Optical Networks

Series Editor: Biswanath Mukherjee

Suresh Subramaniam
Maïté Brandt-Pearce
Piet Demeester
Chava Vijaya Saradhi *Editors*

Cross-Layer Design in Optical Networks

Springer

Chapter 12
Cross-Layer Survivability*

Hyang-Won Lee, Kayi Lee, and Eytan Modiano

1 Introduction

The layered architecture of modern communication networks takes advantage of the flexibility of upper layer technology, such as IP, and the high data rates of lower layer technology, such as WDM. In particular, the WDM technology available today can support up to several terabits per second over a single fiber [9], making networks vulnerable to failures, because a failure for even a short period of time can result in a huge loss of data. The main theme of network survivability is to prevent such data loss by provisioning spare resources for recovery. In this chapter, we focus on the impact of layering on network survivability.

In the layered network, a logical topology is embedded onto a physical topology such that each logical link is spanned by using a path in the physical topology. This is often referred to as *lightpath routing*. Obviously, a single fiber cut can lead multiple logical links sharing the fiber to fail. Due to this correlation between logical link failures, the layered network survivability problem exhibits vastly different characteristics from the single-layer counterpart.

*Based on "Cross-Layer Survivability in WDM-based Networks," by K. Lee, E. Modiano and H. Lee which appeared in IEEE/ACM Transactions on Networking, vol. 19, no. 4, Aug. 2011, and "Reliability in Layered Networks with Random Link Failures," by K. Lee, H. Lee and E. Modiano which appeared in IEEE INFOCOM 2010, Mar. 2010.c 2011 IEEE.

H.-W. Lee (✉)
Konkuk University, Seoul 143-701, Republic of Korea
e-mail: leehw@konkuk.ac.kr

K. Lee
Google Inc, Cambridge, MA, USA
e-mail: kylee@mit.edu

E. Modiano
Massachusetts Institute of Technology, Laboratory for Information and Decision Systems,
77 Mass Ave, Cambridge, MA 02139, USA
e-mail: modiano@mit.edu

S. Subramaniam et al. (eds.), *Cross-Layer Design in Optical Networks*,
Optical Networks 15, DOI 10.1007/978-1-4614-5671-1_12,

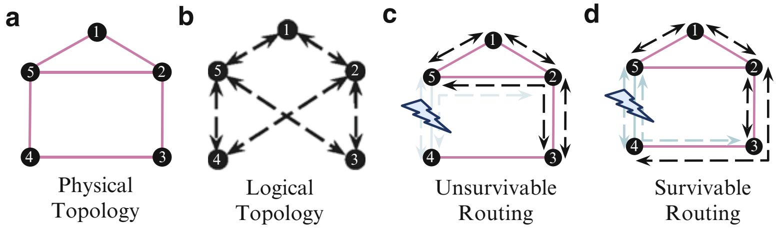

Fig. 12.1 Different lightpath routings can affect survivability. (**a**) Physical Topology, (**b**) Logical Topology, (**c**) Unsurvivable Routing, (**d**) Survivable Routing

The survivability of a layered network is dictated by the underlying lightpath routing. As an example, consider the physical and logical topologies shown in Fig. 12.1a,b. The lightpaths in the logical topology are routed over the physical topology in two different ways in Fig. 12.1c,d. In Fig. 12.1c, a failure of physical fiber (4, 5) would cause lightpaths (4, 5) and (2, 4) to fail. Consequently, node 4 will be disconnected from other nodes in the logical topology. On the other hand, in Fig. 12.1d, the logical topology will remain connected even if one of the fibers fails. This example demonstrates that to design a survivable layered network, one should carefully take into account the network structure across both the logical and physical layers. This is typically referred to as the *cross-layer survivability* problem.

In [2, 6, 7], the impact of physical layer failures on the connectivity of the logical topology was studied in the context of WDM-based networks. The authors proposed heuristic lightpath routing algorithms that minimize the number of disconnected logical node in the presence of a single physical link failure. The work of [15] was the first to introduce the notion of Survivable Lightpath Routing, which is defined to be a lightpath routing such that the logical topology remains connected in the event of a single fiber failure, and developed a mathematical formulation for finding a survivable lightpath routing. These results have been improved with more efficient formulation and extended to account for multiple physical failures [8, 11, 12, 18].

Most works in the literature consider the survivability as a constraint, however this chapter takes a more fundamental approach to addressing cross-layer survivability. In particular, in Sect. 12.2, we study connectivity parameters of a layered network, and observe that they exhibit vastly different properties compared to their single-layer counterparts. This observation motivates a new survivability metric called Min Cross Layer Cut (MCLC). The MCLC quantifies the connectivity of a layered network and is used to develop survivable lightpath routing algorithms. Simulation results show that these algorithms can find a better survivable layered network. Going beyond connectivity, a new survivability metric is introduced and analyzed in order to design a layered network that uses minimal spare capacity for protection against single-fiber failures. In Sect. 12.3 we study cross-layer survivability in the presence of random physical link failures, and in Sect. 12.4 we discuss future directions for cross-layer survivability.

2 Connectivity of Layered Networks

In this section we study key connectivity structures such as flows and cuts in multi-layer graphs in order to develop insights into cross-layer survivability. We will highlight the key differences in combinatorial properties between multi-layer graphs and single-layer graphs. In particular, it turns out 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 meaning in the layered setting. This motivates us to revisit fundamental issues such as quantifying and maximizing survivability in the layered setting.

2.1 Max Flow vs. Min Cut

For single-layer networks, the Max-Flow Min-Cut Theorem [1] states that the maximum number of disjoint paths between two nodes s and t is always the same as the minimum number of edges that need to be removed from the network in order to disconnect the two nodes. Let $\mathrm{MaxFlow_{st}}$ and $\mathrm{MinCut_{st}}$ be integral $s-t$ Max Flow and Min Cut respectively, and let $\mathrm{MaxFlow^R_{st}}$ and $\mathrm{MinCut^R_{st}}$ be their fractional (relaxed) values. The Max-Flow Min-Cut Theorem for single-layer networks can then be stated as follows:

$$\mathrm{MaxFlow_{st}} = \mathrm{MaxFlow^R_{st}} = \mathrm{MinCut^R_{st}} = \mathrm{MinCut_{st}}.$$

Consequently, the term *connectivity* between two nodes can be used unambiguously to refer to different measures such as the maximum number of disjoint paths or the minimum size cut, and this makes it a natural choice as the standard metric for measuring network survivability. The equality among these values has profound implications on survivable network design for single-layer networks. Because all these survivability measures take on the same value, it can naturally be used as the standard survivability metric that is applicable to measuring both disjoint paths or the 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.

Because of its fundamental importance, it is crucial to understand the Max-Flow Min-Cut relationship in layered networks. The following is a generalization of *Max Flow* and *Min Cut* to the layered setting:

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 modeled as $G_L = (V_L, E_L)$ in a similar fashion. The light path 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 $\mathscr{P}_{st}$ be the set of all $s - t$ paths in the logical topology. For each path $p \in \mathscr{P}_{st}$, 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:

$$\begin{aligned} &\text{MaxFlow}_{\text{st}}: \quad \text{Maximize} \sum_{p\in\mathscr{P}_{st}} f_p, \quad \text{subject to :} \\ &\sum_{p:(i,j)\in L(p)} f_p \le 1 \qquad \forall (i,j) \in E_P \\ &f_p \in \{0,1\} \qquad \forall p \in \mathscr{P}_{st} \end{aligned} \tag{12.1}$$

$$\begin{aligned} &\text{MinCut}_{\text{st}}: \quad \text{Minimize} \sum_{(i,j)\in E_P} y_{ij}, \quad \text{subject to :} \\ &\sum_{(i,j)\in L(p)} y_{ij} \ge 1 \qquad \forall p \in \mathscr{P}_{st} \\ &y_{ij} \in \{0,1\} \qquad \forall (i,j) \in E_P \end{aligned} \tag{12.2}$$

The variable f_p in the formulation $\text{MaxFlow}_{\text{st}}$indicates whether the path p is selected for the set of (s, t)-disjoint paths. Constraint (12.1) requires that no selected logical paths share a physical link. Similarly, in the formulation $\text{MinCut}_{\text{st}}$, the variable y_{ij} indicates whether the physical fiber (i, j) is selected for the minimum (s, t)-cut. Constraint (12.2) 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 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 us redefine $\text{MaxFlow}_{\text{st}}$ and $\text{MinCut}_{\text{st}}$ to be the optimal values of the above Max Flow and Min Cut formulations. We also denote $\text{MaxFlow}_{\text{st}}^{\text{R}}$ and $\text{MinCut}_{\text{st}}^{\text{R}}$ to be the optimal values to the linear relaxations of above Max Flow and Min Cut formulations.

First, it is easy to verify that the linear relaxations for the formulations $\text{MaxFlow}_{\text{st}}$ and $\text{MinCut}_{\text{st}}$ maintain a primal–dual relationship, which, by the Duality Theorem [5], implies that $\text{MaxFlow}_{\text{st}}^{\text{R}} = \text{MinCut}_{\text{st}}^{\text{R}}$. In addition, since any feasible solution to an integer program is also a feasible solution to the linear relaxation, the following relationship holds:

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

Therefore, just as with 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 $\text{MaxFlow}_{\text{st}}$,

$\text{MaxFlow}_{\text{st}}^{\text{R}}$ and $\text{MinCut}_{\text{st}}$ 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 parallel lightpaths. For simplicity of exposition, we omit the complete lightpath routing and only show the physical links that are shared by multiple lightpaths. In fact, it can be shown that for a two-node logical topology, any arbitrary fiber–sharing relationship can be realized by reconstructing a physical topology and lightpath routing [14]. Therefore, in the following discussion, we omit the details of the lightpath routing and only show the fiber–sharing relationship of our two-node logical topology.

In Fig. 12.1, 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 shares 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 $\text{MinCut}_{\text{st}}$ is 2 in this case. On the other hand, the value of $\text{MaxFlow}_{\text{st}}$ 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 $\text{MaxFlow}_{\text{st}}^{\text{R}}$ 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, Fig. 12.1 is an example where all three quantities differ.

It was shown in [14] that the gap between $\text{MaxFlow}_{\text{st}}$and $\text{MaxFlow}_{\text{st}}^{\text{R}}$ is $O(|E_L|)$, and the gap between $\text{MinCut}_{\text{st}}$ and $\text{MinCut}_{\text{st}}^{\text{R}}$ is $O(\log|E_L|)$. Thus, 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 necessarily guarantee the 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 Sect. 12.2.3, we introduce two new metrics that can be used to measure the connectivity of multi-layer networks.

2.2 Computational Complexity

In 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 [1]. However, computing and approximating their cross-layer equivalents turns out to be much more difficult. Theorem 1 describes the complexity of computing the Max Flow and Min Cut for multi-layer networks.

Theorem 1 ([14]) *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.*

In summary, the notion of survivability in multi-layer networks bears subtle yet important differences from its single-layer counterpart. Because of that, many new issues arise in the layered setting, including defining, measuring and optimizing survivability metrics. In what follows, our discussion will be focused on appropriate metrics for layered networks and efficient algorithms to maximize the cross-layer survivability.

2.3 Metrics for Cross-Layer Survivability

The previous section demonstrates some fundamental challenges in designing survivable layered network architectures. In particular, choosing the right metric to quantify survivability becomes an important and non-trivial question. Although the right metric will depend on the particular survivability requirement (e.g., disjoint paths or minimum cut), any reasonable metric must be *consistent* in that a network with a higher metric value should be more resilient to failures, *monotonic* in that any addition of physical or logical links to the network should not decrease the metric value, and *compatible* in that the metric should generalize the connectivity metric for single-layer networks.

Next, we introduce two metrics that measure the ability of the network to withstand multiple physical failures, while meeting the above criteria. Although the two metrics appear to measure different aspects of network connectivity, they are in fact closely related, as will be shown later.

2.3.1 Min Cross Layer Cut

The *Min Cross Layer Cut (MCLC)* is a natural generalization of Min Cut in single-layer networks. Similar to the way MinCut_{st} is defined in Sect. 12.2.1 between two given nodes s and t in the network, the Min Cross Layer Cut of a layered network is defined to be the smallest set of physical links whose removal will *globally* disconnect the logical network. A lightpath routing with high Min Cross Layer Cut value implies that the network remains connected even after a relatively large number of physical failures. It is also a generalization of the survivable lightpath routing definition in [15], 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:

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

For each S, computing H_S amounts to finding the Min Cut between the two partitions S and $V_L - S$. Therefore, by Theorem 1, computing H_S is also NP-Hard. Computing the Min Cross Layer Cut, which is defined to be the minimum among all H_S values, is therefore a difficult problem. However, for practical purposes, the MCLC of a large multi-layer network (e.g. 100 nodes) can be computed reasonably fast by solving 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 [14]:

$$\begin{aligned} &\mathrm{M_{MCLC}} : \text{Minimize} \sum\nolimits_{(i,j) \in E_P} y_{ij}, \quad \text{subject to :} \\ &d_t - d_s \le \sum\nolimits_{(i,j) \in E_P} y_{ij} f_{ij}^{st} \qquad \forall (s,t) \in E_L \end{aligned} \tag{12.3}$$

$$\begin{aligned} &\sum_{n \in V_L} d_n \ge 1, d_0 = 0 \\ &d_n, y_{ij} \in \{0,1\} \qquad \forall n \in V_L, (i,j) \in E_P \end{aligned} \tag{12.4}$$

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 (12.3) 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 (12.3) 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 (12.4) 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 Sect. 12.2.4, we will discuss several lightpath routing algorithms to maximize the MCLC value.

2.3.2 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 [10], provides such a measure of survivability. The WLF can be formulated as follows:

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

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 carried by a fiber within a single cut. Intuitively, if we interpret the weight to be the amount of traffic in the lightpath, the value z can be interpreted to be 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 is more capable of spreading its weight within any cut across the fibers.

Recall that in [15], a lightpath routing is defined to be survivable if the resulting layered network survives any single physical link failure. Hence, the survivable lightpath routing ensures that not all of the links in a logical cut share the same fiber, which in turn implies that the Weighted Load Factor $\frac{1}{z}$ is greater than one. Therefore, the Weighted Load Factor captures the connectivity of a layered network, generalizing the survivable lightpath routing. In fact, the Weighted Load Factor is closely related to Min Cross Layer Cut. Given a lightpath routing, let $\mathrm{M_{MCLC}}$ be the ILP formation for its Min Cross Layer Cut, and let *MCLC* and $MCLC^R$ be the optimal values for $\mathrm{M_{MCLC}}$ and its linear relaxation respectively. In addition, let *WLF* be the Weighted Load Factor of the lightpath routing. Then we have the following relationship [14]:

Theorem 2 $MCLC^{\mathrm{R}} \leq WLF \leq MCLC$

Therefore, although the two metrics appear to measure different aspects of network connectivity, they are inherently related. In fact, the two values are often identical as shown in Sect. 12.2.5.

2.4 Lightpath Routing Algorithms for Maximizing MCLC

A natural approach to maximizing the survivability of a layered network is to design a lightpath routing that maximizes the number of failures the network can withstand, i.e., maximizes Min Cross Layer Cut (MCLC). All the lightpath routing algorithms introduced in this section try to maximize the MCLC value. They are all 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:

$$\begin{aligned} &\mathrm{MCF}_{\mathcal{X}}: \quad \text{Minimize } \mathcal{X}(f), \quad \text{subject to:} \\ &f_{ij}^{st} \in \{0,1\} \\ &\{f_{ij}^{st} : (i,j) \in E_P\} \text{ forms an } (s,t)\text{-path}, \ \forall (s,t) \in E_L, \end{aligned}$$

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. The objective function $\mathcal{X}(f)$ depends on the lightpath routing f. Ideally, $\mathcal{X}(f)$ should be the MCLC value, however it turns out to be difficult to express the MCLC as a tractable function. For this reason, it is desired to develop a tractable objective that approximates the MCLC value [14].

2.4.1 Integer Programming Formulations

Let w be a weight assigned to each lightpath. 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.

Such an objective can be formulated as an integer linear program as follows:

$$\begin{aligned} &\mathrm{MCF}_w: \quad \text{Minimize } \rho_w, \quad \text{subject to:} \\ &\rho_w \geq \sum\nolimits_{(s,t)\in E_L} w(s,t) f_{ij}^{st} \ \forall (i,j) \in E_P \\ &f_{ij}^{st} \in \{0,1\} \\ &\{f_{ij}^{st} : (i,j) \in E_P\} \text{ forms an } (s,t)\text{-path}, \ \forall (s,t) \in E_L \end{aligned}$$

The routing strategy of the algorithm is determined by the weight function w, and 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. 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.

It is conceivable that if the weight function somehow reflects the connectivity structure of the logical graph, then it may lead to a better objective toward maximizing the MCLC of the solution. Consider a different weight function w_{MinCut} such that for each edge $(s, t) \in E_L$, the weight $w_{MinCut}(s, t)$ is defined 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 in MCF_w, Lee et al. [14] shows the following relationship between the objective value ρ_w of a feasible solution to $\mathrm{MCF}_w$and the Weighted Load Factor of the associated lightpath routing:

Theorem 3 ([14]) *For any feasible solution f of* $\mathrm{MCF_w}$ with $\mathrm{w_{MinCut}}$ as the weight function, $\frac{1}{\rho_w} \leq WLF$.

As a result of Theorems 2 and 3, the MCLC of a lightpath routing is lower bounded by the value of $\frac{1}{\rho_w}$, which the algorithm will try to maximize.

2.4.2 An Enhanced Multi-Commodity Flow Formulation

As we discussed in Sect. 12.2.3.2, the Weighted Load Factor provides a good lower bound on the MCLC of a lightpath routing. Here we discuss another multi-commodity flow based formulation whose objective function approximates the Weighted Load Factor of a lightpath routing. The formulation, denoted as $\mathrm{MCF_{LF}}$, can be written as follows:

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

Essentially, the formulation optimizes the *unweighted*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 4 ([14]) *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:*

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

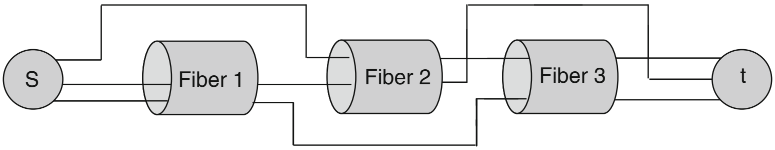

Fig. 12.2 A logical topology with 3 logical links where each pair of links shares a fiber in the physical topology

Therefore, the formulation MCF_{LF}gives a lightpath routing that is optimized for a better lower bound on the MCLC. An ILP is generally difficult to solve, and the above two formulations may not scale to large networks. Nonetheless, randomized rounding technique has been successfully used to solve multi-commodity flow problems lightpath [4, 17]. The formulations $MCF_w$and MCF_{LF}can also be solved via randomized rounding, and more details can be found in [14].

2.5 Simulation

In order to evaluate the performance of the algorithms introduced in the previous section, the NSFNET (Fig. 12.2) is augmented to have connectivity 4, and used as the physical topology. For logical topologies, 350 random graphs are generated such that each of them has connectivity 4 and its size ranges from 6 to 15 nodes. The MCLC values of the lightpath routings generated by the algorithms introduced in Sect. 12.2.4 will be compared as a measure of their survivability performance.

2.5.1 Survivability Performance of Different Lightpath Routing Formulations

We first study the survivability performance of the lightpath routings generated by the different formulations introduced in Sect. 12.2.4.1. Specifically, the following three algorithms are compared:

1. MinCut: Multi-Commodity Flow MCF_w, using the weight function w_{MinCut}
2. LF: Enhanced Multi-Commodity Flow MCF_{LF}.
3. SURVIVE: Survivable lightpath routing algorithm in [15] which computes the lightpath routing that minimizes the total fiber hops, subject to the constraint that the MCLC must be at least two.

Figure 12.4 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. MCLC ≥ 2), a more specialized formulation is required to achieve even higher MCLC values.

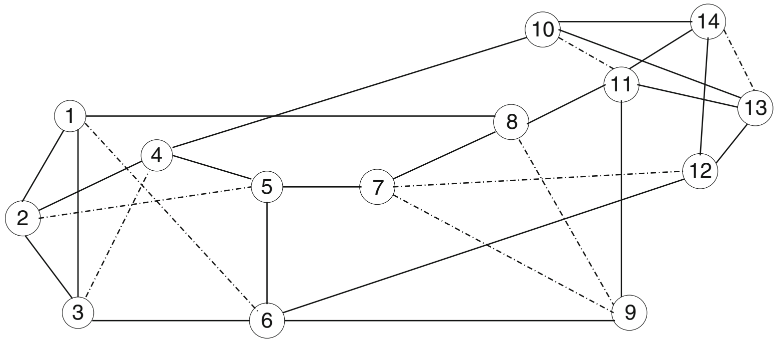

Fig. 12.3 The augmented NSFNET. The *dashed lines* are the new links

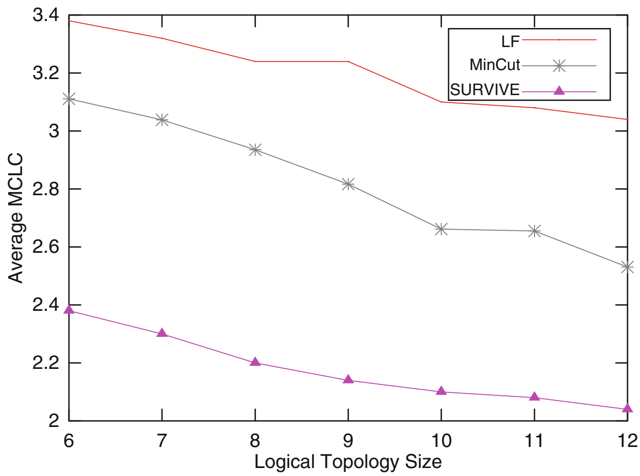

Fig. 12.4 MCLC performance of different lightpath routing formulations

The quality of the lightpath routing also depends on the graph structures captured by the formulations. Compared with MCF_{MinCut}, the 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 logical cut. As a result, the algorithm based on this enhanced formulation is able to provide lightpath routings with higher MCLC values.

2.5.2 Comparison Among Metrics

Recall the lower bounds on the Min Cross Layer Cut in Theorem 4. In this section, we study 350 different lightpath routings and measure these lower bound values for each of the lightpath routing. As Fig. 12.5 shows, the Weighted Load Factor is a

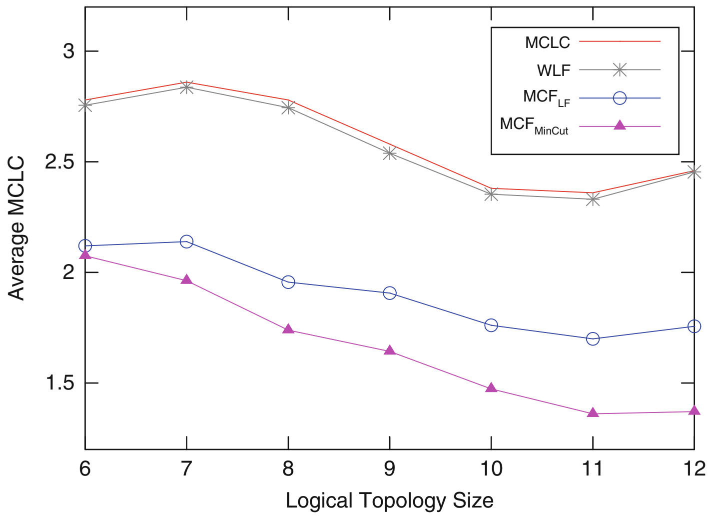

Fig. 12.5 Comparison among Min Cross Layer Cut (MCLC), Weighted Load Factor (WLF), and the optimal values of MCF_{LF} and MCF_{MinCut}

very close approximation of the Min Cross Layer Cut. Among the 350 routings studied, the two metrics are identical in 308 cases. This suggests a very tight connection between the two metrics, and the strong correlation between them 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 Sect. 12.2.4.1. 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 Fig. 12.4. Since there is still a large gap between the MCF_{LF} objective value and the MCLC value, this suggests room for further improvement with formulations that give a better MCLC lower bound. A good formulation that properly captures the cross-layer connectivity structure is essential for generating lightpath routings with high survivability, and it gives a powerful tool for designing highly survivable layered networks.

2.6 Beyond Connectivity

So far, we have assumed that the survivability of a layered network is guaranteed as long as the logical topology remains connected after a failure. Implicitly assumed here is that there is sufficient capacity in the network, so that the disrupted traffic can always be supported over available alternative paths. However, the capacity of the network is finite, and thus it may not always be possible to support the disrupted traffic. Therefore, it is important to take into account spare capacity as well as connectivity. To address this issue, we redefine survivability to meet two conditions: (1) the logical topology remains connected after any physical link failure and (2) there is sufficient capacity in the resulting network to support the traffic requirement.

As mentioned above, the design of survivable lightpath routing has focused only on the connectivity of the logical topology after a physical link failure. While there are a number of works dealing with spare capacity allocation, they assume single-link failures, and hence they are not applicable to the layered network where upon a single physical link failure, multiple logical links can fail simultaneously. In this section, joint survivable lightpath routing and capacity assignment problems are considered for layered networks. We discuss a new metric, first introduced in [10], that can measure the efficiency of spare capacity allocation. This metric is a generalization of the connectivity metrics discussed earlier to account for spare capacity, and can be used to formulate the problem of finding lightpath routings that guarantee efficient use of link capacity for protection.

The new metric, called *Load Factor*[10], quantifies the fraction of working capacity and spare capacity over each logical link. Assume that each link has capacity C. Let $\alpha(\in [0, 1])$ be the fraction of working capacity, i.e., αC is used for working paths, and subsequently, $(1 - \alpha)C$ is reserved for backup paths. Without loss of generality, we assume $C = 1$. For a given pair of logical and physical topologies and its lightpath routing, we define the load factor of the layered network to be the maximum value of α such that the two network survivability conditions mentioned above are satisfied. Clearly, the load factor measures the efficiency of capacity utilization, and it is desirable to find a lightpath routing with maximum load factor. In [10], a necessary and sufficient condition on the load factor was identified and used to develop an MILP formulation for finding a lightpath routing that maximizes the load factor. The results from [10] are discussed below.

Recall that N_L is the set of logical nodes. Denote by $CS(S)$ the cut set corresponding to a partition $\langle S, N - S \rangle$ of N_L. Given a routing of the logical topology denoted by $[f_{ij}^{st}, (i, j) \in E_P, (s, t) \in E_L]$, the following theorem gives a necessary and sufficient condition on the load factor.

Theorem 5 ([10]) *A network is survivable if and only if for every cut-set CS(S) of the logical topology and every physical link failure (i,j), the load factor satisfies the following inequality:*

$$\sum_{(s,t)\in CS(S)} (f_{ij}^{st} + f_{ji}^{st})\alpha \leq \sum_{(s,t)\in CS(S)} [1 - (f_{ij}^{st} + f_{ji}^{st})](1 - \alpha). \tag{12.5}$$

Rearranging the inequality (12.5), it can be shown that the load factor α is given by

$$\alpha = \min_{\substack{S \subset N_L \\ (i,j) \in E_P}} \frac{|CS(S)| - \sum_{(s,t)\in CS(S)} (f_{ij}^{st} + f_{ji}^{st})}{|CS(S)|}. \tag{12.6}$$

If there exists a cut-set $CS(S)$ and fiber (i, j) such that all the links in $CS(S)$ share (i, j), then the value of α is zero. Clearly, in this case, the logical topology is disconnected upon the failure of (i, j), meaning that the given lightpath routing cannot survive a single failure. Note that survivable routing algorithms in the literature such as [15] only guarantee $\alpha > 0$, i.e., the network remains connected after a single fiber failure. Therefore, the problem of finding a lightpath routing with maximum load factor α can be viewed as a generalization of finding a lightpath routing with connectivity guarantee.

The above result can be used to derive an optimal lightpath routing that maximizes the load factor. Let $\{f_{ij}^{st}\}^*$ be the routing that maximizes the load factor, and R denote the set of all possible routings. Then, it follows that

$$\{f_{ij}^{st}\}^* = \underset{f_{ij}^{st} \in R}{\arg\min} \max_{\substack{S \subset N_L \\ (i,j) \in E_P}} \frac{\sum_{(s,t) \in CS(S)} (f_{ij}^{st} + f_{ji}^{st})}{|CS(S)|}. \tag{12.7}$$

Note that the load factor is a special case of the weighted load factor discussed in the previous section where the weight function $w_{st} = 1$ for every logical link (s, t). The ratio in the above optimization is the fraction of logical links in a cut-set that will fail in the event of a fiber failure. Hence, a high ratio implies that the fiber is shared by many logical links, and it can be interpreted as the load on a fiber. The above formulation minimizes the maximum load on each fiber. Intuitively, this will minimize the amount of disrupted traffic in the event of a fiber cut, thereby reducing the demand for spare capacity. Using the representation of an optimal routing in (12.7), an MILP can be formulated for finding a lightpath routing with maximum load factor [10].

3 Extension to Random Failures

So far we considered single physical link failures, and discussed survivability metrics that account for a worst-case failure event. Failures in communication networks can be modeled as random events. It is thus important to understand the impact of random failures on the survivability of layered networks. In this section, we study the relationship between cross-layer connectivity metrics discussed in the previous section and the survivability of a layered network with random physical link failures. Interestingly, maximizing the MCLC value has the effect of maximizing cross-layer reliability in the low failure probability regime, and this observation can be used to develop reliable lightpath routing algorithms.

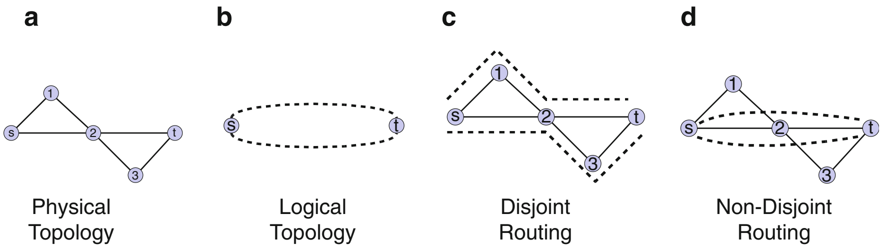

Fig. 12.6 Non-disjoint routings can sometimes be more reliable. (**a**) Physical Topology (**b**) Logical Topology (**c**) Disjoint Routing (**d**) Non-Disjoint Routing

3.1 Cross-Layer Reliability under Random Failures

Consider a layered network that consists of the logical topology $G_L = (V_L, E_L)$ built on top of the physical topology $G_P = (V_P, E_P)$ through a lightpath routing $f = [f_{ij}^{st}, (i, j) \in E_P, (s, t) \in E_L]$. If a physical link (i, j) fails, all of the logical links (s, t) carried over (i, j) (i.e., (s, t) such that $f_{ij}^{st} = 1$) also fail. A set S of physical links is called a *cross-layer cut* if the failure of the links in S causes the logical network to be disconnected. We also define the *network state* to be the subset S of physical links that failed. Hence, if S is a cross-layer cut, the network state S represents a *disconnected* network state. Otherwise, it is a *connected* state.

Each physical link fails independently with probability p. This probabilistic failure model represents a snapshot of a network where links fail and are repaired according to some Markovian process. Hence, p represents the steady-state probability that a physical link is in a failed state. The reliability of a multi-layer network is defined to be the probability that the logical network remains connected. We call this *cross-layer reliability*, and it is a natural survivability metric when the physical topology experiences random failures.

It is important to note that the cross-layer reliability depends on the underlying lightpath routing. For example, in Fig. 12.6, the logical topology consists of two parallel links between nodes s and t. Suppose every physical link fails independently with probability p. The first lightpath routing in Fig. 12.6c routes the two logical links using link-disjoint physical paths $(s, 1, 2, t)$ and $(s, 2, 3, t)$. Under this routing, the logical network will be disconnected with probability $(1 - (1 - p)^3)^2$. On the other hand, the second lightpath routing in Fig. 12.6d, which routes the two logical links over the same shortest physical route $(s, 2, t)$, has failure probability $2p - p^2$. While disjoint path routing is generally considered to be more reliable, it is only true in this example for small values of p. For large values of p(e.g. $p > 0.5$), the second lightpath routing is more reliable. Therefore, whether one lightpath routing is better than another depends on the value of p.

3.2 Cross-Layer Failure Polynomial

In single-layer networks, with random failures, reliability can be expressed as a polynomial in the failure probability p[3]. In [13], this polynomial expression was extended to the layered setting. It turns out that this expression provides important insights to the design of reliable lightpath routings.

Assume that there are m physical links, i.e.,$|E_P| = m$. The probability associated with a network state S with exactly i physical link failures (i.e., $|S| = i$) is $p^i(1-p)^{m-i}$. Let N_i be the number of cross-layer cuts S with $|S| = i$, then the probability that the network is disconnected is simply the sum of the probabilities over all cross-layer cuts, i.e.,

$$F(p) = \sum_{i=0}^{m} N_i p^i (1-p)^{m-i}. \tag{12.8}$$

Therefore, the failure probability of a multi-layer network can be expressed as a polynomial in p. The function $F(p)$ is called the *cross-layer failure polynomial* or simply the *failure polynomial*. The vector $[N_0, \ldots, N_m]$ plays an important role in assessing the reliability of a network. In particular, given the N_i values the reliability of the network can be computed using (12.8) for any value of p.

Clearly, if $N_i > 0$, then $N_j > 0, \ \forall j > i$, because any cut of size i will still be a cut with the addition of more failed links. The smallest i such that $N_i > 0$ is of special importance because it represents the Min Cross Layer Cut (MCLC) of the network, i.e., it is the minimum number of physical link failures needed to disconnect the logical network. Let d be the MCLC value of the network, and assume that it is a constant independent of the physical network size. Note that $N_i = 0, \ \forall i < d$, and the term $N_d p^d (1-p)^{m-d}$ in the failure polynomial dominates all other terms for small values of p. It was shown in [13] that there exists a region of probability p over which a lightpath routing with higher MCLC is more reliable than any lightpath routing with lower MCLC. Consequently, if a lightpath routing maximizes MCLC, i.e., make d as large as possible, it will achieve optimal reliability in the low failure probability regime.

Notice that we already discussed such lightpath routing algorithms in the previous section. In the following, we verify that these algorithms yield good reliability in the low failure probability regime.

3.3 Simulation

We used the augmented NSFNET (Fig. 12.3) as the physical topology, and generated 350 random logical topologies with size from 6 to 12 nodes and connectivity at least 4. We compare the reliability performance of the three lightpath

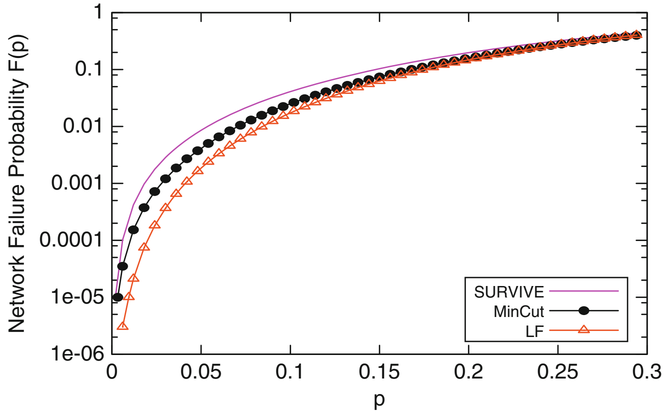

Fig. 12.7 Network failure probabilities of three different lightpath routing algorithms

routing algorithms SURVIVE, MinCut and LF presented in the previous section. For each lightpath routing generated by the algorithms, the failure polynomial is computed and compared.

The network failure probabilities of the three different lightpath routing algorithms are shown in Fig. 12.7, where for each algorithm, network failure probabilities were averaged over 350 different scenarios. When p is small, the two routings MinCut and LF which attempt to maximize the MCLC value are clearly more reliable than the SURVIVE algorithm. Note that in the lower failure probability regime, the algorithm LF whose MCLC value is higher finds a more reliable lightpath routing than the algorithm MinCut. This verifies that maximizing MCLC is a good strategy for maximizing reliability in the low failure probability regime.

4 Future Directions

In this chapter, we reviewed recent advances in cross-layer survivability. In particular, we introduced several metrics that measure the survivability of a layered network, and discussed survivable lightpath routing algorithms based on these metrics. These metrics capture the fundamentals of cross-layer survivability. We believe that many results are yet to be discovered in this context, and envision that the metrics discussed in this chapter will play an important role toward a theory of cross-layer survivability.

While this chapter focused on the role of lightpath routing in cross-layer survivability, the survivable network design problem in a layered setting consists of three components: logical topology design, physical topology design, and lightpath routing algorithm design. Obviously, the connectivity performance of a layered network is limited by the logical and physical topology. For instance, the MCLC value of a layered network is no greater than the min-cut value of either the logical or physical topology. Therefore, for survivable layered network design, it is necessary to have logical and physical topologies that allow *good* light path routing. Note, however, that logical and physical topologies with better connectivity do not necessarily guarantee a more survivable layered network because there may not exist a mapping of the logical topology to physical topology that leads to better survivability. Therefore, when designing a physical topology, the logical topology should be taken into account and vice versa. As a consequence, the results in the survivable single layer network design may not be applicable to the survivable logical and physical topology design problem. This makes the topology design problem an interesting problem for future research.

Indeed, addressing the topology design problem in the layered setting has been largely unexplored. In [16], necessary conditions on physical topologies were developed to ensure that a ring logical topology can be embedded and survive a single fiber failure. These conditions are then used to find lower bounds on the number of physical links needed for such an embedding to exist. Despite this work, the problem of topology design remains largely unexplored.

Acknowledgements This work was supported by NSF grants CNS-0626781 and CNS-0830961 and by DTRA grants HDTRA1-07-1-0004 and HDTRA-09-1-0050. Hyang-Won Lee was supported in part by the Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education, Science and Technology (2012R1A1A1012610).

References

1. Ahuja RK, Magnanti TL, Orlin JB (1993) Network flows: theory, algorithms, and applications. Prentice-Hall, Englewood Cliffs, NJ
2. Armitage J, Crochat O, Boudec JYL (1997) Design of a survivable WDM photonic network. In: Proceedings of the sixteenth IEEE International Conference on Computer Communications, Kobe, Japan
3. Ball M (1980) Complexity of network reliability computations. Networks 10:153–165
4. Banerjee D, Mukherjee B (1996) A practical approach for routing and wavelength assignment in large wavelength-routed optical networks. IEEE J Sel Area Comm 14(5):903–908
5. Bertsimas D, Tsitsiklis JN (1997) Introduction to linear optimization. Athena Scientific, Belmont, Massachusetts
6. Crochat O, Boudec JYL (1998) Design protection for WDM optical networks. IEEE J Sel Area Comm 16(7):1158–1165
7. Crochat O, Boudec JYL, Gerstel O (2000) Protection interoperability for WDM optical networks. IEEE ACM Trans Netw 8(3):384–395
8. Deng Q, Sasaki G, Su CF (2002) Survivable IP over WDM: a mathematical programming problem formulation. In: Proceedings of the 40th Allerton conference on communication, control and computing, Monticello, IL
9. Elhanany I, Kahane M, Sadot D (2001) Packet scheduling in next-generation multiterabit networks. Computer 34(4):104–106

10. Kan DDJ (2003) Design of survivable IP-over-WDM networks: Providing protection and restoration at the electronic layer. Master's Thesis, Massachusetts Institute of Technology
11. Kurant M, Thiran P (2006) Survivable routing in IP-over-WDM networks in the presence of multiple failures. In: EuroNGI workshop on traffic engineering, protection and restoration for NGI, Kraków, Poland
12. Kurant M, Thiran P (2007) Survivable routing of mesh topologies in IP-over-WDM networks by recursive graph contraction. IEEE J Sel Area Comm 25(5):922–933. DOI 10.1109/JSAC.2007.070606
13. Lee K, Lee HW, Modiano E (2010) Reliability in layered networks with random link failures. In: IEEE INFOCOM. San Diego, CA
14. Lee K, Modiano E, Lee HW (2011) Cross-layer survivability in WDM-based networks. In: IEEE/ACM Trans. on Networking 19(4):1000–1013
15. Modiano E, Narula-Tam A (2002) Survivable lightpath routing: A new approach to the design of WDM-based networks. IEEE J Sel Area Comm 20(4):800–809
16. Narula-Tam A, Modiano E, Brzezinski A (2004) Physical topology design for survivable routing of logical rings in wdm-based networks. IEEE J Sel Area Comm 22(8):1525–1583
17. Raghavan P, Tompson CD (1987) Randomized rounding: a technique for provably good algorithms and algorithmic proofs. Combinatorica 7(4):365–374. DOI http://dx.doi.org/10.1007/BF02579324
18. Todimala A, Ramamurthy B (2004) Survivable virtual topology routing under Shared Risk Link Groups in WDM networks. In: BROADNETS '04: Proceedings of the first international conference on broadband networks, pp. 130–139. IEEE Computer Society, Washington, DC, USA. DOI http://dx.doi.org/10.1109/BROADNETS.2004.81