-based networks, the number of lightpaths using a fiber is limited to say W , which is the number of wavelengths supported over a fiber. In this section, we assume that a set of logically disjoint paths with their mapping on the physical topology is given, and the goal is to find a minimum survivable path set among those paths under the WDM restriction. Note that the set of logically disjoint paths can be abstract to a logical topology with two nodes and parallel links (e.g., the one in Fig. 2(a)). Clearly, in this setting, the WDM restriction implies that *each fiber can be used by at most W paths*. Using this property, it can be shown that the MSP problem under the WDM restriction can be solved in polynomial time. To prove this, we need the following lemma.

Lemma 3: Under the wavelength restriction, the minimum number of survivable paths is at most $W + 1$.

Proof: Suppose that the minimum survivable path set contains more than $W + 1$ paths. This implies that there exists a fiber whose failure disconnects at least $W + 1$ paths (so that more than $W + 1$ paths are needed for survivability), which contradicts to the fact that under the WDM restriction, each fiber can be used by at most W paths. ■

Theorem 5: Under the wavelength restriction, the MSP problem can be solved in polynomial time.

Proof: It can be shown that under the WDM restriction, the given path set contains $O(m)$ paths. By Lemma 3, we only need to enumerate path sets of size up to $W + 1$ in order to find a minimum survivable path set. Clearly, this can be done in $O(m^{W+1})$ time. More details can be found in Appendix C. ■

Although there exists a polynomial time optimal algorithm, it requires excessive computation for large values of W and m . As in the case of restricted path length, we have developed approximation algorithms with better running time. Table II shows the summary of our approximation algorithms under the wavelength restriction (See D for details).

Method	Approximation	Running Time	T
ExS	Exact Solution	$O(W^{W+1}m^{W+1})$	D
Greedy	$O(\log m)$	$O(W^2m)$	D
ε -net	$O(\log W \log \xi)$	$O(W \log(W) \log(m) \log(\log(W)))$	P

TABLE II

APPROXIMATION BOUNDS UNDER WAVELENGTH RESTRICTED VERSION:
EXS-EXHAUSTIVE SEARCH, T-TYPE, D-DETERMINISTIC,
P-PROBABILISTIC

IV. MINIMUM NUMBER OF PHYSICAL FIBERS IN SURVIVABLE PATHS (MFSP)

Our focus so far has been on providing protection using the minimum number of paths. In this section, our goal is to find a survivable path set that uses the minimum number of fibers. This problem seems to have a direct connection to the minimum cost survivable path set problem where the cost of a path is the number of fibers used by that path. However, this is not true owing to the fact that costs of paths are not additive, i.e., a fiber that is used by multiple paths only adds one unit of cost. In order to make this point clear, consider Fig. 2. A minimum cost survivable path set problem will find paths 1 and 2 as the set of survivable paths with total cost 7, while the MFSP problem will find paths 2, 3 and 4 as the optimal survivable path which has the total cost 6. In the next section we will develop ILP formulations, and analyze the complexity of MFSP.




Fig. 2. Routing in Multilayer Network

A. ILP Formulation and Complexity

We start with an ILP formulation of the problem. Similar to the MSP problem, the MFSP problem can be formulated in several different ways, but here we only present the path-based formulation which will be used for developing heuristics and approximation algorithms. Given the set of paths and associated fibers, for each path j , assign a binary variable P_j which takes the value 1 if path j is selected and 0 otherwise. Similarly, for each fiber i , assign a binary variable f_i which takes the value 1 if fiber i is selected and 0 otherwise. The matrix A and vector e are defined in the same way as in the MSP formulation (1)-(3).

$$MFSP : \quad \text{minimize} \quad \sum_{i=1}^m f_i \quad (16)$$

$$\text{subject to} \quad A \times P \geq e \quad (17)$$

$$f_i \geq P_j \quad \forall f_i \in P_j \quad (18)$$

$$P_j \in \{0, 1\} \quad \forall P_j \quad (19)$$

In the above, the objective function is the number of fibers used by the selected paths. Again, the constraints in (17) require the selected path to be survivable. The constraints in (18) relate the selected paths and fibers, such that a fiber is selected if at least one of the paths using the fiber is selected. Clearly,

the optimal solution to the above optimization problem gives a set of survivable paths that use the minimum number of fibers. This MFSP problem can be shown to be NP-hard.

Theorem 6: Computing the set of survivable paths using the minimum number of physical fibers is NP-hard.

Proof: We provide a mapping from the Minimum 3-Set Cover problem, which is a special version of the Set Cover problem where each set has exactly 3 elements, to the MFSP problem. The Minimum 3-Set Cover problem is NP-hard, and holds all the inapproximability properties of the Minimum Set Cover problem.




Fig. 3. Physical Topology

Consider an instance of the Minimum Set Cover problem with the ground set E and a family of subsets F . Suppose that each subset in F contains only 3 elements. To show a mapping, we construct a physical topology as shown in Fig. 3, such that each node on the left corresponds to a subset in $F = \{C_1, \dots, C_{|F|}\}$ and the nodes on the right are the elements of $E = \{e_1, \dots, e_m\}$. Node j on the left is connected to node i on the right if and only if $e_i \in C_j$. Note that a node on the left is connected to only three nodes on the right (i.e., each set contains only three elements).

We can construct a logical topology and its lightpath routing over the physical topology; such that for protection, we need to have m paths from s to t that pass through all the nodes on the right. Moreover, since each path between s and the nodes on the left uses a large number of fibers, we should select a survivable path set that uses the minimum number of nodes on the left. Consequently, the minimum fiber survivable path set for the aforementioned layered network gives a minimum set cover for the given instance of E and F , which shows the NP-hardness of the MFSP problem. For the complete proof, see Appendix F. ■

Since the MFSP problem is a reduction from the minimum 3-set cover problem, it is unlikely that the MFSP problem has an efficient optimal algorithm. For this reason, we develop new heuristics and approximation algorithms. In particular, as in the previous section, we focus on the practical scenario where the number of paths on a fiber is at most W , i.e., the wavelength restricted setting. We first present a greedy algorithm, and then a randomized rounding algorithm based on the path-based formulation for MFSP.

B. Additive Cost Greedy Algorithm (ACG)

Recall that the goal is to find a survivable path set that uses the minimum number of fibers. Hence, it is desired to select a path that uses a small number of fibers while surviving many new fibers (i.e., fibers not survived by already selected paths) as possible. Note that this is clearly different from the MSP problem where the number of fibers does not matter. The Additive Cost Greedy algorithm requires the set of paths and associated fibers as input. We define a new cost metric in order to take into account the two factors simultaneously. Let C_j be the number of fibers used by path j . The ‘‘amortized cost’’ AC_j of path j , which is updated for every iteration, is defined as follows:

$$AC_j = \frac{C_j}{\#\text{newly survived fibers by } P_j},$$

where the denominator is the number of fibers survived by path j and not survived by the previously selected paths. Our greedy algorithm selects a path with minimum amortized cost, updates the amortized costs of the remaining paths, and continue until all the fibers are survived. This greedy algorithm, which we call the Additive Cost Greedy algorithm, gives an approximate solution.

Theorem 7: The Additive Cost Greedy algorithm provides an $O(W \log m)$ approximation to the MFSP problem.

Note that the number C_j in the additive cost of path j does not change over iterations. That is, the additive cost implicitly assumes that selecting path j will add C_j fibers to the total cost, while only the number of new fibers is added to the total cost. Therefore, one can better take into account the actual change to the cost by updating C_j as the number of fibers that are used by path j and not used by the previously selected paths. In Section V, we will show that this Non-additive Cost Greedy (NACG) algorithm works better than the ACG, by finding survivable path sets with fewer fibers.

C. Randomized Rounding Algorithm

Randomized rounding is a widely used technique to solve difficult integer optimization problems. In general, randomized rounding scheme solves the Linear Program (LP) relaxation of the original ILP formulation, and rounds the solution randomly. In our case, the LP relaxation of the MFSP problem is given as

$$\text{LP relaxation: minimize } \sum_{i=1}^m f_i \quad (20)$$

$$\text{subject to } A \times P \geq e \quad (21)$$

$$f_i \geq P_j \quad \forall f_i \in P_j \quad (22)$$

$$0 \leq P_j \leq 1. \quad (23)$$

Let P_j^* and f_i^* be the optimal values of path j and fiber i . Note that the above path-based LP uses the set of paths and associated fibers as input. Our randomized rounding algorithm to solve the MFSP problem works as follows:

- 1) Initialize $S = \emptyset$. Solve the relaxed problem.

- 2) Select each path j with probability P_j^* , and add it to S if selected.
- 3) Repeat step 2 for T times.

Since paths are selected randomly, some fibers may not be survived in one iteration. Clearly, as the number of iterations T increases, the probability of surviving all of the fibers increases. On the other hand, it may increase the number of selected paths and thus fibers. Therefore, the parameter T determines the survivability probability and the approximation quality of the solution. The following theorem characterizes this relationship.

Theorem 8: With $T = O(\log \frac{m}{1-q})$ iterations, the randomized rounding algorithm gives an $O(W \log \frac{m}{1-q})$ approximation with probability at least q .

Proof: We first find an upper bound on the expected number of fibers selected in each iteration (which gives the approximation quality of the solution), and then, the probability of survivability is derived.

1) *Expected Number of Selected Fibers:* Note that fiber i is selected if any of the paths using the fiber is added to the path set S . Moreover, in each iteration, each path j is added with probability P_j^* . To count the number of selected fibers, define a random variable F_i for each fiber i such that $F_i = 1$ if fiber i is selected and 0 otherwise. The expected number of fibers selected in each iteration can be written as

$$E[\sum_{i=1}^m F_i] = \sum_{i=1}^m \Pr(F_i = 1) = \sum_{i=1}^m (1 - \Pr(F_i = 0)) \quad (24)$$

Therefore, we need to compute $\Pr(F_i = 0)$, which is the probability of a fiber not being selected. Note again that a fiber is not selected if none of the paths using the fiber are selected. It follows that

$$\Pr(F_i = 0) = \prod_{j:a_{ij}=0} (1 - P_j^*) \quad (25)$$

$$\geq \prod_{j:a_{ij}=0} (1 - f_i^*) \quad (\text{by constraint (22)}) \quad (26)$$

where the equality is due to the independence of path selections. Let w_i be the number of paths that use fiber f_i , i.e., $w_i = |\{j : a_{ij} = 0\}|$. Then, we can obtain

$$\Pr(F_i = 0) \geq \prod_{j:a_{ij}=0} (1 - f_i^*) = (1 - f_i^*)^{w_i}, \quad (27)$$

$$\Pr(F_i = 1) \leq 1 - (1 - f_i^*)^{w_i}. \quad (28)$$

Finally, by using the fact that $1 - (1 - x)^n \leq nx$, the probability of selecting a fiber can be upper-bounded as

$$\Pr(F_i = 1) \leq w_i f_i^*. \quad (29)$$

Combining (24) and (29) yields the following bound on the

expected number of fibers selected in each iteration:

$$\begin{aligned} E[\sum_{i=1}^m F_i] &\leq \sum_{i=1}^m w_i f_i^* \\ &\leq W \times \sum_{i=1}^m f_i^* \quad (\text{by wavelength restriction}) \\ &= W \times LP^*, \end{aligned} \quad (30)$$

where LP^* is the optimal value of the LP relaxation.

2) *Probability of Survivability:* Next, we derive an upper bound on the probability that the selected path set is *not* survivable, by applying the idea of the feasibility argument in [8]. First, for each fiber i , the probability that the selected path set cannot survive the failure of fiber i can be written as follows:

$$\Pr(\text{fiber } i \text{ not survived in one iteration}) \quad (31)$$

$$= \Pr(\text{none of paths surviving fiber } i \text{ are picked}) \quad (32)$$

$$= \prod_{j:a_{ij}=1} (1 - P_j^*) \quad (33)$$

$$\leq \prod_{j:a_{ij}=1} e^{-P_j^*} \quad \text{using } (1 - x \leq e^{-x}) \quad (34)$$

$$\leq e^{-\sum_{j:a_{ij}=1} P_j^*} \leq \frac{1}{e}. \quad (\text{using constraint 21}) \quad (35)$$

Since the randomized rounding runs for T iterations with $T = \log \frac{m}{1-q}$, we can obtain

$$\Pr(f_i \text{ not covered in all iterations}) \leq \frac{1}{e^{\log \frac{m}{1-q}}} = \frac{1-q}{m}. \quad (36)$$

Thus, by the union bound,

$$\Pr(\text{there exist an unsurvived fiber}) \leq m \times \frac{1-q}{m} = 1 - q. \quad (37)$$

3) *Approximation Result:* By (30), the total expected number of fibers after T iterations is bounded as

$$E[\text{Total \# fibers}] \leq W \log \frac{m}{1-q} LP^* \quad (38)$$

Since the solution is in integer form, it is an upperbound for the ILP solution. Thus, with probability at least q ,

$$\frac{E[\text{Total \# fibers}]}{W \log \frac{m}{1-q}} \leq ILP \leq E[\text{Total \# fibers}]. \quad (39)$$

■

D. Random-Sweep Greedy (RSG)

Next, we present a new Greedy algorithm for the MFSP problem, which is called the Random-Sweep greedy. Unlike the Greedy algorithm discussed in Section III, the RSG removes a path (from the selected path set) which survives the fibers covered by other paths; so that the size of the path set can be further reduced while maintaining the survivability. Although we could not quantify the performance of this algorithm, it performs near optimally in some scenarios as will be shown in Section V.

Method	Approximation	Running Time	T
ExS	Exact	$O(W^{W+1}m^{W+1})$	D
ACG	$W \log m$	$O(W^2m)$	D
ε -net	$W \log W \log \xi$	$O(W \log(W) \log(m) \log(\log(W)))$	P
RR	$W \log \frac{m}{1-p}$	$\log \frac{m}{1-p}$	P
RSG	nearly Opt.	$O(W^2m)$	D

TABLE III

APPROXIMATION BOUNDS UNDER WAVELENGTH RESTRICTED VERSION OF MFSP: EXS-EXHAUSTIVE SEARCH, RSG-RANDOM SWEEP GREEDY, ACG-ADDITIVE COST GREEDY, RR-RANDOMIZED ROUNDING, T-TYPE, D-DETERMINISTIC, P-PROBABILISTIC

The RSG algorithm also requires the knowledge of the set of paths and associated fibers. Let S_j be the set of fibers that are survived by path j . Moreover, let the cost C_j of path j in each iteration be the number of fibers that are used by path j , and not used by the previously selected paths. Using the cost function C_j , define the amortized cost AC_j as the ratio of C_j to the number of newly survived fibers by path j . The first two iterations of RSG are the same as the Non-Additive Cost greedy algorithm. That is, in each iteration, it selects a path with minimum amortized cost. If the first two paths survive all of the fibers, the algorithm terminates. Otherwise, it continues as follows.

Suppose the RSG algorithm is in the i^{th} iteration. First, find a path, say i , with minimum amortized cost among the remaining paths. Then, pick a path, say j , randomly from the previously selected paths and find $S_j \cup S_i$, which is the set of fibers that are survived by either path i or path j . If there exists a path k among the previously selected paths such that $S_k \subset S_i \cup S_j$, remove path k from the selected paths. Note that removing such a path does not affect the survivability of the selected path set, i.e., the same set of fibers are still survived after the removal. More importantly, we can possibly decrease the number of fibers used by the selected paths.

Table III shows the summary of our algorithms for the MFSP problem. Note that we have also developed an ε -net algorithm and its details can be found in H.

V. SIMULATION RESULTS

We compare the performance of our algorithms using both large-scale random network topologies, as well as the US backbone network topology. In particular, we compare the following algorithms:

- ILP-based optimal algorithm computed by CPLEX; denoted by ILP
- Simple Greedy algorithm from Section III-B1; denoted by MSPG
- Additive Cost Greedy algorithm from Section IV-B; denoted by ACG
- Non-additive Cost Greedy algorithm from Section IV-B; denoted by NACG
- Random-Sweep Greedy algorithm from Section IV-D; denoted by RSG
- Randomized rounding algorithm from Section IV-C; denoted by RR

- ε -net algorithm from Section III-B2; denoted by EPS

A. Performance in Large-scale Random Topologies

We first consider a random layered network where the logical topology consists of 50 paths between nodes s and t . This layer is mapped onto the physical topology containing 100 fibers, using the mapping structure shown in [14]. In the wavelength restricted version of the problem, at most W paths can be assigned to each fiber. For each value of W , we generate 1000 random topologies each with 50 paths that are randomly routed on the physical topology. We then apply our algorithms to each network in order to find a survivable path set using the minimum number of fibers (i.e., to solve the MFSP problem). Note that for Randomized Rounding the performance depends on the survivability guarantee of the algorithm, which is 99.9% for the results shown below.




Fig. 4. Comparison of algorithms for MFSP problem: Approximation quality in random networks

Fig. 4 compares the average number of fibers in the survivable path set found by each algorithm. It can be seen that as the value of W increases, the number of used fibers increases. This is due to the fact that when W is large, more logical paths can share a fiber, and therefore more logical paths are needed since a single physical link failure can lead to a large number of logical path failures. Note that the Random-Sweep Greedy (RSG) algorithm is nearly optimal, and the performance of ε -net algorithm is better than RSG for large values of W . Fig. 5 compares the logarithm of the running time of the algorithms. It can be seen that the Randomized Rounding algorithm is the fastest, while the RSG algorithm which gives the closest to optimal solution, and the ε -net algorithm which performs nearly optimally for networks with large values of W , have larger running times. Note also that the running times are nearly independent of W for all of the proposed algorithms. In contrast, obtaining the exact optimal solution using the ILP formulation becomes quickly impractical as W increases.

Next, we consider larger networks where there are 1000 fibers in the physical topology and 500 paths in the logical topology, with W ranging from 1 to 40. Fig. 6 shows the performance of the various algorithms as a function of W . The performance of the ILP-based algorithm is omitted since CPLEX often fails to find a solution within a reasonable




Fig. 5. Run Time Comparison of Different heuristics with respect to Optimal

amount of time. Again we see that the RSG algorithm considerably outperforms the rest of algorithms.




Fig. 6. Comparison of Approximation Algorithms in Large Networks

B. Performance in Real Networks




Fig. 7. Physical Topology

Next, we examine the performance of the approximation algorithms over the US backbone topology shown in Fig. 7, with the objective of finding a minimum survivable path set between nodes 4 and 22 [15]. For the logical topology, we generated random graphs with eight nodes (including 4 and 22) each of degree 4. We use shortest path lightpath routing for the logical links.

Table IV shows the average number of paths and average running time of each algorithm. It can be seen that the RSG and randomized rounding algorithms are nearly optimal, and furthermore, the randomized rounding gives a solution almost instantly. We also note that the survivability guarantee of the

Method	Number of Paths	Running Time (ms)
ILP	2.0069	7.2133
RSG	2.0160	2.0167
RR	2.0482	0.0272
MSPG	2.2241	0.1911
EPS	2.551	1.6000

TABLE IV
COMPARISON OF ALGORITHMS FOR MSP IN REAL NETWORKS

Randomized Rounding algorithm is 99% for the results shown in the table.

VI. CONCLUSION

We considered the problem of finding survivable paths in layered networks. The traditional disjoint paths approach for protection cannot be directly applied to layered networks, since physically disjoint paths may not always exist in such networks. To address this issue, we introduced the new notion of *survivable path set*. We showed that in general the problem of finding the minimum size survivable path set (MSP) and the problem of finding the minimum fiber survivable path set (MFSP) are NP-hard and inapproximable. However, under practical constraints, we are able to develop both optimal and approximation algorithms for the MSP and MFSP problems. An important future direction is to develop backup routing schemes based on survivable path sets.

REFERENCES

- [1] R. Bhandari, "Optimal diverse routing in telecommunication fiber networks," in *IEEE INFOCOM*, June 1994.
- [2] J. Q. Hu, "Diverse routing in optical mesh networks," *IEEE Transactions on Communications*, vol. 51, no. 3, pp. 489–494, March 2003.
- [3] D. Xu et al., "Trap avoidance and protection schemes in networks with shared risk link groups," *Lightwave Technology*, vol. 21, no. 11, pp. 2683–2693, Nov. 2003.
- [4] D. Xu, Y. Xiong, and C. Qiao, "A new promise algorithm in networks with shared risk link groups," in *IEEE Globecom*, Dec. 2003.
- [5] P. Datta and A. K. Somani, "Diverse routing for shared risk resource group (srrg) failures in wdm optical networks," in *First Annual International Conference on Broadband Networking*, Oct. 2004, pp. 120–129.
- [6] D. Xu et al., "On finding disjoint paths in single and dual link cost networks," in *IEEE INFOCOM*, 2004.
- [7] S. Yuan, S. Varma, and J. P. Jue, "Minimum color path problems for reliability in mesh networks," in *IEEE INFOCOM*, 2005.
- [8] P. Raghavan and C. D. Thompson, "Randomized rounding: a technique for provably good algorithms and algorithmic proofs," in *Combinatorica*, vol. 7, no. 4, 1987, pp. 365–374.
- [9] C. Lund and M. Yannakakis, "On the hardness of approximating minimization problems," *ACM*, vol. 41, no. 5, pp. 960–1071, Sep. 1994.
- [10] K. Varadarajan, "Epsilon nets and union complexity," in *25th annual symposium on Computational geometry*, Jun. 2009.
- [11] J. Matousek, *Lectures on Discrete Geometry*. NJ: Springer-Verlag, 2002.
- [12] G. Even, D. Rawitz, and S. Shahar, "Hitting set when the vc-dimension is small," *Information Processing Letters*, vol. 95, no. 2, pp. 358–362, Jul. 2005.
- [13] H. Bronnimann and M. Goodrich, "Almost optimal set covers in finite vc-dimension," *Discrete Computational Geometry*, vol. 14, no. 1, pp. 463–479, Oct. 1995.
- [14] K. Lee, E. Modiano, and H. Lee, "Cross layer survivability in wdm-based networks," *IEEE/ACM Transactions on Networking*, 2011.
- [15] [Online]. Available: <http://www.optical-network.com/topology.php>
- [16] D. Haussler and E. Welzl, "Epsilon-nets and simplex range queries," in *Second Annual Symposium on Computational Geometry*, 1986.
- [17] V. Chvatal, "A greedy heuristic for the set covering problem," *Mathematics of Operations Research*, vol. 4, no. 3, pp. 233–235, Aug. 1979.

APPENDIX A
PROOF OF COMPLEXITY OF MSP

The Minimum survivable paths problem can be reduced from the NP-hard Minimum Set Cover problem. Given an instance of Minimum Set Cover Problem with ground set E and family of subsets R , we construct a physical topology $E = \{f_1, \dots, f_m\}$ containing all m fibers and a logical topology $R = \{P_1, \dots, P_n\}$, where each P_j corresponds to the set of fibers that survive in the failure of path j , i.e. all fibers that are *not* used by path j . It follows that the minimum number of logical paths that survives all the physical fibers is equal to the size of a minimum set cover. As the last step of proof, we need to show we can construct a physical topology with the given routing. Given the set of paths and the fibers used by each path (complement of fibers survived by each path), we can use the physical topology in [14]. The inapproximability result follows immediately from the inapproximabilities of the Minimum Set Cover problem.

APPENDIX B
PROOF OF THEOREM 4

In the procedure of ε -net algorithm, the “path-selection” algorithm will be applied iteratively, and checks the survivability of the selected path set after each iteration. If not all fibers are survived, the algorithm doubles the weight of all paths that survive the failure of fibers in \bar{S} , where \bar{S} is the set all the fibers that are not survived yet, and repeat the random path selection.

Let ξ be the optimal solution of MSP. Based on the results in [12], [13], it can be shown that if in each iteration the selected subset of paths survive a “good” subset of fibers, in $O(\xi \log(\frac{m}{\xi}))$ iterations, the algorithm will return a set of survivable paths, with high probability. A subset is “good” if it is an ε -net.

Definition 2: Consider a set system $F = (X, R)$, where X is the set of elements and R is the set of subsets of X . A set $H \subset X$ is an “ ε -net” of F if $S \cap H \neq \emptyset$, for every subset $S \in R$ for which $|S| \geq \varepsilon|X|$.

Lemma 4 claims that it is guaranteed that in each iteration the selected paths survive a “good” subset of fibers

Lemma 4: For all $\varepsilon \in (0, \frac{1}{2})$, if $s = c \frac{\log K}{\varepsilon} \log \frac{\log K}{\varepsilon}$, where c is a constant, the path-selection algorithm selects a subset of paths that survives all of the ε -Survivable fibers with high probability.

For the proof of Lemma 4, we use the new techniques found by Haussler and Welzl in [16]. Theorem 9 is an improvement on their work [11], [12].

Before presenting Theorem 9, we need to define VC-dimension.

Definition 3: Let R be a set system on a set X . Let us say that a subset $A \subset X$ is shattered by R if each of the subsets of A can be obtained as the intersection of some $S \in R$ with A , i.e. if $R|_A = 2^A$.

Define the VC-dimension of R , denoted by $\dim(R)$, as the supremum of the sizes of all finite shattered subsets of X . If

arbitrarily large subsets can be shattered, the VC-dimension is ∞ .

Theorem 9: Let $F = (X, R)$ denote a set system with weights $w(u)$. For every $\varepsilon \in (0, \frac{1}{2})$, a random sample of X according to the probability distribution $w(u) = w(X)$ is likely to be an ε -net with respect to $w(u)$, if the sample contains $O(\frac{d}{\varepsilon} \log(\frac{d}{\varepsilon}))$ elements, where d is the VC-dimension of the set system.

To prove Lemma 4, it is enough to show that $O(\frac{\log K}{\varepsilon} \log(\frac{\log K}{\varepsilon}))$ paths are needed to cover all ε -Survivable fibers.

Let $X = \{P_1, \dots, P_n\}$ be the set of paths in our problem, and $R = \{f_1, \dots, f_m\}$ be the set of subsets of X , where each fiber-set f_i corresponds to the set of paths that survives the failure of fiber i . Therefore, an ε -net will cover all ε -Survivable fibers.

In this setting, a subset A of paths is shattered if intersection of A with every fiber-set produces all the subsets of A . For instance, $A = \{P_1, P_2, P_3\}$ is shattered by R if there exist 8 fibers $f_1, f_2, f_3, f_4, f_5, f_6, f_7, f_8$ such that $\{A \cap f_1 = P_1, A \cap f_2 = P_2, A \cap f_3 = P_3, A \cap f_4 = \{P_1, P_2\}, A \cap f_5 = \{P_1, P_3\}, A \cap f_6 = \{P_2, P_3\}, A \cap f_7 = A, A \cap f_8 = \emptyset\}$.

Lemma 5: In path length restricted version of MSP, VC-dimension d is less than $\log K$.

Proof: Suppose VC-dimension is d . Then, by definition, there exist a subset of paths A of size d which intersection of A with all fiber-sets generates all subsets of A . In particular, for every $P_j \in A$, half of the subsets created by $A \cap R$ should contain P_j and the other half should not contain it.

Under the path length restricted assumption, each path uses at most K fibers, and survives at least $m - K$ remaining fibers. Therefore, at least $m - K$ fibers contain a particular path j . Thus,

$$2^{d-1} \leq m - K \text{ and } d \leq 1 + \log(m - K). \quad (40)$$

On the other hand, for each P_j at most K fibers do not contain it. Hence,

$$2^{d-1} \leq K \text{ and } d \leq 1 + \log K. \quad (41)$$

By combining both equations 40 and 41, we will have the following result:

$$d \leq 1 + \log K. \quad (42)$$

■

APPENDIX C
PROOF OF WDM POLYNOMIALTY

Under the assumption in Section III-C, we are given a set of logically disjoint paths that satisfy the WDM restriction. As explained, this setting implies that each fiber can be used by at most W paths. Therefore, we have the following Lemma:

Lemma 6: Under the WDM restriction, the number of given paths can be at most $W \cdot m$.

Proof: By the Max-Flow Min-Cut Theorem, in the physical topology the number of disjoint paths between nodes s and t is equal to the minimum $s - t$ cut (MC). On the

other hand, since a fiber can be used by at most W logical links, each physical path can carry at most W logical links. Therefore, the maximum possible number of logical paths is $W \cdot MC \leq W \cdot m$. ■

APPENDIX D EPSILON NET IN WDM

Using the same techniques discussed in the Section B, we have the following Theorem:

Theorem 10: The ε -net algorithm finds a set of survivable paths of size $O(\log W \log \xi)\xi$, with high probability and terminates in $O(\xi \log(\frac{m}{\varepsilon}))$ iterations.

To prove Theorem 10, we need Lemma 7. Then, by an argument similar to the one to prove Thm 4, the ε -net algorithm will find a set of survivable paths with a $\log W \log \xi$ approximation bound.

Lemma 7: $\forall \varepsilon \in (0, 1)$, if $s = c \frac{\log W}{\varepsilon} \log \frac{\log W}{\varepsilon}$ where c is a constant, the path-selection algorithm selects a subset of paths that survives all of the s -Survivable fibers with high probability.

Proof: Similar to the argument for the proof of Theorem 4, it is enough to prove Lemma 8. ■

Lemma 8: In wavelength restricted version of MSP, VC-dimension d is less than W .

Proof: Let $X = \{P_1, \dots, P_n\}$ be the set of all paths, and $R = \{f_1, \dots, f_m\}$ be the set of all fiber sets where a fiber is associated i.e., $P_j \in f_i$ if and only if path j survives fiber i 's failure.

Let VC-dimension be d . Then, by the definition of VC-dimension, there exist a subset A of paths such that $|A| = d$ and intersection of A with all fiber sets generates all the subsets of A . In particular, there exist a fiber set i such that $f_i \cap A = \emptyset$, which means $A \subset X - f_i$. Thus,

$$d = |A| \leq |X - f_i|. \quad (43)$$

On the other hand, under the wavelength restricted assumption, each fiber can be used by at most W paths. Therefore,

$$n - W \leq |f_i|, \quad \forall f_i. \quad (44)$$

Combining inequalities (43) and (44) results in $d \leq W$. ■

APPENDIX E OTHER ILP FORMULATIONS OF MFSP

A. Link-Based Formulation

The idea of Link-Based formulation for MFSP problem is the same as formulation for MSP. The only difference is that we do not need to find the paths in the main logical topology using flow constraints and minimize the number of paths.

For each fiber r , let f_r be a binary variable which takes the value 1 if fiber r is selected, and 0 otherwise. Similar to MSP link-based formulation, variables x_{ijk} refers to the logical links (i, j) and constraint (46) refers to the flow constraints in the remaining logical topology E_L^k correspondent to the failure of fiber k .

$$\text{minimize} \quad \sum_{r=1}^m f_r \quad (45)$$

$$\text{subject to} \quad \left. \begin{aligned} \sum_{(s,j) \in E_L^k} x_{sjk} &= 1 && \forall \text{ Fiber } k \\ \sum_{(i,t) \in E_L^k} x_{itk} &= 1 && \forall \text{ Fiber } k \\ \sum_{(i,j) \in E_L^k} x_{ijk} - \sum_{(j,i) \in E_k} x_{jik} &= 0 && \forall k, \forall i \neq s, t \end{aligned} \right\} \quad (46)$$

$$f_r \geq x_{ijk} \quad \forall i, j, k, \forall f_r \in x_{ijk} \quad (47)$$

$$x_{ijk} \in \{0, 1\} \quad \forall i, j, k \quad (48)$$

Constraint (47) shows the relation between the selected logical links and fibers, such that fiber i is selected if at least a logical link using f_i is selected. The objective function is the sum of selected fibers. Hence, solving this ILP formulation will find a set of survivable paths that uses minimum number of fibers.

B. Cut-Based Formulation

A set of paths will survive any single failure, if in the occurrence of any fiber failure there exist at least one path from s to t .

Let f_r be a binary variable for each fiber r , and y_{ij} be a binary variable for each logical link ij . Let N be the set of all nodes in the logical topology. The objective function is minimizing the total number of selected fibers.

$$\text{minimize} \quad \sum_{r=1}^n f_r \quad (49)$$

$$\text{subject to} \quad \sum_{(i,j) \in E_L^k: i \in S, j \in \bar{S}} y_{ij} \geq 1 \quad \forall k, S \subset N, S \neq N, \emptyset$$

$$s \in S, t \in \bar{S} \quad (50)$$

$$f_r \geq y_{ij} \quad \forall i, j, \forall f_r \in y_{ij} \quad (51)$$

$$y_{ij} \in \{0, 1\} \quad (52)$$

Define “ $s-t$ cut” as a cut $[S, \bar{S}]$ such that $s \in S$ and $t \in \bar{S}$. Constraint (50) shows that for every “ $s-t$ cut” in remaining graph E_L^k , there exist at least one logical link from S to \bar{S} . This will guarantee the survivability of network in the failure of fiber k . Constraint (51) builds the relation between selected logical links and fibers, such that a fiber i will be selected if at least one logical link using f_i is selected. Constraints (50) and (51) select the fibers used by a set of survivable paths and the objective function will find the solution to the MFSP problem.

APPENDIX F
PROOF OF MFSP COMPLEXITY

In the physical topology shown in Figure 3, nodes s and t are the starting and ending nodes. Each node on the left side (n nodes) is connected to 3 nodes on the right such that all the nodes on right are covered by the nodes on left. There are L nodes between s and each node on the left where L is a large number (say $L \geq 3m + 3n$ so that the left hand side should be the first priority when minimizing the used number of fibers) and there are $m \geq 3$ nodes on the tail of the graph such that every node on right connects to the tail through the first node r .

The logical topology is shown in Figure 8.




Fig. 8. Logical Topology

In the logical topology, from s to the nodes on right, each fiber is also a lightpath, while from nodes on right side to t , there are m parallel lightpaths with a specific routing. The first lightpath will be routed on fibers f_1, U_1 and L_2 to L_m , the second lightpath will be routed on fibers f_2, L_1, U_2 and L_3 to L_m and so on. Therefore, lightpath i will use fibers f_i, U_i and all the other L_j s ($j \neq i$).

To survive any single failure in the fibers from right nodes to node t , we need to have at least m paths, each going through one of the parallel logical links. These m paths will not share any fiber from nodes on left to nodes on right, thus any single failure on the fibers between left and right nodes will be survived. Finally, to survive any fiber failure from node s to nodes on the left, it is enough that at least two of paths use disjoint logical links from node s to left nodes.

Consequently, it is enough just to have m paths covering all nodes on the right hand side. On the other hand, paths between s and left nodes use a large number of fibers. To have a set of paths using the minimum number of fibers, we need to pick the minimum number of nodes from left, to cover all the nodes on right which is a mapping from minimum 3-set cover problem to our problem. The remaining of the proof is explained in the main text.

APPENDIX G
PROOF OF THEOREM 7

Before proving Theorem 7, we need to prove two other Lemmas. Consider a set S of survivable paths and let F be

the total number of fibers used by all paths in S .

Lemma 9: Under the wavelength restricted assumption, the following inequality holds: $\frac{1}{W} \sum_{j=1}^{|S|} C_j P_j \leq F$

Proof: Figure 9 shows the relation between paths and the fibers used by them. There exist an edge between node j on left (P_j) and node i on right (f_i) if $f_i \in P_j$. Since each path j is using C_j fibers, the total number of edges is $\sum_{j=1}^{|S|} C_j P_j$. On the other hand, by assumption, each fiber can be used by at most W paths, therefore each right node can be incident to at most W edges. Thus, we have $\sum_{j=1}^{|S|} C_j P_j \leq WF$, which completes the proof.




Fig. 9. relation between paths and fibers

Following theorem is a direct consequence of Lemma 9. ■

Theorem 11: Let S be a set of survivable paths and R be the set of all possible survivable path sets ($S \in R$). Let ξ be the minimum number of fibers used by a survivable path set. Then, we have: $\frac{1}{W} \min_{S \in R} \sum_{j=1}^{|S|} C_j P_j \leq \xi \leq \min_{S \in R} \sum_{j=1}^{|S|} C_j P_j$.

Proof: Lemma 9 gives the left inequality. In the right inequality, $\min_{S \in R} \sum_{j=1}^{|S|} C_j P_j$ outputs a set S of survivable paths, therefore it is feasible and gives an upperbound for the optimal solution. ■

By Theorem 11, the optimal solution to the problem $\min_{S \in R} \sum_{j=1}^{|S|} C_j P_j$ provides a W approximation to the MFSP problem. Note that this problem seeks to find a set of survivable paths with minimum cost where the cost of a path is the number of fibers used by that path, and these costs are assumed to be additive. Clearly, this problem is a reduction from minimum cost set cover problem. Since the minimum cost set cover problem is NP-hard, finding a set of survivable paths with minimum additive costs is also NP-hard. Therefore, we use the explained additive cost greedy algorithm to approximate the additive cost survivable path set problem. Now we can prove the $O(W \log m)$ bound stated in Theorem 7.

Proof: Let ξ be optimal value of MFSP problem. By the argument in [17], the additive cost of paths selected by ACG is not larger than $O(\log m)\xi$, i.e.,

$$\frac{Greedy}{\log m} \leq \min_{P \in \mathcal{S}} \sum_{j=1}^n C_j P_j \quad (53)$$

where *Greedy* denotes the additive cost of ACG. Combining equation (53) and Theorem 11 gives the following inequality:

$$\frac{Greedy}{W \log m} \leq \xi \quad (54)$$

■

APPENDIX H EPSILON-NET IN MFSP

Combining the results of Randomized algorithm described in subsection III-B and Theorem 11 results in the following corollary.

Corollary 1: Using the ε -net algorithm, one can find an $O(W \log W \log \xi)$ approximation.

Proof: Similar to the proof of Theorem 7. ■