

Cross-Layer Survivability in WDM-Based Networks

Kayi Lee, *Member, IEEE*, Eytan Modiano, *Senior Member, IEEE*, Hyang-Won Lee, *Member, IEEE*,

Abstract—In layered networks, a single failure at a lower layer may cause multiple failures in the upper layers. As a result, traditional schemes that protect against single failures may not be effective in cross-layer networks. In this paper, we introduce the problem of maximizing the connectivity of layered networks. We show that connectivity metrics in layered networks have significantly different meaning than their single-layer counterparts. Results that are fundamental to survivable single-layer network design, such as the Max-Flow Min-Cut theorem, are no longer applicable to the layered setting. We propose new metrics to measure connectivity in layered networks and analyze their properties. We use one of the metrics, Min Cross Layer Cut, as the objective for the survivable lightpath routing problem, and develop several algorithms to produce lightpath routings with high survivability. This allows the resulting cross-layer architecture to be resilient to failures between layers.

Index Terms—Cross-layer survivability, max-flow, min-cut, connectivity, disjoint paths, survivable path set, lightpath routing, multi-commodity flow

I. INTRODUCTION

Modern communication networks are constructed using a layered approach, as shown in Figure 1. Such a network typically consists of an electronic packet switched network (such as IP); often this packet-switched network is built on top of one or more electronic circuit switched transport networks (e.g., ATM, SONET; sometimes neither or both); and these in turn are built upon a fiber network. This multitude of layers is used in order to simplify network design and operations. However, this layering also leads to certain inefficiencies and interoperability issues. In this paper, we focus on the impact of layering on network survivability.

We examine this problem in the context of Wavelength Division Multiplexing (WDM) based networks; although the concepts discussed are equally applicable to other layered architectures (e.g., IP over ATM, ATM over SONET, etc.). In a WDM-based network, the logical topology is defined by a set of nodes and lightpaths connecting the nodes, while the physical topology is defined by a (possibly different) set of nodes and the fibers connecting them. For example, an IP-over-WDM network consists of IP routers that are connected using (WDM) lightpaths as shown in Figure 1. Each lightpath is realized by setting up a physical connection using one of the wavelength channels in the optical fibers. For networks where wavelength conversion capability [1] is unavailable, the lightpaths are also subject to the wavelength continuity constraint, which requires the lightpath to use the same wavelength channel along the physical route [2].

The authors are with the Massachusetts Institute of Technology, Cambridge, MA 02139. (e-mail: {kylee, modiano, hwlee}@mit.edu)

This work was supported by NSF grants CNS-0626781 and CNS-0830961 and by DTRA grant number HDTRA1-07-1-0004.

This work was presented in part at the IEEE INFOCOM conference, April 2009.

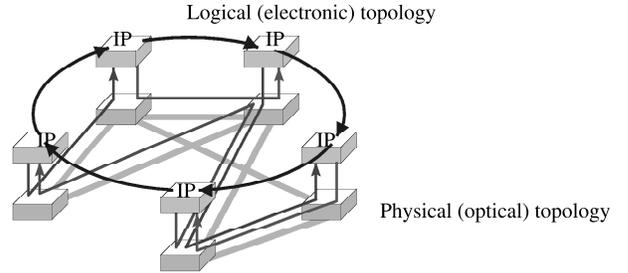


Fig. 1. An IP-over-WDM network where the IP routers are connected using optical lightpaths. The logical links (arrowed lines on top) are formed using lightpaths (arrowed lines at the bottom) that are routed on the physical fiber (thick gray lines at the bottom). In general, the logical and physical topologies are not the same.

Networks often rely on the logical layer for providing protection and restoration services. However, even when the logical topology is designed to tolerate single logical link failures, once the logical topology is embedded on the physical topology, the logical topology may no longer be survivable to single physical (fiber) link failures. This is because each physical fiber link may carry multiple lightpaths. Hence, the failure of a single fiber link can lead to the failure of multiple links in the logical topology, which may subsequently leave the logical topology disconnected.

As a simple illustrative example, consider the physical and logical topologies shown in Figures 2(a) and (b). The lightpaths in the logical topology are routed over the physical topology in two different ways in Figures 2(c) and (d). In Figure 2(c), a failure of physical fiber (1, 5) would cause lightpaths (1, 5) and (3, 5) to fail. Consequently, node 5 will be disconnected from other nodes in the logical topology. On the other hand, in Figure 2(d), the logical topology will remain connected even if one of the fibers fails. The above example demonstrates that in a multi-layer network, a physical link failure can result in multiple logical link failures, and that the routing of the logical links on the physical topology has a big impact on the connectivity of the multi-layer network.

In contrast to the simplified example of Figure 2, real-life networks are highly intertwined and layered. However, due to the lack of general understanding of the issues in cross-layer survivability, most existing protection and restoration mechanisms are based on principles that are applicable only to single-layer network environments, and are subject to cross-layer issues as illustrated above.

Nearly all previous literature in cross-layer survivability focused on single fiber failure, where the problem of interest generally falls into two broad categories: finding primary and protection routes for a single lightpath [3]–[6], or routing multiple lightpaths jointly for a given logical topology [7]–[16]. The problem studied in this paper belongs to the second

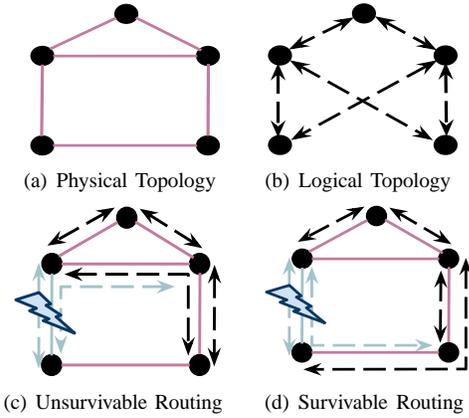


Fig. 2. Different lightpath routings can affect survivability

category, with the focus on designing lightpath routings that are tolerant against the maximum number of physical failures. Some of the concepts introduced in this paper are generalizations of the *survivable lightpath routing* concept, which was first introduced in [7]. The same paper also developed an ILP formulation for survivable routing of arbitrary logical topologies, which was subsequently improved in [10] and [15]. The problem of routing logical rings survivably on the physical network was studied in [7], [13], [14]. In [11] the authors introduced the notion of piecewise survivable mapping and developed an algorithm to compute survivable routings based on piecewise survivable components. The same technique was extended to compute lightpath routings that are survivable against k failures, for a fixed value of k [17].

To the best of our knowledge, this is the first paper that formally studies classical survivability theory in the context of layered networks. We show that standard survivability metrics, such as the minimum cut and maximum disjoint paths, that have been widely used in characterizing the survivability properties of single-layer networks, lose much of their meaning in the context of cross-layer architecture. In particular, the Max-Flow Min-Cut Theorem, which constitutes the foundations of network survivability theory and provides the mathematical justification of the aforementioned metrics, no longer holds in the cross-layer context. Such a fundamental difference suggests that many basic issues of cross-layer survivability are largely not understood.

In Section II, we investigate some combinatorial properties of layered graphs related to network survivability, and highlight the key difference from their single-layer counterparts. In Section III, we specify the requirements for cross-layer survivability metrics, and propose two new metrics, Min Cross Layer Cut and Weighted Load Factor, that measure the connectivity of multi-layer networks. In Section IV, we consider the survivable lightpath routing problem using the Min Cross Layer Cut as the objective, and develop several lightpath routing algorithms based on the multi-commodity flow formulation in order to maximize the cross-layer connectivity of the network. In Section V, we present the simulation results for these algorithms, along with some empirical studies of the metrics introduced in Section III.

II. FLOWS, CUTS AND PATH SETS IN LAYERED GRAPHS

In this section, we will study connectivity structures such as flows, cuts and paths in multi-layer graphs from a theoretical standpoint, in order to develop insights into cross-layer survivability. We will highlight the key difference in combinatorial properties between multi-layer graphs and single-layer graphs. In particular, we will show that fundamental survivability results, such as the Max Flow Min Cut Theorem, are no longer applicable to multi-layer networks. Consequently, metrics such as “connectivity” have significantly different meanings in the cross-layer setting. Such fundamental differences make it much more challenging to design survivable multi-layer networks.

A. Max Flow vs Min Cut

For single-layer networks, the Max-Flow Min-Cut Theorem [18] states that the maximum amount of flow passing from the source s to the sink t always equals the minimum capacity that needs to be removed from the network so that no flow can pass from s to t . In addition, if all links have integral capacity, then there exists an integral maximum flow. This implies the maximum number of disjoint paths between s and t is the same as the minimum cut between the two nodes. Hence, the term *connectivity* between two nodes can be used unambiguously to refer to different measures such as maximum disjoint paths or minimum cut, and this makes it a natural choice as the standard metric for measuring network survivability.

Because of its fundamental importance, we would like to investigate the Max-Flow Min-Cut relationship for multi-layer networks. We first generalize the definitions of *Max Flow* and *Min Cut* for layered networks:

Definition 1: In a multi-layer network, the *Max Flow* between two nodes s and t in the logical topology is the maximum number of physically disjoint $s - t$ paths in the logical topology. The *Min Cut* between two nodes s and t in the logical topology is the minimum number of physical links that need to be removed in order to disconnect the two nodes in the logical topology.

We model the physical topology as a network graph $G_P = (V_P, E_P)$, where V_P and E_P are the nodes and links in the physical topology. The logical topology is modelled as $G_L = (V_L, E_L)$, where $V_L \subseteq V_P$. The lightpath routing is represented by a set of binary variables f_{ij}^{st} , where a logical link (s, t) uses physical fiber (i, j) if and only if $f_{ij}^{st} = 1$. Let \mathcal{P}_{st} be the set of all $s - t$ paths in the logical topology. For each path $p \in \mathcal{P}_{st}$, let $L(p)$ be the set of physical links used by the logical path p , that is, $L(p) = \cup_{(s,t) \in p} \{(i, j) | f_{ij}^{st} = 1\}$. Then the Max Flow and Min Cut between nodes s and t can be formulated mathematically as follows:

$$\begin{aligned} \text{MaxFlow}_{st} : \quad & \text{Maximize} \sum_{p \in \mathcal{P}_{st}} f_p, \quad \text{subject to:} \\ & \sum_{p: (i,j) \in L(p)} f_p \leq 1 \quad \forall (i, j) \in E_P \quad (1) \\ & f_p \in \{0, 1\} \quad \forall p \in \mathcal{P}_{st} \end{aligned}$$

$$\begin{aligned} \text{MinCut}_{st} : \text{Minimize } & \sum_{(i,j) \in E_P} y_{ij}, \quad \text{subject to:} \\ & \sum_{(i,j) \in L(p)} y_{ij} \geq 1 \quad \forall p \in \mathcal{P}_{st} \\ & y_{ij} \in \{0, 1\} \quad \forall (i,j) \in E_P \end{aligned} \quad (2)$$

The variable f_p in the formulation MaxFlow_{st} indicates whether the path p is selected for the set of (s, t) -disjoint paths. Constraint (1) requires that no selected logical paths share a physical link. Similarly, in the formulation MinCut_{st} , the variable y_{ij} indicates whether the physical fiber (i, j) is selected for the minimum (s, t) -cut. Constraint (2) requires that all logical paths between s and t traverse some physical fiber (i, j) with $y_{ij} = 1$.

Note that the above formulations generalize the the Max Flow and Min Cut for single-layer networks. In particular, the formulations model the classical Max Flow and Min Cut of a graph G if both G_P and G_L are equal to G , and $f_{ij}^{st} = 1$ if and only if $(s, t) = (i, j)$.

Let MaxFlow_{st} and MinCut_{st} be the optimal values of the above Max Flow and Min Cut formulations. We also denote MaxFlow_{st}^R and MinCut_{st}^R to be the optimal values to the linear relaxations of above Max Flow and Min Cut formulations. The Max-Flow Min-Cut Theorem for single-layer networks can then be written as follows:

$$\text{MaxFlow}_{st} = \text{MaxFlow}_{st}^R = \text{MinCut}_{st}^R = \text{MinCut}_{st}.$$

The equality among these values has profound implications on survivable network design for single-layer networks. Because all these survivability measures converge to the same value, it can naturally be used as the standard survivability metric that is applicable to measuring both disjoint paths or minimum cut. Another consequence of this equality is that linear programs (which are polynomial time solvable) can be used to find the minimum cut and disjoint paths in the network.

It is therefore interesting to see whether the same relationship holds for multi-layer networks. First, it is easy to verify that the linear relaxations for the formulations MaxFlow_{st} and MinCut_{st} maintain a primal-dual relationship, which, by Duality Theorem [19], implies that $\text{MaxFlow}_{st}^R = \text{MinCut}_{st}^R$. In addition, since any feasible solution to an integer program is also a feasible solution to the linear relaxation, we can establish the following relationship:

$$\text{Observation 1: } \text{MaxFlow}_{st} \leq \text{MaxFlow}_{st}^R = \text{MinCut}_{st}^R \leq \text{MinCut}_{st}.$$

Therefore, like single-layer networks, the maximum number of disjoint paths between two nodes cannot exceed the minimum cut between them in a multi-layer network.

However, unlike the single-layer case, the values of MaxFlow_{st} , MaxFlow_{st}^R and MinCut_{st} are not always identical, as illustrated in the following example. In our examples throughout the section, we use a logical topology with two nodes s and t that are connected by multiple lightpaths. For simplicity of exposition, we omit the complete lightpath routing and only show the physical links that are shared by multiple lightpaths. Theorem 1 states that this simplification can be made without loss of generality.

Theorem 1: Let G_L be a logical topology with two nodes s and t , connected by n lightpaths $E_L = \{e_1, e_2, \dots, e_n\}$, and let $\mathcal{R} = \{R_1, R_2, \dots, R_k\}$ be a family of subsets of E_L , where each $|R_i| \geq 2$, that captures the fiber-sharing relationship of the logical links. There exist a physical topology $G_P = (V_P, E_P)$ and lightpath routing of G_L over G_P , such that:

- 1) there are exactly k fibers in E_P , denoted by $F = \{f_1, f_2, \dots, f_k\}$, that are used by multiple lightpaths;
- 2) for each fiber $f_i \in F$, the set of lightpaths using f_i is R_i .

Proof: See Appendix A. ■

Theorem 1 implies that for a two-node logical topology, any arbitrary fiber-sharing relationship \mathcal{R} can be realized by reconstructing a physical topology and lightpath routing. Therefore, in the following discussion, we can simplify our examples by only giving the fiber-sharing relationship of our two-node logical topology without showing the details of the lightpath routing.

In Figure 3, the two nodes in the logical topology are connected by three lightpaths. The logical topology is embedded on the physical topology in such a way that each pair of lightpaths share a fiber. It is easy to see that no single fiber can disconnect the logical topology, and that any pair of fibers would. Hence, the value of MinCut_{st} is 2 in this case. On the other hand, the value of MaxFlow_{st} is only 1, because any two logical links share some physical fiber, so none of the paths in the logical network are physically disjoint. Finally, the value of MaxFlow_{st}^R is 1.5 because a flow of 0.5 can be routed on each of the lightpaths without violating the capacity constraints at the physical layer. Therefore, all three quantities are different in this example. We will study the integrality gaps for the formulations more carefully.

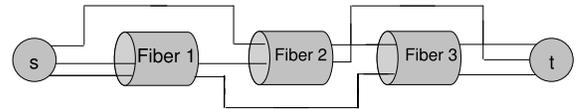


Fig. 3. A logical topology with 3 links where each pair of links shares a fiber in the physical topology

1) *Integrality Gap for MaxFlow_{st} :* The above example can be generalized to show that the ratio between MaxFlow_{st} and MaxFlow_{st}^R is $O(n)$, where n is the number of paths between s and t . Consider an instance of lightpath routing where the two nodes in the logical network are connected by n logical links, and every pair of logical links share a separate fiber. In this case, the value of MaxFlow_{st} will be 1, and the value of MaxFlow_{st}^R will be $\frac{n}{2}$, using the same arguments as above. Therefore, the ratio $\frac{\text{MaxFlow}_{st}^R}{\text{MaxFlow}_{st}}$ is $O(n)$. Note that this is an asymptotically tight bound since $\text{MaxFlow}_{st} \geq 1$ and $\text{MaxFlow}_{st}^R \leq n$ for all lightpath routings.

2) *Integrality Gap for MinCut_{st} :* The ratio between MinCut_{st} and MinCut_{st}^R can be shown to be at most $O(\log n)$ as a direct application of the result by Lovász [20], who showed that the integrality gap between integral and fractional set cover is $O(\log n)$. We can construct a lightpath routing

where the gap between the two values is $O(\log n)$, thereby showing the tightness of the bound.

Consider a layered network consisting of a two-node logical topology, and a set of k fibers $F = \{f_1, \dots, f_k\}$ that are shared by multiple logical links. For every subset T of $\lfloor \frac{k}{2} \rfloor + 1$ fibers in F , we add a logical link between the two logical nodes that uses only the fibers in T . Hence, for every set of $\lfloor \frac{k}{2} \rfloor - 1$ fibers, there is a logical link that does not use any of the fibers. This implies the Min Cut is at least $\lfloor \frac{k}{2} \rfloor$.

On the other hand, since each logical link uses exactly $\lfloor \frac{k}{2} \rfloor + 1$ fibers, the assignment where each $y_{ij} = \frac{1}{\lfloor \frac{k}{2} \rfloor + 1}$ satisfies Constraint (2), and is therefore a feasible solution to MinCut_{st}^R . The objective value of this solution is $\frac{k}{\lfloor \frac{k}{2} \rfloor + 1}$, which is at most

2. Therefore, the integrality gap $\frac{\text{MinCut}_{st}^R}{\text{MinCut}_{st}^I}$ is at least $\frac{k}{4}$.

Therefore, for the two-node logical network with $n = \binom{k}{\lfloor \frac{k}{2} \rfloor + 1}$ logical links, the ratio between the integral and relaxed optimal values for the Min Cut is $O(k) = O(\log n)$. We summarize our observation as follows:

Observation 2: In a layered network, the values of MaxFlow_{st} , MaxFlow_{st}^R and MinCut_{st} can be all different. In addition, the gaps among the three values are not bounded by any constant.

Therefore, a multi-layer network with high connectivity value (i.e. that tolerates a large number of failures) does not guarantee existence of physically disjoint paths. This is in sharp contrast to single-layer networks where the number of disjoint paths is always equal to the minimum cut.

It is thus clear that network survivability metrics across layers are not trivial extensions of the single layer metrics. New metrics need to be carefully defined in order to measure cross-layer survivability in a meaningful manner. In Section III, we will specify the requirements for cross-layer survivability metrics, and propose two new metrics that can be used to measure the connectivity of multi-layer networks.

B. Minimum Survivable Path Set

In this section, we introduce another graph structure, called *Survivable Path Set*, that is useful in describing connectivity in layered networks. A survivable path set for two logical nodes s and t is the smallest set of $s - t$ logical paths such that at least one of the paths in the set survives for any single physical link failure. The Minimum Survivable Path Set, denoted as MinSPS_{st} , is the size of the smallest survivable path set. For convenience, MinSPS_{st} is defined to be ∞ if no survivable path set exists.

In a single layer network, the value of MinSPS_{st} reveals nothing more than the existence of disjoint paths, as its value is either two or ∞ , depending on whether disjoint paths between s and t exist. However, for multi-layer networks, MinSPS_{st} can take on other values. For example, in Figure 3, the minimum survivable path set for s and t has size three because any pair of logical links can be disconnected by a single fiber failure. In fact, it is easy to verify that:

- $\text{MinSPS}_{st} = 2$ if and only if $\text{MaxFlow}_{st} \geq 2$;
- $\text{MinSPS}_{st} = \infty$ if and only if $\text{MinCut}_{st} = 1$.

Therefore, the value of MinSPS_{st} provides a different perspective about the connectivity between two nodes in the cross-layer setting. It is particularly interesting in the regime where $\text{MaxFlow}_{st} = 1$ and $\text{MinCut}_{st} \geq 2$, i.e., there is a gap between the Max Flow and the Min Cut. The following theorem reveals a connection between survivable path sets and the relaxed maxflow MaxFlow_{st}^R .

Theorem 2: $\text{MinSPS}_{st} \leq \left\lfloor \frac{\log |E_P|}{\log \text{MaxFlow}_{st}^R} \right\rfloor + 1$.

Proof: See Appendix B ■

It is worth noting that the theorem provides a sufficient condition for the existence of disjoint paths in the layered networks, in terms of the optimal value of MaxFlow_{st}^R :

Corollary 3: Disjoint paths between two nodes s and t exist in a layered network if the relaxed Max Flow, MaxFlow_{st}^R , is greater than $\sqrt{|E_P|}$.

Proof: By Theorem 2, a survivable path set of size two exists if $\text{MaxFlow}_{st}^R > \sqrt{|E_P|}$. This implies the existence of $s - t$ disjoint paths in the layered network. ■

Therefore, survivable path sets not only are interesting graph structures that describe connectivity of layered networks, they can also be useful in revealing the relationship between integral and fractional flows in the layered network.

C. Computational Complexity

For single-layer networks, because the integral Max Flow and Min Cut values are always identical to the optimal relaxed solutions, these values can be computed in polynomial time [18]. However, computing and approximating their cross-layer equivalents turns out to be much more difficult. Theorem 4 describes the complexity of computing the Max Flow and Min Cut for multi-layer networks.

Theorem 4: Computing Max Flow and Min Cut for multi-layer networks is NP-hard. In addition, both values cannot be approximated within any constant factor, unless $P=NP$.

Proof: The Max Flow can be reduced from the NP-hard Maximum Set Packing problem [21]:

Maximum Set Packing: Given a set of elements $E = \{e_1, e_2, \dots, e_n\}$ and a family $\mathcal{F} = \{C_1, C_2, \dots, C_m\}$ of subsets of E , find the maximum value k such that there exist k subsets $\{C_{j_1}, C_{j_2}, \dots, C_{j_k}\} \subseteq \mathcal{F}$ that are mutually disjoint.

Given an instance of Maximum Set Packing, we construct a 2-node logical topology connected by multiple lightpaths as described in Theorem 1, so that the optimal value of the Maximum Set Packing instance equals the maximum number of physically disjoint paths in the 2-node logical topology. This means that Maximum Set Packing is polynomial time reducible to the 2-node disjoint path problem. Theorem 1 implies that any instance of the 2-node disjoint path problem is polynomial time reducible to an instance of the multi-layer Max Flow problem. It follows that Maximum Set Packing is polynomial time reducible to the multi-layer Max Flow problem, which proves the theorem.

Given an instance of Maximum Set Packing with ground set E and a family \mathcal{F} of subsets of E , we construct a logical topology with two nodes, s and t , connected by $|\mathcal{F}|$ logical

links, where each logical link corresponds to a subset in \mathcal{F} . The logical links are embedded on the physical network in a way that two logical links share a physical fiber if and only if their corresponding subsets share a common element in the Maximum Set Packing instance. It immediately follows that a set of physically disjoint $s-t$ paths in the logical topology corresponds to a family of mutually disjoint subsets of E .

Similarly, the Min Cut can be reduced from the NP-hard Minimum Set Cover problem [21]:

Minimum Set Cover: Given a set $E = \{e_1, e_2, \dots, e_n\}$ and a family $\mathcal{F} = \{C_1, C_2, \dots, C_m\}$ of subsets of E , find the minimum value k such that there exist k subsets $\{C_{j_1}, C_{j_2}, \dots, C_{j_k}\} \subseteq \mathcal{F}$ that cover E , i.e., $\bigcup_{C_{j_i} \in \mathcal{C}} = E$.

Given an instance of Minimum Set Cover with ground set E and family of subsets \mathcal{F} , we construct a logical topology that contains two nodes connected by a set of $|E|$ logical links, where each logical link l_i corresponds to the element e_i . The logical links are embedded on the physical network in a way that exactly $|\mathcal{F}|$ fibers, namely $\{f_1, \dots, f_{|\mathcal{F}|}\}$, are used by multiple logical links, and the logical link l_i uses physical fiber f_j if and only if $e_i \in C_j$. It follows that the minimum number of physical fibers that forms a cut between the two logical nodes equals the size of a minimum set cover.

The inapproximability result follows immediately from the inapproximabilities of the Maximum Set Packing and Minimum Set Cover problems [22]–[24]. ■

In summary, multi-layer connectivity exhibits fundamentally different structural properties from its single-layer counterpart. Because of that, it is important to reinvestigate issues of quantifying, measuring as well as optimizing survivability in multi-layer networks. In the rest of the paper, we will focus on designing appropriate metrics for layered networks, and developing algorithms to maximize the cross-layer survivability.

III. METRICS FOR CROSS-LAYER SURVIVABILITY

The previous section demonstrates the new challenges in designing survivable layered network architectures. Insights into quantifying and optimizing survivability are fundamentally different between the single-layer and multi-layer settings. In this section, we focus on the issue of quantifying survivability in multi-layer networks. Not only should such metrics have natural physical meaning in the cross-layer setting, they should also be mathematically consistent and compatible with the conventional single-layer connectivity metric. Hence, we first define formal requirements for metrics that can be used to quantify cross-layer survivability:

- **Consistency:** A network with a higher metric value should be more resilient to failures.
- **Monotonicity:** Any addition of physical or logical links to the network should not decrease the metric value.
- **Compatibility:** The metric should generalize the connectivity metric for single-layer networks. In particular, when applied to the degenerated case where the physical and logical topologies are identical, the metric should be equivalent to the connectivity of the topology.

A metric that carries all the above properties would give us a meaningful and consistent measure of survivability for layered networks. We propose two metrics, the *Min Cross Layer Cut* and the *Weighted Load Factor*, that can be used as cross-layer survivability metrics. It is easy to verify that both metrics satisfy the above requirements.

A. Min Cross Layer Cut

In Section II, we defined MinCut_{st} to be the minimum number of physical failures that would disconnect logical nodes s and t . One can easily generalize this by taking the minimum over all possible node pairs to obtain a global connectivity metric. We define the *Min Cross Layer Cut (MCLC)* to be the minimum number of physical failures that would disconnect the logical topology.

A lightpath routing with high Min Cross Layer Cut value implies that the network remains connected even after a large number of physical failures. It is also a generalization of the survivable lightpath routing definition in [7], since a lightpath routing is survivable if and only if its Min Cross Layer Cut is greater than 1.

Let S be a subset of the logical nodes V_L , and $\delta(S)$ be the set of the logical links with exactly one end point in S . Let H_S be the minimum number of physical links failures required to disconnect all links in $\delta(S)$. The Min Cross Layer Cut can be defined as follows:

$$MCLC = \min_{S \subseteq V_L} H_S.$$

For each S , computing H_S can be considered as finding the Min Cut between the two partitions S and $V_L - S$. In the proof of Theorem 4, we have shown that computing the value of MinCut_{st} is NP-Hard even if the logical topology contains just two nodes. This immediately implies that computing the global MCLC value is NP-Hard:

Theorem 5: Computing the MCLC is NP-Hard.

In practice, however, the MCLC is bounded by the node degree of the logical topology, which is usually a small constant d . In that case, the MCLC can be computed in polynomial time by enumerating all fiber sets with up to d fibers. To compute the MCLC in a general setting, it can be modelled by the following integer linear program.

Given the physical and logical topologies (V_P, E_P) , and (V_L, E_L) , let f_{ij}^{st} be binary constants that represent the lightpath routing, such that logical link (s, t) uses physical fiber (i, j) if and only if $f_{ij}^{st} = 1$. The MCLC can be formulated as the integer program below:

$$\begin{aligned} M_{MCLC} : \text{Minimize} \quad & \sum_{(i,j) \in E_P} y_{ij}, \quad \text{subject to:} \\ d_t - d_s \leq \quad & \sum_{(i,j) \in E_P} y_{ij} f_{ij}^{st} \quad \forall (s, t) \in E_L \end{aligned} \quad (3)$$

$$\begin{aligned} \sum_{n \in V_L} d_n \geq 1, d_0 = 0 \quad & (4) \\ d_n, y_{ij} \in \{0, 1\} \quad & \forall n \in V_L, (i, j) \in E_P \end{aligned}$$

The integer program contains a variable y_{ij} for each physical link (i, j) , and a variable d_k for each logical node

k . Constraint (3) maintains the following property for any feasible solution: if $d_k = 1$, the node k will be disconnected from node 0 after all physical links (i, j) with $y_{ij} = 1$ are removed. To see this, note that since $d_k = 1$ and $d_0 = 0$, any logical path from node 0 to node k contains a logical link (s, t) where $d_s = 0$ and $d_t = 1$. Constraint (3) requires that such a logical link traverse at least one of the fibers (i, j) with $y_{ij} = 1$. As a result, all paths from node 0 to node k must traverse one of these fibers, and node k will be disconnected from node 0 if these fibers are removed from the network. Constraint (4) requires node 0 to be disconnected from at least one node, which ensures that the set of fibers (i, j) with $y_{ij} = 1$ forms a global Cross Layer Cut.

In Section IV, we will use MCLC as the objective for the survivable lightpath routing problem, and develop algorithms to maximize such an objective.

B. Weighted Load Factor

Another way to measure the connectivity of a layered network is by quantifying the ‘‘impact’’ of each physical failure. The *Weighted Load Factor (WLF)*, an extension of the metric *Load Factor* introduced in [25], provides such a measure of survivability.

Given the physical topology (V_P, E_P) and logical topology (V_L, E_L) , let f_{ij}^{st} be binary constants that represent the lightpath routing, such that logical link (s, t) uses physical fiber (i, j) if and only if $f_{ij}^{st} = 1$. We formulate the WLF as follows:

$$\begin{aligned} \text{M}_{\text{WLF}} : \text{Maximize } & \frac{1}{z}, \quad \text{subject to:} \\ z \cdot \sum_{(s,t) \in \delta(S)} w_{st} & \geq \sum_{(s,t) \in \delta(S)} w_{st} f_{ij}^{st} \\ & \forall S \subset V_L, (i, j) \in E_P \\ \sum_{(s,t) \in \delta(S)} w_{st} & > 0 \quad \forall S \subset V_L \\ 0 \leq z, w_{st} & \leq 1 \quad \forall (s, t) \in E_L, \end{aligned}$$

where $\delta(S)$ is the cut set of S , i.e., the set of logical links that have exactly one end point in S .

The variables w_{st} are the weights assigned to the lightpaths. Over all possible logical cuts, the variable z measures the maximum fraction of weight inside a cut carried by a fiber. Intuitively, if we interpret the weight to be the amount of traffic in the lightpath, the value z can be interpreted as the maximum fraction of traffic across a set of nodes disrupted by a single fiber cut. The Weighted Load Factor formulation, defined to maximize the reciprocal of this fraction, thus tries to compute the logical edge weights that minimize the maximum fraction. This effectively measures the best way of spreading the weight across the fibers for the given lightpath routing. A lightpath routing with a larger Weighted Load Factor value means that it is more capable of spreading its weight within any cut across the fibers.

The Weighted Load Factor also generalizes the survivable lightpath routing defined in [7], since its value will be greater than 1 if and only if the lightpath routing is survivable.

Although the formulation M_{WLF} contains the quadratic terms zw_{st} , the optimal value of z can be obtained by iteratively solving the linear program with different fixed values of z . Using binary search over the range of z , we can find the minimum z where a feasible solution exists.

Unfortunately, the formulation M_{WLF} contains an exponential number of constraints, and may not be polynomial time solvable. In fact, Theorem 6 states that finding objective value for M_{WLF} is NP-Hard, even if the weights of the logical links w_{st} are given.

Theorem 6: Computing the Weighted Load Factor for a lightpath routing is NP-Hard even if the weight assignment w_{st} for the logical links is fixed.

Proof: The NP-Hardness proof is based on the reduction from the NP-Hard *Uniform Sparsest Cut* [26] problem. For details, see Appendix C. ■

Finally, Theorem 7 describes the relationship between the WLF and the MCLC. Given a lightpath routing, let M_{MCLC} be the ILP formation for its Min Cross Layer Cut, and let MCLC and MCLC^R be the optimal values for M_{MCLC} and its linear relaxation respectively. In addition, let WLF be the Weighted Load Factor of the lightpath routing. Then we have the following relationship:

$$\text{Theorem 7: } \text{MCLC}^R \leq \text{WLF} \leq \text{MCLC}.$$

Proof: See Appendix D. ■

Therefore, although the two metrics appear to measure different aspects of network connectivity, they are inherently related. As we will show in Section V, the two values are often identical.

IV. LIGHTPATH ROUTING FOR MCLC MAXIMIZATION

In this section, we consider the survivable lightpath routing problem using MCLC as the objective. At an abstract level, the optimal lightpath routing can be expressed as the following optimization problem:

$$\max_{f \in \mathcal{F}} \min_{S \subset V_L} \text{MFC}(f, S),$$

where \mathcal{F} is set of all possible lightpath routings, V_L is the logical node set, and $\text{MFC}(f, S)$ is the minimum number of fibers whose removal will disconnect all logical links in the cut set $\delta(S)$ given the lightpath routing f . This is a Max-Min problem that may not have a simple formulation. In [27], we present an ILP formulation for this optimization problem. However, the formulation contains an exponential number of variables and constraints, which makes it infeasible even for small networks.

Therefore, in this section, we consider ILP formulations whose objective values are *lower bounds* to the MCLC. These formulations are much simpler than the exact formulation presented in [27]. This makes it possible to develop survivable lightpath routing algorithms based on these simpler formulations. In particular, in Section IV-C we discuss how to use randomized rounding [28] based on these formulations as a heuristic to approximate MCLC maximization. Note that since MCLC is $O(\log n)$ inapproximable, polynomial time algorithms with approximation guarantees within this factor

are unlikely to exist. Therefore, we will instead evaluate the performance of our algorithms via simulation in Section V.

All the formulations introduced in this section are based on multi-commodity flows, where each lightpath is considered a commodity to be routed over the physical network. Given the physical network $G_P = (V_P, E_P)$ and the logical network $G_L = (V_L, E_L)$, the multi-commodity flow for a lightpath routing can be generally formulated as follows:

MCF $_{\mathcal{X}}$: Minimize $\mathcal{X}(f)$, subject to:

$$\begin{aligned} f_{ij}^{st} &\in \{0, 1\} \\ \{f_{ij}^{st} : (i, j) \in E_P\} &\text{ forms an } (s, t)\text{-path, } \forall (s, t) \in E_L, \end{aligned} \quad (5)$$

where f is the variable set that represents the lightpath routing, such that $f_{ij}^{st} = 1$ if and only if lightpath (s, t) uses physical fiber (i, j) in its route. $\mathcal{X}(f)$ is an objective function that depends on f .

For WDM networks where the wavelength continuity constraint is present, the above formulation can be extended to capture the wavelength assignment aspect. In that case, the wavelength assignment can be modelled by replacing the variable set f_{ij}^{st} by $f_{ij\lambda}^{st}$, which equals 1 if and only if lightpath (s, t) uses wavelength λ on physical link (i, j) . Constraint (5) can be easily extended to restrict that, for each logical link (s, t) , $\{f_{ij\lambda}^{st} = 1\}$ forms an (s, t) physical path along one of the wavelengths. To make sure that any wavelength λ on a physical fiber is used by at most one lightpath, the following constraint will be added:

$$\sum_{(s,t) \in E_L} f_{ij\lambda}^{st} \leq 1 \quad \forall (i, j) \in E_P, \forall \lambda. \quad (6)$$

Similar formulations based on multi-commodity flows with wavelength continuity constraint have been proposed to solve the RWA problem of WDM networks [29], [30], where the objective is to minimize the number of lightpaths that traverse the same fiber. The key difference in the problem studied in this paper is in the objective function \mathcal{X} , which should instead describe the survivability of the lightpath routing. To focus on the survivability aspect of the problem, the wavelength continuity constraint will be omitted in the formulations below. However, in cases where the wavelength continuity constraint is necessary, all these formulations can be extended as discussed above.

A. Simple Multi-Commodity Flow Formulations

Ideally, to ensure that the lightpath routing is survivable against the largest number of failures, the objective function $\mathcal{X}(f)$ should express the MCLC value of the lightpath routing given by f . However, since there is no simple way to express the lightpath routing problem that maximizes the MCLC as an integer linear program, we use an objective that approximates the MCLC value. In our formulation, each lightpath is assigned a weight w . The objective function ρ_w measures the maximum load of the fibers, where the load is defined to be the total lightpath weight carried by the fiber. The intuition is that the multi-commodity flow formulation will try to spread the weight of the lightpaths across multiple fibers, thereby

minimizing the impact of any single fiber failure. We can formulate an ILP with such an objective as follows:

MCF $_w$: Minimize ρ_w , subject to:

$$\begin{aligned} \rho_w &\geq \sum_{(s,t) \in E_L} w(s, t) f_{ij}^{st} \quad \forall (i, j) \in E_P \\ f_{ij}^{st} &\in \{0, 1\} \\ \{f_{ij}^{st} : (i, j) \in E_P\} &\text{ forms an } (s, t)\text{-path, } \forall (s, t) \in E_L \end{aligned}$$

As we will prove in Theorem 8, with a careful choice of the weight function w , the value $\frac{1}{\rho_w}$ gives a lower bound on the MCLC. Therefore, a lightpath routing with a low ρ_w value is guaranteed to have a high MCLC.

The routing strategy of the algorithm is determined by the weight function w . For example, if w is set to 1 for all lightpaths, the integer program will minimize the number of lightpaths traversing the same fiber. Effectively, this will minimize the number of disconnected lightpaths in the case of a single fiber failure.

In order to customize MCF $_w$ towards maximizing the MCLC of the solution, we propose a different weight function w_{MinCut} that captures the connectivity structure of the logical topology. For each edge $(s, t) \in E_L$, we define $w_{MinCut}(s, t)$ to be $\frac{1}{|MinCut_L(s, t)|}$, where $MinCut_L(s, t)$ is the minimum (s, t) -cut in the logical topology. Therefore, if an edge (s, t) belongs to a smaller cut, it will be assigned a higher weight. The algorithm will therefore try to avoid putting these small cut edges on the same fiber.

If w_{MinCut} is used as the weight function used in MCF $_w$, we can prove the following relationship between the objective value ρ_w of a feasible solution to MCF $_w$ and the Weighted Load Factor of the associated lightpath routing:

Theorem 8: For any feasible solution f of MCF $_w$ with w_{MinCut} as the weight function, $\frac{1}{\rho_w} \leq WLF$.

Proof: By definition of the weight function w_{MinCut} , given any $S \subset V_L$, every edge in $\delta(S)$ has weight at least $\frac{1}{|\delta(S)|}$. Therefore, we have:

$$\sum_{(s,t) \in \delta(S)} w(s, t) \geq \sum_{(s,t) \in \delta(S)} \frac{1}{|\delta(S)|} = 1. \quad (7)$$

Now consider the lightpath routing associated with f . For any logical cut $\delta(S)$, the maximum fraction of weight inside the cut carried by a fiber is:

$$\begin{aligned} &\max_{(i,j) \in E_P} \frac{\sum_{(s,t) \in \delta(S)} w(s, t) f_{ij}^{st}}{\sum_{(s,t) \in \delta(S)} w(s, t)} \\ &\leq \max_{(i,j) \in E_P} \sum_{(s,t) \in \delta(S)} w(s, t) f_{ij}^{st}, \quad \text{by Equation (7)} \\ &\leq \max_{(i,j) \in E_P} \sum_{(s,t) \in E_L} w(s, t) f_{ij}^{st} \leq \rho_w. \end{aligned}$$

In other words, no fiber in the network is carrying more than a fraction ρ_w of the weight in any cut. This gives us a feasible solution to the Weighted Load Factor formulation M_{WLF} , where each variable w_{st} is assigned the value of

$w_{MinCut}(s, t)$, and the variable z is assigned the value of ρ_w . As a result, the Weighted Load Factor, defined to be the maximum value of $\frac{1}{z}$ among all feasible solutions to M_{WLF} , must be at least $\frac{1}{\rho_w}$. ■

As a result of Theorems 7 and 8, the MCLC of a lightpath routing is lower bounded by the value of $\frac{1}{\rho_w}$, which the algorithm will try to maximize.

B. Enhanced Multi-Commodity Flow Formulation

As we have discussed in Section III-B, the Weighted Load Factor provides a good lower bound on the MCLC of a lightpath routing. Here we propose another multi-commodity flow based formulation whose objective function approximates the Weighted Load Factor of a lightpath routing. The formulation, denoted as MCF_{LF} , can be written as follows:

$$\begin{aligned} MCF_{LF}: \quad & \text{Minimize } \gamma, \quad \text{subject to:} \\ & \gamma |\delta(S)| \geq \sum_{(s,t) \in \delta(S)} f_{ij}^{st} \quad \forall (i, j) \in E_P, S \subset V_L \\ & f_{ij}^{st} \in \{0, 1\} \\ & \{f_{ij}^{st} : (i, j) \in E_P\} \text{ forms an } (s, t)\text{-path, } \forall (s, t) \in E_L \end{aligned}$$

Essentially, the formulation optimizes the *unweighed* Load Factor of the lightpath routing, (i.e., all weights equal one), by minimizing the maximum fraction of a logical cut carried by a single fiber. As this formulation provides a constraint for each logical cut, it captures the impact of a single fiber cut on the logical topology in much greater detail. The following theorem shows that for any lightpath routing, its associated Load Factor value $\frac{1}{\gamma}$ gives a tighter lower bound than $\frac{1}{\rho_w}$, given by the MCF_w formulation.

Theorem 9: For any lightpath routing, let ρ_w be its associated objective value in the formulation MCF_w with w_{MinCut} as the weight function, and let γ be its associated objective value in the formulation MCF_{LF} . In addition, let WLF be its Weighted Load Factor. Then:

$$\frac{1}{\rho_w} \leq \frac{1}{\gamma} \leq WLF.$$

Proof: The value $\frac{1}{\gamma}$ is the objective value for the formulation M_{WLF} in Section III-B when all logical links have weight 1. This gives a feasible solution to M_{WLF} , and implies that $WLF \geq \frac{1}{\gamma}$.

To prove that $\frac{1}{\rho_w} \leq \frac{1}{\gamma}$, we consider the physical link (i, j) and logical cut set $\delta(S)$ where (i, j) carries a fraction γ of the logical links in $\delta(S)$. Let L_{ij} be the set of logical links in E_L carried by (i, j) . Therefore, we have $\gamma = \frac{|L_{ij} \cap \delta(S)|}{|\delta(S)|}$. In addition, by the definition of ρ_w , we have

$$\rho_w \geq \sum_{(s,t) \in L_{ij}} w(s, t) \geq \sum_{(s,t) \in L_{ij} \cap \delta(S)} \frac{1}{|\delta(S)|} = \gamma.$$

This implies $\frac{1}{\rho_w} \leq \frac{1}{\gamma}$. ■

Therefore, the formulation MCF_{LF} gives a lightpath routing that is optimized for a better lower bound on the MCLC. However, this comes at the cost of a larger number of constraints and solving such an integer program may not be feasible in

practice. Therefore, we next introduce a randomized rounding technique that approximates the optimal lightpath routing by solving the linear relaxation of the integer program. As we will see in Section V, the randomized rounding technique significantly speeds up the running time of the algorithm without observable degradation in the MCLC performance. This offers a practical alternative to solving the integer program formulations introduced in this section.

C. Randomized Rounding for Lightpath Routing

While the multi-commodity flow integer program formulations discussed in the previous section introduce a novel way to route lightpaths in a survivable manner, such an approach may not scale to large networks, due to the inherent complexity of solving integer programs. In order to circumvent the computational difficulty, we apply the *randomized rounding* technique, which is able to quickly obtain a near-optimal solution to the integer program. Randomized rounding has previously been used to solve multi-commodity flow problems to minimize the link load [28], [29], and its performance guarantee is studied in [28].

Through randomized rounding, we solve the linear relaxation of the ILP, and choose a random physical path between s and t for each lightpath (s, t) with probability based on the optimal fraction solution obtained from the linear relaxation. We consider a variant, called $RANDOM_k$, where the process of choosing random lightpath routing is repeated for k times. The lightpath routing with the highest MCLC value is selected as the final output. The repetition increases the probability of obtaining a good solution. As we will see in the next section, randomized rounding provides a much more efficient way of obtaining a survivable lightpath routing without sacrificing the quality of solution.

V. SIMULATION

In this section, we discuss our simulation results for the algorithms introduced in Section IV. We first compare the lightpath routing algorithms by solving the ILP directly and by randomized rounding. Next, we compare the survivability performance among different formulations. Finally, we investigate the different lower bounds of MCLC, and their effects on the MCLC value of the lightpath routing when used as an optimization objective.

1) *ILP vs Randomized Rounding:* In this experiment, we use the NSFNET (Figure 4) as the physical topology. The network is augmented to have connectivity 4, which makes it possible to study the performance of the algorithms where a higher MCLC value is possible. We generated 350 random logical topologies with connectivity at least 4, and size ranging from 6 to 12 nodes. Using the formulation MCF_w with weight function $w_{MinCut}(s, t)$ introduced in Section IV-A as our benchmark, we compare the performance of $RANDOM_{10}$ against solving the ILP optimally.

Table I compares the average running time between the algorithms ILP and $RANDOM_{10}$ on various logical topology size. All simulations are run on a Xeon E5420 2.5GHz workstation with 4GB of memory, using CPLEX to solve

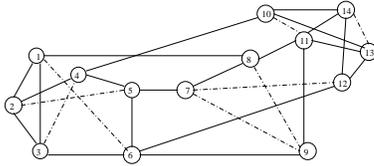


Fig. 4. The augmented NSFNET. The dashed lines are the new links.

the integer and linear programs. As the number of logical nodes increases, the running time for the integer program algorithm ILP increases tremendously. On the other hand, there is no observable growth in the average running time for the algorithm RANDOM_{10} , which is less than a minute. In fact, our simulation on larger networks shows that the algorithm ILP often fails to terminate within a day when the network size goes beyond 12 nodes. On the other hand, the algorithm RANDOM_{10} for MCF_w is able to terminate consistently within 2 hours for very large instances with a 100-node physical topology and 50-node logical topology. This shows that the randomized approach is a much more scalable solution to compute survivable lightpath routings.

Logical Topology Size	Average Running Time (seconds)	
	ILP	RANDOM_{10}
6	33.2	31.9
7	50.5	33.9
8	660.0	30.1
9	1539.0	26.4
10	3090.6	32.3
11	8474.5	32.0
12	15369.7	29.7

TABLE I. Average Running Time of ILP and RANDOM_{10}

In Figure 5, the survivability performance of the randomized algorithm is compared with its ILP counterpart. Each data point in the figure is the MCLC average of 50 random instances with the given logical network size. As our result shows, the lightpath routings produced by RANDOM_{10} have higher MCLC values than solving the ILP optimally. This is because the objective value for ILP MCF_w is a *lower bound* on MCLC. As we will see in Section V-3, this lower bound is often not tight enough to accurately reflect the MCLC value, and the optimal solution to the ILP does not necessarily yield a lightpath routing with maximum MCLC. On the other hand, the randomized algorithm generates lightpath routings non-deterministically based on the optimal fractional solution of MCF_w . Therefore, it approximates the lightpath routing given by the ILP, with an additional randomization component to explore better solutions. When the randomized rounding process is repeated many times, the algorithm often encounters a solution that is even better than the one given by the ILP.

To sum up, randomized rounding provides an efficient alternative to solving integer programs without observable quality degradation. This allows us to experiment with more complex formulations in larger networks where solving the integer programs optimally is infeasible. In the next section, we will compare the different formulations introduced in Section IV-A, using randomized rounding to compute the lightpath routings.

2) *Lightpath Routing with Different Formulations*: In this experiment, we study the survivability performance of the

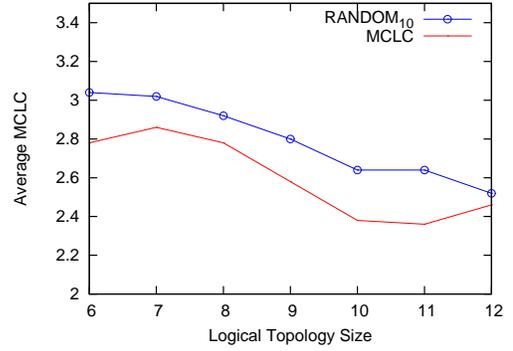


Fig. 5. MCLC Performance of Randomized Rounding vs ILP

lightpath routings generated by the formulations introduced in Section IV. We use the 24-node USIP network (Figure 6), augmented to have connectivity 4, as the physical topology. We generate 500 random graphs with connectivity 4 and size ranging from 6 to 15 nodes as logical topologies.

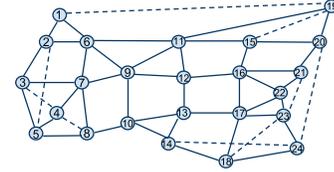


Fig. 6. The augmented USIP network. The dashed lines are the new links.

We compare the MCLC performance of the lightpath routings generated by the randomized rounding algorithm, RANDOM_{100} , on the following formulations:

- 1) Multi-Commodity Flow MCF_w , with the identity weight function, i.e., $w(s, t) = 1$ for all $(s, t) \in E_L$ (Identity);
- 2) Multi-Commodity Flow MCF_w , with the weight function w_{MinCut} introduced in Section IV-A (MinCut);
- 3) Enhanced Multi-Commodity Flow MCF_{LF} (LF).

For comparison, we also run randomized rounding on the Survivable Lightpath Routing formulation (SURVIVE), introduced in [7], which computes the lightpath routing that minimizes the total fiber hops, subject to the constraint that the MCLC must be at least two.

Figure 7 compares the average MCLC values of the lightpath routings computed by the four different algorithms. Overall, the formulations introduced in this paper achieve better survivability than SURVIVE. This is because these formulations try to maximize the MCLC in their objective functions, whereas SURVIVE minimizes the physical hops. Therefore, even though SURVIVE does well in finding a survivable routing (i.e. $\text{MCLC} \geq 2$), the new formulations are able to achieve even higher MCLC values, which allow more physical failures to be tolerated.

To further verify the survivability performance of the lightpath routings from a different perspective, for each lightpath routing, we simulated the scenario where each physical link fails independently with probability 0.01. Figure 8 shows the average probability that the logical topology becomes disconnected under this scenario. The

result is consistent with Figure 7, as lightpaths routings with higher MCLC values can tolerate more physical failures, and the logical topologies are thus more likely to stay connected.

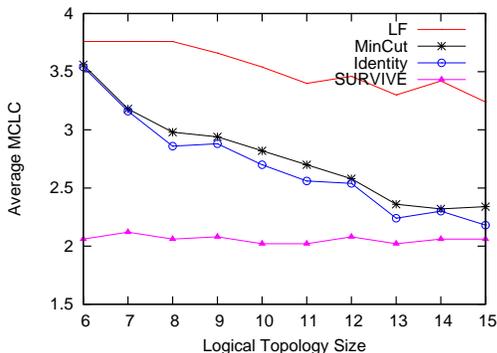


Fig. 7. MCLC Performance of Different Formulations

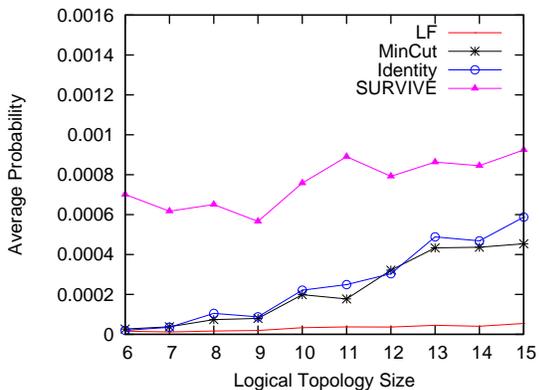


Fig. 8. Probability that Logical Topology Becomes Disconnected if Physical Links Fail Independently with Probability 0.01

The quality of the lightpath routing also depends on the graph structures captured by the formulations. Compared with $MCF_{Identity}$, the formulation MCF_{MinCut} uses a weight function that captures the connectivity structure of the logical topology. As a result, the algorithm will try to avoid putting edges that belong to smaller cuts onto the same physical link, thereby minimizing the impact of a physical link failure on these critical edges. This allows the algorithm MCF_{MinCut} to produce lightpath routings with higher MCLC values than $MCF_{Identity}$.

As the number of logical links increases with the size of the logical topology, the maximum fiber link load becomes less effective in capturing the impact of fiber failures for the logical cuts. As a result, the performance of the simple multi-commodity formulations degrades somewhat rapidly. On the other hand, the enhanced formulation MCF_{LF} captures the connectivity structure of the logical topology in much greater detail, by having a constraint to describe the impact of a physical link failure to each logic cut. Therefore, this formulation is able to provide lightpath routings with the highest MCLC values. This observation is supported by Theorem 9, which states that the objective maximized by MCF_{LF} is closer to the actual MCLC value.

3) *Lower Bound Comparison:* In Theorem 9 we establish different lower bounds for the MCLC. In this experiment, we

measure these lower bound values for 500 different lightpath routings, and compare them to the actual MCLC values.

As Figure 9 shows, the Weighted Load Factor is a very close approximation of the Min Cross Layer Cut. Among the 500 routings being investigated, the two metrics are identical in 368 cases. This suggests a tight connection between the two metrics, which also justifies the choice of such metrics as survivability measures.

The figure also reveals a strong correlation between the MCLC performance and the tightness of the lower bounds given by the multi-commodity flow formulations in Section IV-A. Compared to MCF_w , the formulation MCF_{LF} provides an objective value that is closer to the actual MCLC value of the lightpath routing. This translates to better lightpath routings, as we saw in Figure 7. Since there is still a large gap between the MCF_{LF} objective value and the MCLC value, this suggests room for further improvement with a formulation that gives a better MCLC lower bound.

To summarize this section, a good formulation that properly captures the cross-layer connectivity structure is essential for generating lightpath routings with high survivability. Combined with randomized rounding, it gives a powerful tool for designing highly survivable layered networks.

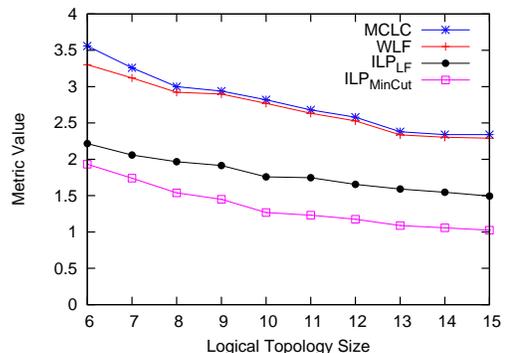


Fig. 9. Comparison among Min Cross Layer Cut (MCLC), Weighted Load Factor (WLF), and the Optimal Values of ILP_{LF} and ILP_{MinCut}

VI. CONCLUSION

In this paper, we introduce the problem of maximizing the connectivity of layered networks. We show that survivability metrics in multi-layer networks have significantly different meaning than their single-layer counterparts. We propose two survivability metrics, the Min Cross Layer Cut and the Weighted Load Factor, that measure the connectivity of a multi-layer network, and develop linear and integer formulations to compute these metrics. In addition, we use the metric Min Cross Layer Cut as the objective for the survivable lightpath routing problem, and develop multi-commodity flow formulations to approximate this objective. We show, through simulations, that our algorithms produce lightpath routings with significantly better Min Cross Layer Cut values than existing survivable lightpath routing algorithm.

Our simulation shows that a good formulation, combined with the randomized rounding technique, provides a powerful tool for generating highly survivable layered networks. Therefore, an important direction for future research is to establish

a better formulation for the lightpath routing problem that maximizes the Min Cross Layer Cut. The multi-commodity flow formulation introduced in this paper approximates the Min Cross Layer Cut by using its lower bound as the objective function. However, this lower bound is often not very close to the actual Min Cross Layer Cut value. A better objective function, such as the Weighted Load Factor, would significantly improve the proposed lightpath routing algorithms. We are currently exploring the possibilities in this direction.

The similarity between the Min Cross Layer Cut and the Weighted Load Factor is also intriguing. Our simulation result demonstrated a very tight connection between the two metrics. This observation might reflect certain property of cross-layer network connectivity that are yet to be discovered and formalized. A better understanding of how these metrics relate to each other will possibly lead to important insights into the cross-layer survivability problem.

APPENDIX A PROOF OF THEOREM 1

Theorem 1: Let G_L be a logical topology with two nodes s and t , connected by n lightpaths $E_L = \{e_1, e_2, \dots, e_n\}$, and let $\mathcal{R} = \{R_1, R_2, \dots, R_k\}$ be a family of subsets of E_L where each $|R_i| \geq 2$. There exists a physical topology $G_P = (V_P, E_P)$ and lightpath routing of G_L over G_P , such that:

- 1) there are exactly k fibers in E_P , denoted by $F = \{f_1, f_2, \dots, f_k\}$, that are used by multiple lightpaths;
- 2) for each fiber $f_i \in F$, the set of lightpaths using the fiber f_i is R_i .

Proof: Given a logical topology $G_L = (V_L, E_L)$ with two nodes s and t connected by n lightpaths $E_L = \{e_1, e_2, \dots, e_n\}$, and $\mathcal{R} = \{R_1, R_2, \dots, R_k\}$ be the family of subsets of E_L , we construct a physical topology and lightpath routing that satisfy the conditions specified in the theorem.

• Physical Topology:

The physical topology contains the two end nodes s and t in the logical network. In addition, between the two end nodes, there are n groups of nodes. Each group i containing $k + 1$ nodes, namely $x_0^i, x_1^i, \dots, x_k^i$. For any $i \in \{1, \dots, n\}, j \in \{1, \dots, k\}$, there is an edge connecting nodes x_{j-1}^i and x_j^i . In addition, s is connected to x_0^i and x_k^i is connected to t for all $i \in \{1, \dots, n\}$. In other words, in the physical network we have constructed so far, there are n edge disjoint paths connecting s and t , and each path has $k + 2$ edges.

Next, we add k pairs of nodes $\{(y_1, z_1), \dots, (y_k, z_k)\}$ to the physical network, where each node pair (y_j, z_j) is connected by an edge. Finally, we connect x_{j-1}^i to y_j and z_j to x_j^i for all $i \in \{1, \dots, n\}, j \in \{1, \dots, k\}$.

• Lightpath Routing:

We will define a route in the physical topology for each lightpath e_i . Each route l_i will contain $k + 2$ segments:

$$s \rightsquigarrow x_0^i \rightsquigarrow x_1^i \rightsquigarrow \dots \rightsquigarrow x_k^i \rightsquigarrow t.$$

Segments $s \rightsquigarrow x_0^i$ and $x_k^i \rightsquigarrow t$ will take the direct edges $s \rightarrow x_0^i$ and $x_k^i \rightarrow t$ respectively as their routes. The routes for other segments depend on whether e_i is in R_j :

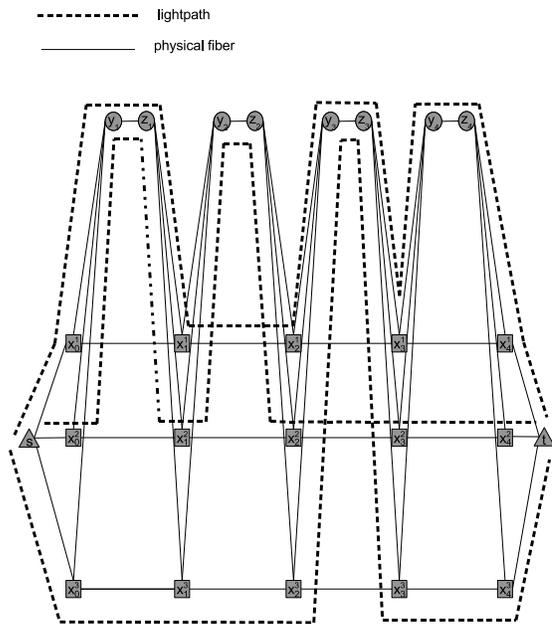


Fig. 10. The physical topology and lightpath routing on three lightpaths between two logical nodes s and t , and lightpath-sharing relationship $\mathcal{R} = \{\{1, 2\}, \{2\}, \{1, 3\}, \{1\}\}$.

- If $e_i \in R_j$, the route for $x_{j-1}^i \rightsquigarrow x_j^i$ is $x_{j-1}^i \rightarrow y_j \rightarrow z_j \rightarrow x_j^i$;
- If $e_i \notin R_j$, the route for $x_{j-1}^i \rightsquigarrow x_j^i$ is $x_{j-1}^i \rightarrow x_j^i$.

Figure 10 shows the physical topology and lightpath routing constructed from a two-node logical topology with $\mathcal{R} = \{\{1, 2\}, \{2\}, \{1, 3\}, \{1\}\}$.

By construction, all fibers except $\{(y_1, z_1), \dots, (y_k, z_k)\}$ are used by at most one lightpath. Also, a lightpath e_j uses fiber (y_i, z_i) if and only if e_j is in R_i . In other words, there are exactly k fibers, $(y_1, z_1), \dots, (y_k, z_k)$, that are used by multiple lightpaths, and each fiber (y_i, z_i) is used by the lightpaths in R_i . ■

APPENDIX B PROOF OF THEOREM 2

Let MinSPS_{st} be the size of the minimum survivable path set between the logical nodes s and t . Theorem 2 describes the relationship between the value of MinSPS_{st} and the relaxed Max Flow, MaxFlow_{st}^R , between the two nodes:

$$\text{Theorem 2: } \text{MinSPS}_{st} \leq \left\lfloor \frac{\log |E_P|}{\log \text{MaxFlow}_{st}^R} \right\rfloor + 1.$$

Proof: Let \mathcal{P}_{st} and E_P be the set of logical $s - t$ paths and the set of physical links respectively. For each $s - t$ path $p \in \mathcal{P}_{st}$, denote the set of physical links used by p as $L(p)$. We first construct a bipartite graph on the node set (\mathcal{P}_{st}, E_P) . There is an edge $(p, l) \in \mathcal{P}_{st} \times E_P$ if and only if the $s - t$ path p does *not* use physical link l , i.e., $l \notin L(p)$. In other words, the edge (p, l) is in the bipartite graph if and only if the path p survives the failure of physical link l .

We prove the theorem by explicitly constructing a survivable path set with size at most $\left\lfloor \frac{\log |E_P|}{\log \text{MaxFlow}_{st}^R} \right\rfloor + 1$, using the bipartite graph. Algorithm $\text{SPS}_{\text{GREEDY}}$ describes a greedy algorithm that constructs the path set by repeatedly selecting $s - t$ paths

and removing physical links whose failures the selected path can survive. When the algorithm terminates, every physical link failure is survived by a selected path in the output. Therefore, the algorithm gives a survivable path set.

Algorithm 1 SPS_{GREEDY}

```

1:  $P := \emptyset, S := E_P$ 
2: while  $S \neq \emptyset$ : do
  - Select  $p \in \mathcal{P}_{st}$  with the largest node degree in the
    bipartite graph.
  -  $P := P \cup \{p\}, S := S \setminus L(p)$ 
  - Remove nodes  $p$  and  $L(p)$  from the bipartite graph.
3: Return  $P$ 

```

The key observation for this algorithm is that, every iteration of the algorithm removes a constant fraction of remaining nodes in E_P . We state this result as the following lemma:

Lemma 10: Let B^i be the bipartite graph at the beginning of the i^{th} iteration of the algorithm, where the remaining node sets for E_P and \mathcal{P}_{st} are E_P^i and \mathcal{P}_{st}^i respectively. There exists a node in \mathcal{P}_{st}^i with node degree at least $\frac{|E_P^i|(\alpha-1)}{\alpha}$, where α is the optimal value for the formulation MaxFlow_{st}^R .

Proof: Suppose $\{f_p^* | p \in \mathcal{P}_{st}\}$ is the optimal solution for MaxFlow_{st}^R , such that $\sum_{p \in \mathcal{P}_{st}} f_p^* = \alpha$. For the purpose of analysis, for each edge $(p, l) \in \mathcal{P}_{st}^i \times E_P^i$ in the bipartite graph, we assign the edge a weight f_p^* .

For each node v in the bipartite graph, let $d(v)$ be its node degree, and we define its weight $w(v)$ to be sum of the weight of its incident edges. Then we have:

$$\sum_{p \in \mathcal{P}_{st}^i} \frac{w(p)}{d(p)} = \sum_{p \in \mathcal{P}_{st}^i} f_p^* \leq \alpha. \quad (8)$$

For each node l in E_P^i , its neighbors in \mathcal{P}_{st}^i are the same as its neighbors in \mathcal{P}_{st} , since otherwise it should have already been removed from the bipartite graph. Its node weight is:

$$\begin{aligned} w(l) &= \sum_{p \in \mathcal{P}_{st}^i: l \notin L(p)} f_p^* = \sum_{p \in \mathcal{P}_{st}: l \notin L(p)} f_p^* \\ &= \sum_{p \in \mathcal{P}_{st}} f_p^* - \sum_{p \in \mathcal{P}_{st}: l \in L(p)} f_p^* \\ &\geq \alpha - 1, \quad \text{since } \sum_{p: l \in L(p)} f_p^* \leq 1, \text{ by Equation (1).} \end{aligned}$$

Therefore the total weight for the nodes in E_P^i is at least $|E_P^i|(\alpha - 1)$, which implies:

$$\sum_{p \in \mathcal{P}_{st}^i} w(p) \geq |E_P^i|(\alpha - 1). \quad (9)$$

Let d_{max} be the largest node degree among the nodes in \mathcal{P}_{st}^i . We have:

$$d_{max} \geq \frac{\sum_{p \in \mathcal{P}_{st}^i} w(p)}{\sum_{p \in \mathcal{P}_{st}^i} \frac{w(p)}{d(p)}} \geq \frac{|E_P^i|(\alpha - 1)}{\alpha}, \text{ by (8) and (9).}$$

Therefore, the set \mathcal{P}_{st}^i contains a node with degree at least $\frac{|E_P^i|(\alpha-1)}{\alpha}$. ■

As a result of Lemma 10, every iteration of the algorithm removes a fraction of $\frac{\alpha-1}{\alpha}$ nodes of E_P^i from the bipartite graph.

After the i^{th} path is selected, the number of nodes in E_P that remain in the bipartite graph is at most $(1 - \frac{\alpha-1}{\alpha})^i |E_P|$. The algorithm will terminate as soon as $(1 - \frac{\alpha-1}{\alpha})^i |E_P| < 1$, which implies $i > \frac{\log |E_P|}{\log \alpha}$. Therefore, the algorithm returns a survivable path set with size $\lceil \log_{\alpha} |E_P| \rceil + 1$. ■

APPENDIX C
PROOF OF THEOREM 6

Theorem 6: Computing the Weighted Load Factor for a lightpath routing is NP-Hard even if the weight assignment w_{st} for the logical links is fixed.

Proof: We construct a reduction from the NP-Hard *Uniform Sparsest Cut* [26] problem:

• **Uniform Sparsest Cut:**

Given an undirected graph $G = (V, E)$, compute the value of $\min_{S \subset V_L} \frac{|\delta(S)|}{|S||V-S|}$.

Given the graph $G = (V, E)$ in an instance of Uniform Sparsest Cut problem, we construct an instance of the Weighted Load Factor problem, with the weight assignment w_{st} fixed, such that the optimal values of the two problems are identical. Without loss of generality, we assume G is connected. We will construct a physical topology, logical topology, lightpath routing f_{st}^{ij} and weight assignment w_{st} of the logical links based on the graph $G = (V, E)$ in the Uniform Sparsest Cut instance.

- **Logical Topology:** The logical topology is a complete graph on the vertex set $V_L = V$. Each logical link (s, t) has weight $w_{st} = 1$.
- **Physical Topology:** The physical topology is a complete graph on the vertex set $V_P = V \cup \{u, v\}$, where u and v are two new vertices not in V .
- **Lightpath Routing:** For each logical link (s, t) , if (s, t) is an edge of G in the Uniform Sparsest Cut instance, the logical link takes on the physical route $s \rightarrow u \rightarrow v \rightarrow t$. Otherwise, it takes on the physical route $s \rightarrow t$.

Let S be an arbitrary subset of V . Let $\delta_{SC}(S)$ be the cut set of S with respect to graph G of the Uniform Sparsest Cut instance, and let $\delta_L(S)$ be the cut set of S with respect to the logical topology G_L , which is a complete graph on $V_L = V$. We claim the following equality:

$$\frac{|\delta_{SC}(S)|}{|S||V-S|} = \max_{(i,j) \in E_P} \frac{\sum_{(s,t) \in \delta_L(S)} w_{st} f_{st}^{ij}}{\sum_{(s,t) \in \delta_L(S)} w_{st}}. \quad (10)$$

This is because every physical link not attached to u or v is used by at most one logical link. In addition, any logical link that uses a physical link in the form (x, u) or (v, x) , for any x in V_P , also uses (u, v) in the lightpath routing. Since G is connected, for each $S \subset V$, there is at least one logical link in $\delta_{SC}(S)$ that uses the physical link (u, v) . Therefore, for any $S \subset V_L$, the physical link (u, v) carries the largest number of logical links in $\delta_L(S)$. Since a logical link uses (u, v) if and only if the corresponding edge exists in G , the number of logical links in $\delta_{SC}(S)$ using (u, v) is $|\delta_{SC}(S)|$. Therefore, the fraction of weight carried by the physical link (u, v) is $\frac{|\delta_{SC}(S)|}{|\delta_L(S)|} = \frac{|\delta_{SC}(S)|}{|S||V-S|}$. This implies the sparsest cut value equals the Weighted Load Factor value. ■

APPENDIX D
PROOF OF THEOREM 7

Let $MCLC$ and $MCLC^R$ be the optimal objective values for formulation M_{MCLC} and its linear relaxation M_{MCLC}^R respectively. And let WLF be the Weighted Load Factor of the lightpath routing. Theorem 7 declares the following:

Theorem 7: $MCLC^R \leq WLF \leq MCLC$.

Proof: Recall that the ILP formulation for $MCLC$ is:

$$M_{MCLC} : \text{Minimize } \sum_{(i,j) \in E_P} y_{ij}, \quad \text{subject to:}$$

$$d_t - d_s \leq \sum_{(i,j) \in E_P} y_{ij} f_{ij}^{st}, \quad \forall (s,t) \in E_L \quad (11)$$

$$\sum_{n \in V_L} d_n \geq 1 \quad (12)$$

$$d_0 = 0, \quad d_n, y_{ij} \in \{0, 1\}, \quad \forall n \in V_L, (i,j) \in E_P$$

where f_{ij}^{st} are binary constants such that logical link (s,t) traverses physical fiber (i,j) if and only if $f_{ij}^{st} = 1$.

For the rest of the proof, for any $S \subset V_L$, we denote $\delta(S)$ to be the set of logical links with exactly one end point in S .

We first prove that $MCLC^R \leq WLF$. We construct the dual [19] of M_{MCLC}^R :

$$M_{MCLC}^{\text{Dual,R}} : \text{Maximize } q, \quad \text{subject to:}$$

$$\sum_{(s,t) \in E_L} g^{st} f_{ij}^{st} \leq 1, \quad \forall (i,j) \in E_P \quad (13)$$

$$q + \sum_{(s,t) \in E_L} g^{st} - \sum_{(t,s) \in E_L} g^{ts} \leq 0, \quad \forall s \neq 0 \quad (14)$$

$$q, g^{st} \geq 0, \quad \forall (s,t) \in E_L$$

The variables y_{ij} in the primal M_{MCLC}^R correspond to Constraint (13) in the dual. Similarly, the variables d_s , where $s \neq 0$, in the primal correspond to Constraint (14) in the dual. For Constraints (11) and (12) in the primal, the corresponding variables in the dual are g^{st} and q respectively. We can interpret the variable g^{st} as the flow value assigned to logical link (s,t) . Then Constraint (13) requires that the total flow on each physical fiber be at most 1. Constraint (14) requires at least q units of incoming flow for all nodes other than node 0. Intuitively, the dual program tries to maximize the value q such that the node 0 sends at least q units of flow to every other node, subject to the capacity constraint for each fiber.

We will use Lemma 11 to prove a lower bound on WLF :

Lemma 11: Let (q, g) be a feasible solution for $M_{MCLC}^{\text{Dual,R}}$, and let $g(S) = \sum_{(s,t) \in E_L: s \notin S, t \in S} g^{st} - \sum_{(s,t) \in E_L: s \in S, t \notin S} g^{st}$ be the net flow into the cut set S . Then $g(S) \geq kq$, for any $S \subseteq V_L \setminus \{0\}$ with $k = |S|$.

Proof: Consider an arbitrary node set $S \subseteq V_L \setminus \{0\}$, and let $k = |S|$. We prove by induction on k that $g(S) \geq kq$.

- **Base case:** $k = 0$: S is an empty set, so $g(S) = 0 = kq$.
- **Inductive case:** Suppose for some $0 \leq k < |V_L| - 1$, $g(S) \geq kq$ for all S with $|S| = k$ and $0 \notin S$. Now let S' be any subset of $k+1$ nodes that does not contain node

0, let b be an arbitrary node in S' , and let $S'_b = S' \setminus \{b\}$. Since S'_b is a set of k nodes, by induction hypothesis, we have $g(S'_b) \geq kq$. It follows that:

$$g(S') = g(S'_b) + \sum_{(t,b) \in E_L} g^{tb} - \sum_{(b,t) \in E_L} g^{bt}$$

$$\geq g(S'_b) + q, \quad \text{by Constraint (14)}$$

$$\geq (k+1)q.$$

By induction, $g(S) \geq kq \forall S \subseteq V_L \setminus \{0\}$ and $k = |S|$. ■

Now we are ready to prove that $MCLC^R \leq WLF$. Given an optimal solution (q^*, g^*) to the formulation $M_{MCLC}^{\text{Dual,R}}$, the value of g^{st^*} is a feasible assignment of the variable w_{st} in the Weighted Load Factor formulation M_{WLF} . The corresponding objective value for this assignment is:

$$\min_{S \subset V_L, (i,j) \in E_P} \frac{\sum_{(s,t) \in \delta(S)} g^{st^*}}{\sum_{(s,t) \in \delta(S)} g^{st^*} f_{ij}^{st}}$$

$$\geq \frac{q^*}{\sum_{(s,t) \in \delta(S)} g^{st^*} f_{ij}^{st}}, \quad \text{by Lemma 11}$$

$$\geq q^*, \quad \text{by Constraint (13)}$$

which implies $WLF \geq q^*$. On the other hand, by Duality Theorem [19], the optimal value for M_{MCLC}^R is exactly q^* . Therefore we have $MCLC^R \leq WLF$.

Next, we prove that $WLF \leq MCLC$. Let C be the set of physical fibers that constitute a Min Cross Layer Cut, and let a be an arbitrary node in the logical network. Let $S_C \subset V_L$ be the set of nodes reachable from a after C has been removed from the physical network. It follows that all logical links in $\delta(S_C)$ use fibers in C .

Let w be the weight function on E_L that achieves the optimal Weighted Load Factor, and let $w(S_C)$ be the total weight of the logical links in $\delta(S_C)$. Also, let (i^*, j^*) be the physical fiber that carries the most weight for lightpaths in $\delta(S_C)$. The definition of WLF implies that:

$$WLF = \min_{S \subset V_L, (i,j) \in E_P} \frac{\sum_{(s,t) \in \delta(S)} w_{st}}{\sum_{(s,t) \in \delta(S)} w_{st} f_{ij}^{st}}$$

$$\leq \frac{\sum_{(s,t) \in \delta(S_C)} w_{st}}{\sum_{(s,t) \in \delta(S_C)} w_{st} f_{i^*j^*}^{st}}. \quad (15)$$

Next, since all logical links in $\delta(S_C)$ use fibers in C , we have:

$$\sum_{(s,t) \in \delta(S_C)} w_{st} \leq \sum_{(i,j) \in C} \sum_{(s,t) \in \delta(S_C)} w_{st} f_{ij}^{st}$$

$$\leq |C| \sum_{(s,t) \in \delta(S_C)} w_{st} f_{i^*j^*}^{st}. \quad (16)$$

Finally, combining inequalities (15) and (16), we have:

$$WLF \leq \frac{\sum_{(s,t) \in \delta(S_C)} w_{st}}{\sum_{(s,t) \in \delta(S_C)} w_{st} f_{i^*j^*}^{st}} \leq |C| = MCLC. \quad \blacksquare$$

REFERENCES

- [1] S. Subramaniam, M. Azizoglu, and A. K. Somani, "All-optical networks with sparse wavelength conversion," *IEEE/ACM Trans. Netw.*, vol. 4, no. 4, pp. 544–557, 1996.
- [2] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: a novel approach to high bandwidth optical WANs," *IEEE Transactions on Communications*, vol. 40, no. 7, pp. 1171–1182, 1992.
- [3] J. Q. Hu, "Diverse routing in optical mesh networks," *IEEE Transactions on Communications*, vol. 51, no. 3, Mar. 2003.
- [4] S. Yuan and J. P. Jue, "Dynamic lightpath protection in wdm mesh networks under wavelength-continuity and risk-disjoint constraints," *Comput. Netw.*, vol. 48, no. 2, pp. 91–112, 2005.
- [5] D. Xu, Y. Chen, Y. Xiong, C. Qiao, and X. He, "On the complexity of and algorithms for finding the shortest path with a disjoint counterpart," *IEEE/ACM Trans. Netw.*, vol. 14, no. 1, pp. 147–158, 2006.
- [6] R. Andersen, F. Chung, A. Sen, and G. Xue, "On disjoint path pairs with wavelength continuity constraint in WDM networks," in *IEEE INFOCOM*, 2004, pp. 524–535.
- [7] E. Modiano and A. Narula-Tam, "Survivable lightpath routing: A new approach to the design of WDM-based networks," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 4, May 2002.
- [8] J. Armitage, O. Crochat, and J.-Y. L. Boudec, "Design of a survivable WDM photonic network," in *IEEE INFOCOM*, Apr. 1997.
- [9] O. Crochat and J.-Y. L. Boudec, "Design protection for WDM optical networks," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 7, Sep. 1998.
- [10] Q. Deng, G. Sasaki, and C.-F. Su, "Survivable IP over WDM: a mathematical programming problem formulation," in *Proc 40th Allerton Conference on Communication, Control and Computing*, Oct. 2002.
- [11] M. Kurant and P. Thiran, "Survivable Routing of Mesh Topologies in IP-over-WDM Networks by Recursive Graph Contraction," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 5, 2007.
- [12] H. Lee, H. Choi, S. Subramaniam, and H.-A. Choi, "Survivable embedding of logical topologies in wdm ring networks," *Inf. Sci. Inf. Comput. Sci.*, vol. 149, no. 1-3, pp. 151–160, 2003.
- [13] A. Sen, B. Hao, and B. H. Shen, "Survivable routing in WDM networks," in *ISCC '02: Proceedings of the Seventh International Symposium on Computers and Communications (ISCC'02)*. Washington, DC, USA: IEEE Computer Society, 2002, p. 726.
- [14] —, "Minimum cost ring survivability in WDM networks," in *Workshop on High Performance Switching and Routing, 2003*, 2003.
- [15] A. Todimala and B. Ramamurthy, "Survivable virtual topology routing under Shared Risk Link Groups in WDM networks," in *BROADNETS '04: Proceedings of the First International Conference on Broadband Networks*. Washington, DC, USA: IEEE Computer Society, 2004.
- [16] H. Zang, C. Ou, and B. Mukherjee, "Path-protection routing and wavelength assignment in WDM mesh networks under duct-layer constraints," *IEEE/ACM Transactions on Networking*, vol. 11, Apr. 2003.
- [17] M. Kurant and P. Thiran, "Survivable Routing in IP-over-WDM Networks in the Presence of Multiple Failures," in *EuroNGI Workshop on Traffic Engineering, Protection and Restoration for NGI*, May 2006.
- [18] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, 1993.
- [19] D. Bertsimas and J. N. Tsitsiklis, *Introduction to Linear Optimization*. Athena Scientific, 1997.
- [20] L. Lovász, "On the ratio of optimal integral and fractional covers," *Discrete Mathematics*, vol. 13, pp. 383–390, 1975.
- [21] M. Garey and D. Johnson, *Computers and Intractability, a Guide to the Theory of NP-Completeness*. W.H. Freeman and Co., 1979.
- [22] C. Lund and M. Yannakakis, "On the hardness of approximating minimization problems," *J. ACM*, vol. 41, no. 5, pp. 960–981, 1994.
- [23] G. Ausiello, A. D'Atri, and M. Protasi, "Structure preserving reductions among convex optimization problems," *J. Comput. System Sci.*, vol. 21, pp. 136–153, 1980.
- [24] J. Hastad, "Clique is hard to approximate within $n^{(1-\epsilon)}$," *Acta Math.*, vol. 182, pp. 105–142, 1999.
- [25] D. D.-J. Kan, "Design of survivable IP-over-WDM networks: Providing protection and restoration at the electronic layer," *Master's Thesis, Massachusetts Institute of Technology*, Jun. 2003.
- [26] S. Arora, S. Rao, and U. Vazirani, "Expander flows, geometric embeddings and graph partitioning," in *STOC '04: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. New York, NY, USA: ACM, 2004, pp. 222–231.
- [27] K. Lee and E. Modiano, "Cross-layer survivability in WDM-based networks," Massachusetts Institute of Technology, Tech. Rep., 2008, available at <http://web.mit.edu/kylee/www/CrossLayerSurvivability.pdf>.
- [28] P. Raghavan and C. D. Tompson, "Randomized rounding: a technique for provably good algorithms and algorithmic proofs," *Combinatorica*, vol. 7, no. 4, pp. 365–374, 1987.
- [29] D. Banerjee and B. Mukherjee, "A practical approach for routing and wavelength assignment in large wavelength-routed optical networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 5, pp. 903–908, Jun. 1996.
- [30] A. E. Ozdaglar and D. P. Bertsekas, "Routing and wavelength assignment in optical networks," *IEEE/ACM Transactions on Networking*, vol. 11, no. 2, pp. 259–272, Apr. 2003.

PLACE
PHOTO
HERE

Kayi Lee (S'99 M'05) received the B.S. degrees in Computer Science and Mathematics and the M.Eng. degree in Computer Science from the Massachusetts Institute of Technology (MIT), Cambridge, in 1999 and 2000, respectively. He is currently a Software Engineer at Google Inc., while working towards his Ph.D. degree in Computer Science at MIT part-time. His research interests include routing and network survivability.

PLACE
PHOTO
HERE

Eytan Modiano (S'90-M'93-SM'00) received his B.S. degree in Electrical Engineering and Computer Science from the University of Connecticut at Storrs in 1986 and his M.S. and PhD degrees, both in Electrical Engineering, from the University of Maryland, College Park, MD, in 1989 and 1992 respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Post Doctoral Fellow during 1992–1993. Between 1993 and 1999 he was with MIT Lincoln Laboratory where he was the project leader

for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. Since 1999 he has been on the faculty at MIT; where he is an Associate Professor. His research is on communication networks and protocols with emphasis on satellite, wireless, and optical networks. He is the co-recipient of the Sigmetrics 2006 Best paper award for the paper "Maximizing Throughput in Wireless Networks via Gossiping", and the Wiopt 2005 best student paper award for the paper "Minimum Energy Transmission Scheduling Subject to Deadline Constraints."

He is currently an Associate Editor for The IEEE/ACM Transactions on Networking. He had served as Associate Editor for The IEEE Transactions on Information Theory, and as a guest editor for IEEE JSAC special issue on WDM network architectures; the Computer Networks Journal special issue on Broadband Internet Access; the Journal of Communications and Networks special issue on Wireless Ad-Hoc Networks; and for IEEE Journal of Lightwave Technology special issue on Optical Networks. He was the Technical Program co-chair for Wiopt 2006, IEEE Infocom 2007, and ACM MobiHoc 2007.

PLACE
PHOTO
HERE

Hyang-Won Lee (S'02-M'08) received his B.S., M.S. and Ph.D. degrees all in Electrical Engineering and Computer Science from Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea, in 2001, 2003 and 2007, respectively. He is currently a Postdoctoral Research Associate at the Massachusetts Institute of Technology, Cambridge, MA. His research interests are in the areas of congestion control, wireless resource allocation and network survivability.