This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE Globecom 2011 proceedings.

Maximizing Reliability in WDM Networks through Lightpath Routing

Hyang-Won Lee, Kayi Lee and Eytan Modiano
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: {hwlee, kylee, modiano}@mit.edu

***Abstract*—We study the reliability maximization problem in WDM networks with random link failures. Reliability in these networks is defined as the probability that the logical network is connected, and it is determined by the underlying lightpath routing and the link failure probability. We show that in general the optimal lightpath routing depends on the link failure probability, and characterize the properties of lightpath routings that maximize the reliability in different failure probability regimes. In particular, we show that in the low failure probability regime, maximizing the "cross-layer" min cut of the (layered) network maximizes reliability, whereas in the high failure probability regime, minimizing the spanning tree of the network maximizes reliability. Motivated by these results, we develop lightpath routing algorithms for reliability maximization.**

I. Introduction

Modern communication networks are constructed using a layered approach, with one or more electronic layers (e.g., IP, ATM, SONET) built on top of an optical fiber network. The survivability of such networks under fiber failures largely depends on how the logical electronic topology is embedded onto the physical fiber topology. In the context of WDM networks, this is known as *lightpath routing*. However, finding a reliable lightpath routing is rather challenging because it must take into account the sharing of physical fibers by logical links and its impact on the connectivity of the logical topology. In this paper, we focus on characterizing and developing reliable lightpath routings for layered networks assuming that physical links experience random failures; as adopted by several works in the literature [1]–[3].

A natural survivability metric in this context is the probability that given a lightpath routing, the logical topology remains connected; we call this probability the *cross-layer (network) reliability*. The cross-layer reliability reflects the survivability performance achieved by the lightpath routing. Hence, it is desirable to design a lightpath routing that maximizes the reliability. This layered network reliability problem remains largely unexplored. Although there has been a large body of work in the context of survivable lightpath routings for layered networks [4]–[11], most of these works address the problem of finding a lightpath routing that can survive a single physical link failure. In [12], we studied lightpath routing algorithms that maximize the connectivity of a layered network, i.e., maximize the number of physical link failures that the logical network can survive. Our work in [13] was the first study that deals with reliability in a cross-layer setting. In particular, we extended the polynomial expression for single-layer network reliability to the layered setting, and developed approximation algorithms for reliability computation.

In contrast to the layered case, network reliability has been extensively studied in the single-layer setting, with a focus on reliability computation [14] and reliable network design [15]. However, the layered network design problem is fundamentally different from the single-layer network design problem because in the layered setting, logical topology, physical topology and lightpath routing algorithm should be jointly designed. Consequently, the results from the single-layer setting are not applicable to the layered setting. Our work in this paper focuses on reliable lightpath routing design assuming that the logical and physical topologies are given.

Our contributions can be summarized as follows: In Section II, we present the network model and *show that in general the optimal lightpath routing depends on the link failure probability*. In Section III, we *identify the conditions for optimal lightpath routings in different failure probability regimes*. Namely, in the low probability regime, maximizing the min cut of the (layered) network maximizes reliability, whereas in the high probability regime, minimizing the spanning tree of the network maximizes reliability. In Section IV, we *develop lightpath routing algorithms that maximize the reliability in both the low and high failure probability regimes*, and in Section V, we present extensive simulation results.

II. Model and Background

We consider a layered network $\mathcal{G}$ 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, where V and E are the set of nodes and links respectively. In the context of WDM network, a logical link is called a *lightpath*, and each lightpath is routed over the physical topology. This *lightpath routing* is denoted by $f = [f_{ij}^{st}, (i,j) \in E_P, (s,t) \in E_L]$, where f_{ij}^{st} takes the value 1 if logical link (s,t) is routed over physical link (i,j), and 0 otherwise.

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

This work was supported by NSF grants CNS-0830961 and CNS-1017800, and by DTRA grants HDTRA1-07-1-0004 and HDTRA-09-1-0050.

some Markovian process. Hence, p represents the steady-state probability that a physical link is in a failed state. This model has been adopted by several previous works [1]–[3].

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* as 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.

A. Importance of Lightpath Routing

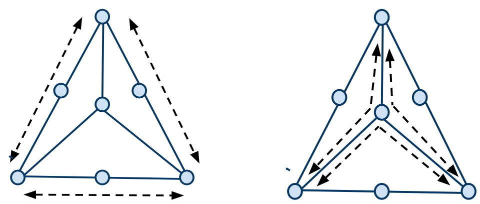

(a) Optimal Routing in Low Regime

(b) Optimal Routing in High Regime

Fig. 1. Example showing that optimal routings depend on the value of p. Physical topology is solid line, logical topology is the triangle formed by the 3 corner nodes and 3 edges, and lightpath routing is dashed line.

The reliability of a multi-layer network is defined as the probability that the logical network remains connected. It is important to note that the reliability depends on the underlying lightpath routing. For example, Fig. 1 shows two different lightpath routings. In Fig. 1(a), the logical links are routed over physically disjoint paths, and its reliability is given by $3(1-p)^4 - 2(1-p)^6$. In contrast, in Fig. 1(b), every pair of logical links share a physical link, and its reliability is $(1-p)^3$. While disjoint path routing is considered to be more reliable, it is easy to see that in this example the disjoint routing has better reliability only for small values of p whereas for large p (e.g., $p > 0.7$) the non-disjoint routing is more reliable. Therefore, whether one lightpath routing is better than another depends on the value of p. In Section III, we will investigate how the reliability and the connectivity parameters of a layered network are related for different values of p. The key to this study is the polynomial expression for the cross-layer reliability, which we introduce below.

B. Failure Polynomial and Connectivity Parameters

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.,

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

Therefore, the failure probability of a multi-layer network can be expressed as a polynomial in p. The function $F(p)$ will be called the *cross-layer failure polynomial* or simply the *failure polynomial*. The coefficients N_i's contain the information on the structure of a layered graph, determined by the underlying lightpath routing. Below we introduce some important coefficients related to connectivity.

Each N_i represents the number of cross-layer cuts of size i in the network. Define a *Min Cross Layer Cut (MCLC)* as a smallest set of physical links needed to disconnect the logical network. Denote by d the size of MCLC, then d is the smallest i such that $N_i > 0$, meaning that the logical network will not be disconnected by fewer than d physical link failures. The MCLC is a generalization of single-layer min-cut to the multi-layer setting [12].

A *cross-layer spanning tree* is defined as a *minimal* set of fibers whose survival keeps the logical network connected. Hence, if T is a cross-layer spanning tree, then the survival of just $T \setminus \{(i, j)\}$ renders the logical network disconnected for any fiber $(i, j) \in T$. Note that this is a generalization of the single-layer spanning tree. However, unlike a single-layer graph where all spanning trees have the same size, in a layered graph, spanning trees can have different sizes. Thus, we define a *Min Cross Layer Spanning Tree (MCLST)* as a cross-layer spanning tree with minimum number of physical links. Let b be the size of MCLST, then b is the smallest i such that $N_{m-i} < \binom{m}{m-i}$, meaning that there is a set of physical links of size i whose survival would keep the logical network connected.

Note that for given logical and physical topologies, MCLC and MCLST are all determined by the lightpath routing. For example, in Fig. 1, the disjoint routing has $d = 2$ and $b = 4$, whereas the non-disjoint routing has $d = 1$ and $b = 3$. As mentioned above, the disjoint routing has better reliability for small p, and the non-disjoint routing has better reliability for large p. This example suggests that maximizing MCLC may lead to better reliability for small p, while minimizing MCLST may lead to better reliability for large p. It turns out that this is true in general, and will be further discussed in Section III.

III. Properties of Optimal Lightpath Routings

Next, we study the properties of optimal lightpath routings for different failure probability regimes. These properties will give insight on how routings should be designed for better reliability. Proofs are omitted for brevity and can be found in [16].

A. Uniformly Optimal Routings

An important question in this study is the existence of routings that are most reliable for all failure probabilities. We begin with the following definition:

Definition 1: For given logical and physical topologies, a lightpath routing is said to be *uniformly optimal* if its reliability is greater than or equal to that of any other lightpath routing for every value of p.

Therefore, a uniformly optimal lightpath routing yields the best reliability for any value $p \in [0, 1]$. While it is desirable to design a uniformly optimal routing, such a routing does not always exist. Consider the example in Fig. 1 again. It can be shown that in the low failure probability regime a disjoint routing such as the one in Fig. 1(a) is optimal as it maximizes the MCLC. In contrast, in the high failure probability regime, the routing in Fig. 1(b) can be shown to be optimal as it minimizes the MCLST. In brief, there does not always exist a lightpath routing that minimizes $F(p)$ uniformly over all values of p.

Due to this difficulty, we focus on non-uniformly (or locally) optimal routings, which will give more tractable design criteria for lightpath routing algorithms. Theorem 1 below is a crucial result in this study; namely, it reveals a connection between local optimality and uniform optimality. First, a lightpath routing is said to be *locally optimal* if there exist p_1, $p_2 \in [0, 1]$, where $p_1 < p_2$, such that it yields the best reliability for every $p \in [p_1, p_2]$.

Theorem 1: Consider a pair of logical and physical topologies (G_L, G_P) for which there exists a uniformly optimal routing. Then, any locally optimal lightpath routing for (G_L, G_P) is also uniformly optimal.

Motivated by this result, we study locally optimal lightpath routings. In particular, we develop the conditions for a lightpath routing to be optimal for both the low failure probability regime (small p) and high failure probability regime (large p).

B. Low Failure Probability Regime

It is easy to see that in the failure polynomial, the terms corresponding to small cross-layer cuts dominate when p is small. Hence, for reliability maximization in the low failure probability regime, it is desirable to minimize the number of small cross-layer cuts. We use this intuition to derive the properties of optimal routings for small p. We begin with the following definition:

Definition 2: Consider two lightpath routings 1 and 2. Routing 1 is said to be *more reliable* than routing 2 *in the low failure probability regime* if there exists a positive number p_0 such that the reliability of routing 1 is higher than that of routing 2 for $0 < p < p_0$. A lightpath routing is said to be *locally optimal in the low failure probability regime* if it is more (or equally) reliable than any other routing in the low failure probability regime.

Let d_j be the size of the MCLC under routing $j(= 1, 2)$. Let N_i and M_i be the numbers of cross-layer cuts of size i under routings 1 and 2 respectively. We call the vector $N = [N_i, \forall i]$ the *cut vector*. The following is an example of cut vectors N and M with $d_1 = 4$ and $d_2 = 3$:

$$\begin{array}{cccccccc} i & 0 & 1 & 2 & 3 & 4 & 5 & \cdots & m \\ N_i & 0 & 0 & 0 & 0 & 20 & 26 & \cdots & 1 \\ M_i & 0 & 0 & 0 & 9 & 19 & 30 & \cdots & 1. \end{array}$$

Using cut vectors of lightpath routings, we define *lexicographical ordering* as follows: Routing 1 is lexicographically smaller than routing 2 if $N_{i^*} < M_{i^*}$ where i^* is the smallest i at which N_i and M_i differ. In the above example, we have $i^* = 3$ and $N_{i^*} < M_{i^*}$, hence routing 1 is lexicographically smaller. Therefore, if a lightpath routing is lexicographically smaller than another, it has fewer small cross-layer cuts and thus yields better reliability for small p.

Theorem 2: If routing 1 is lexicographically smaller than routing 2, then routing 1 is more reliable than routing 2 in the low failure probability regime.

Clearly, Theorem 2 leads to a local optimality condition; that is, if a lightpath routing minimizes the cut vector lexicographically, then it is locally optimal in the low failure probability regime. An interesting case is when routing 1 has larger MCLC than routing 2 (as in the above example). In this case, routing 1 is lexicographically smaller than routing 2 and hence it follows from Theorem 2 that:

Corollary 1: If $d_1 > d_2$, then routing 1 is more reliable than routing 2 in the low failure probability regime.

Consequently, a lightpath routing with the maximum size MCLC yields the best reliability for small p. Similarly, routing 1 is also lexicographically smaller than routing 2 when they have the same size of MCLC but routing 1 has fewer MCLCs. This leads to the following result:

Corollary 2: If $d_1 = d_2$ and $N_{d_1} < M_{d_2}$, then routing 1 is more reliable than routing 2 in the low probability regime.

Therefore, for reliability maximization in the low failure probability regime, it is desirable to maximize the size of the MCLC while minimizing the number of such MCLCs. This condition will be used to develop lightpath routing algorithms in Section IV.

C. High Failure Probability Regime

We have seen that small cuts are dominant when p is small, and it is important to minimize the number of small cuts. Analogously, for large p, large cuts are dominant, and hence, minimizing the number of large cuts would result in maximum reliability. In other words, the cut vector should be minimized for large cuts for better reliability in the high failure probability regime (A similar lexicographical ordering can be defined but is omitted for brevity). Similar to the case of low probability regime, we define the following:

Definition 3: Consider two lightpath routings 1 and 2. Routing 1 is said to be *more reliable* than routing 2 *in the high failure probability regime* if there exists a number $p_0 < 1$ such that the reliability of routing 1 is higher than that of routing 2 for $p_0 < p < 1$.

An important parameter in this case is the Min Cross Layer Spanning Tree (MCLST), because logical networks with small MCLST may remain connected even if only a small number of physical links survive due to high failure probability. Let b_j be the size of MCLST for routing $j(= 1, 2)$. The following two theorems characterize the routings with better reliability in the high failure probability regime.

Theorem 3: If $b_1 < b_2$, then routing 1 is more reliable than routing 2 in the high failure probability regime.

Theorem 4: If $b_1 = b_2$ and $N_{m-b_1} < M_{m-b_2}$, then routing 1 is more reliable than routing 2 in the high failure probability regime.

Therefore, for reliability maximization in the high failure probability regime, it is desirable to find a lightpath routing

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE Globecom 2011 proceedings.

that minimizes the size of MCLST and maximizes the number of MCLSTs.

IV. Maximizing Reliability through Lightpath Routing

We have seen in the previous section that optimizing reliability for low and high failure probability regimes requires different objectives. In this section, we discuss lightpath routing algorithms optimized for these different objectives.

A. *Low Failure Probability Regime*

In the low failure probability regime, reliability is maximized by a lightpath routing that minimizes the cut vector lexicographically. However, expressing such objective precisely as a mathematical formulation can be difficult since the elements N_i in the cut vector are correlated. One can simplify the problem by focusing only on maximizing the MCLC, which is a necessary condition for lexicographical minimization. This simplified version of the problem has been previously studied in [12], where Multi-Commodity Flow (MCF) based formulations were developed to route the logical links so as to maximize lower bounds on the MCLC. However, developing a simple lightpath routing formulation that maximizes the MCLC remains an open problem. Moreover, joint routing of the logical links to minimize the lexicographical ordering appears to be intractable; as it is known that just computing the MCLC is NP-hard [12].

In this section, we introduce a novel *lightpath rerouting* approach, which incrementally reroutes a single logical link in each step in order to reduce the number of small cross-layer cuts. Fig. 2 shows a simple example of how rerouting can eliminate small cuts. Initially, the Min Cross Layer Cut size of the lightpath routing is 1 and there are three cross-layer cuts of this size. The logical links are then rerouted sequentially so that the MCLCs can be converted into non-cuts. At the end, the MCLC value of the lightpath routing is increased to 2. Note that it can be further improved by switching the two-hop path in Fig. 2(c) to the single-hop path.

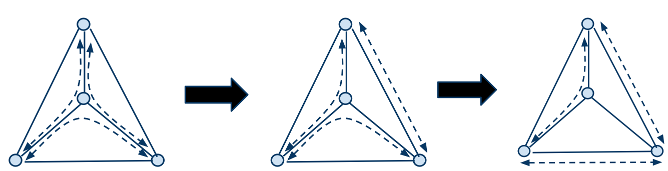

(a) $d = 1, N_d = 3$ (b) $d = 1, N_d = 1$ (c) $d = 2, N_d = 5$

Fig. 2. Improving reliability via rerouting. The physical topology is in solid lines, logical topology is the triangle formed by the 3 corner nodes and 3 edges, and the lightpath routing is in dashed lines. The MCLC value and the number of MCLCs in the lightpath routings are denoted by d and N_d.

Fig. 3 is a general description of our rerouting framework, which works as follows: Given any initial lightpath routing,

(1) Select a logical link, say (s, t), and reroute (s, t) to reduce the number of MCLCs.
(2) Repeat (1) until no further improvement is possible.

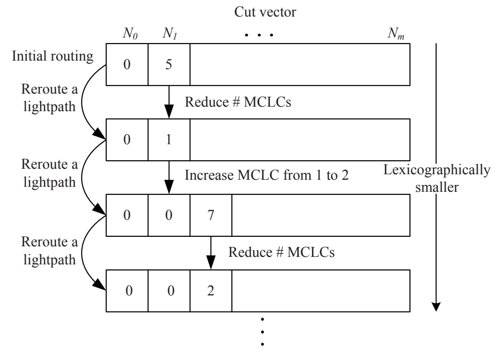

Fig. 3. The lightpath rerouting framework

Therefore, each iteration will reduce the number of MCLCs, and possibly increase the size of the MCLC if every MCLC happens to be converted into non-cuts, thereby minimizing the cut vector lexicographically. The rerouting terminates, if no further improvement is possible by rerouting a single lightpath. Since each iteration computes a physical route for only one logical link, this approach effectively breaks down the joint lightpath routing problem into multiple smaller steps, which helps to improve the overall running time. As we will see in Section V, this rerouting approach is very effective in obtaining lightpath routings with better reliability than existing methods such as the MCF approach of [12].

Suppose that an initial lightpath routing is given, and let d be the size of the MCLC under the initial routing. When the physical route of a logical link changes, some of the cross-layer cuts will be converted into non-cuts, and some non-cuts will be converted into cross-layer cuts. In the low failure probability regime, the reliability will be improved by the rerouting if the conversion of small cross-layer cuts to non-cuts outnumbers the conversion in the opposite direction. Therefore, we can formulate the lightpath rerouting as an optimization problem to maximize the reduction in the number of MCLCs.

It is, however, important to make sure that the non-cuts of size smaller than d remain non-cuts after rerouting, because otherwise, the new routing can decrease the MCLC value. It can be shown that only the non-cuts of size $d - 1$ can be possibly converted into cuts by a single-link rerouting, hence we only need to guarantee that the non-cuts of size $d - 1$ are not converted into cuts. Therefore, the rerouting problem can be described as follows:

Find	A lightpath and its physical path that
Maximize	Net reduction in # MCLCs
Subject to	Each non-cut of size $d - 1$ remains a non-cut

We can formulate this problem as an ILP, and the detailed ILP formulation can be found in [16]. The optimal solution to the ILP gives the logical link whose rerouting will maximize the net reduction in the number of MCLCs, and the optimal reroute. We can then iteratively solve the ILP based on the improved solution until no further improvement can be obtained.

B. High Failure Probability Regime

The key idea behind rerouting is that the size of the MCLC can be increased by converting all MCLCs into non-cuts. While in theory the same idea can be applied to convert large cross-layer cuts into large non-cuts, (which correspond to small spanning trees), the set of large cuts to consider can be prohibitively large, making the rerouting approach infeasible for minimizing spanning trees. On the other hand, unlike MCLCs, MCLST minimization can be formulated by a simple ILP. Define lightpath routing variable f_{ij}^{st} to be 1 if logical link (s,t) is routed over physical link (i,j), and 0 otherwise. Let $y_{ij} = 1$ if physical link (i,j) survives, and 0 otherwise. Let $z^{st} = 1$ if lightpath (s,t) survives, 0 otherwise.

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

$$(|V_L| - 1) \cdot z^{st} \geq x^{st}, \quad \forall (s,t) \in E_L \quad (2)$$

$$y_{ij} \geq z^{st} + f_{ij}^{st} - 1 \quad \forall (s,t) \in E_L, \forall (i,j) \in E_P \quad (3)$$

$$\sum_{t\in V_L} x^{st} - \sum_{t\in V_L} x^{ts} = \begin{cases} |V_L| - 1, & \text{if } s = 0 \\ -1, & \text{if } s \in V_L - \{0\} \end{cases}$$

$$\{(i,j) : f_{ij}^{st} = 1\} \text{ forms an } (s,t)\text{-path in } G_P, \forall (s,t) \in E_L$$

$$0 \leq y_{ij} \leq 1; \quad 0 \leq x^{st}; \quad z_{ij}, f_{ij}^{st} \in \{0,1\}$$

The variables x^{st} represent a flow on the logical topology where 1 unit of flow is sent from logical node 0 to every other logical node. Constraint (2) requires these flows to be carried only on the surviving logical links, which implies the surviving links form a connected logical subgraph. Constraint (3) ensures the survival of physical links that are used by any surviving logical links. Therefore, the set of physical links (i,j) with $y_{ij} = 1$ forms a cross-layer spanning tree. As a result, the optimal solution to the above ILP yields a lightpath routing that minimizes the size of the MCLST.

V. Simulation Results

All simulations are based on the augmented NSFNET (Figure 4) as the physical topology, and 350 random logical topologies with size ranging from 6 to 12 nodes and connectivity at least 4. We use the algorithms from [12] and the rerouting algorithm introduced in Section IV to generate different sets of lightpath routings, and study their properties such as MCLC, MCLST and reliability.

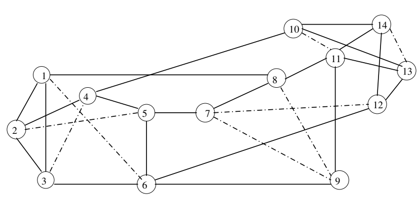

Fig. 4. The Augmented NSFNET. The dashed lines are the new links.

A. Lightpath Routings Optimized for Different Probability Regimes

We first compare the lightpath routings optimized for different probability regimes. We use the rerouting algorithm in Section IV-A, and the formulation MCLST in Section IV-B to generate two sets of lightpath routings. The two sets, namely $\mathsf{LPR}_{\mathsf{Low}}$ and $\mathsf{LPR}_{\mathsf{High}}$, represent routings optimized for the low and high failure probability regimes.

Figure 5 shows the average MCLC and MCLST values for the two sets of lightpath routings. There are noticeable differences in the values between the two sets, suggesting that the two objectives can lead to vastly different lightpath routings. In Figure 6, the survivability, both in terms of reliability and unreliability (i.e., 1 - reliability), of the pair over different link failure probabilities is shown. As expected, when the link failure probability is small, the lightpath routings in $\mathsf{LPR}_{\mathsf{Low}}$ achieve higher reliability. In particular, when the link failure probability approaches 0, there is an order of magnitude difference in terms of unreliability, meaning that maximizing the size of MCLC can have significant impact in the network reliability. As the link failure probability increases, it becomes more important to minimize the size of MCLST, so $\mathsf{LPR}_{\mathsf{High}}$ is able to achieve higher reliability in that regime. However, the difference in reliability is not as prominent as in the previous case.

In practical settings, the failure probability of individual physical links is typically very small. Therefore, our simulation result suggests that minimizing the lexicographic ordering of the lightpath routings can often lead to meaningful improvement in network survivability.

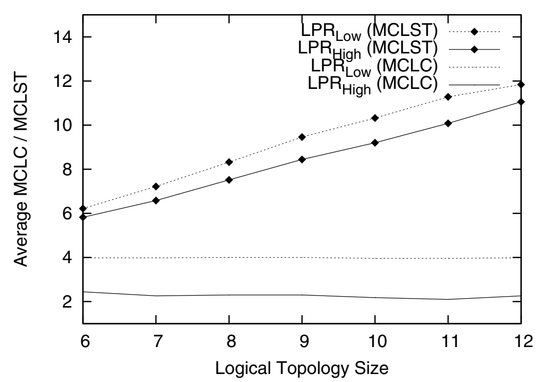

Fig. 5. Lightpath routings optimized for different probability regimes have different MCLC and MCLST values. $\mathsf{LPR}_{\mathsf{Low}}$ are lightpath routings optimized for MCLC, and $\mathsf{LPR}_{\mathsf{High}}$ are lightpath routings optimized for MCLST.

B. Robustness of Rerouting to Initial Lightpath Routings

As discussed in Section IV, we can repeatedly apply lightpath rerouting to any initial lightpath routing to obtain a reliable routing. Next, we investigate the performance of the rerouting algorithm by applying to two different initial lightpath routings generated by MCF of [12], and Shortest Path, which routes each lightpath over its shortest physical path.

Figures 7(a) and 7(b) show the average MCLC and reliability values of the two sets of lightpath routings before and after the repeated rerouting steps. Initially, the lightpath routings generated by Shortest Path have significantly lower MCLC

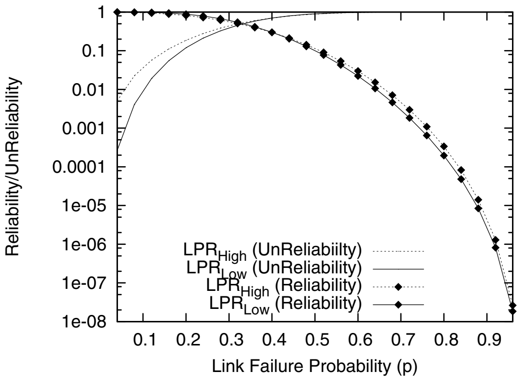

Fig. 6. Reliability (or Unreliability) of Lightpath Routings optimized for different probability regimes. Unreliability is defined to be (1 - Reliability).

and reliability than the ones generated by MCF. However, the lightpath rerouting algorithm is able to improve both sets of lightpath routings to similar MCLC and reliability values. This illustrates the robustness of the lightpath rerouting approach with respect to the initial choice of lightpath routing. Furthermore, it raises the average MCLC of the initial lightpath routings to almost 4, which is the connectivity of the logical topologies and is therefore an upper bound on the MCLC value. In other words, in terms of MCLC, the rerouting algorithm provides near-optimal performance regardless of the initial routings.

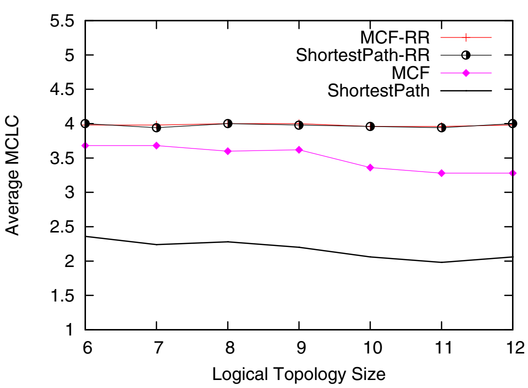

(a) Average MCLC

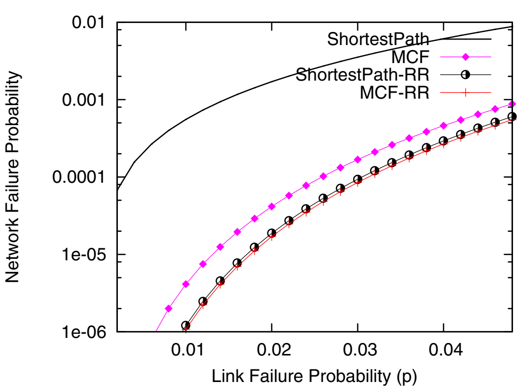

(b) Network Failure Probability

Fig. 7. Lightpath Rerouting with Different Initial Lightpath Routings

VI. Conclusion

We studied lightpath routing algorithms for reliability maximization in WDM networks. The key to this study is the polynomial expression for reliability which relates structural properties of the network graph and the lightpath routing to the reliability. Using this polynomial, we showed that reliable routings depend on the link failure probability, and identified optimality conditions for reliability maximization in different failure probability regimes. In particular, we showed that a routing with the maximum size of Min Cross Layer Cuts (MCLC) and the minimum number of MCLCs is most reliable in the low failure probability regime. On the other hand, in the high failure probability regime, a routing with the minimum size of Min Cross Layer Spanning Tree (MCLST) and the maximum number of MCLSTs maximizes reliability. Using these results and a novel rerouting technique, we developed lightpath routing algorithms that can maximize reliability in the desired probability regime. We demonstrated the performance of our algorithms through extensive simulations.

References

[1] J. Zhang and B. Mukherjee, “A review of fault management in wdm mesh networks: Basic concepts and research challenges,” *IEEE Network*, pp. 41–48, Mar 2004.

[2] M. Tornatore, G. Maier, and A. Pattavina, “Availability design of optical transport networks,” *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 8, pp. 1520–1532, Aug. 2005.

[3] J. Segovia, E. Calle, P. Vila, J. Marzo, and J. Tapolcai, “Topology-focused availability analysis of basic protection schemes in optical transport networks,” *Journal of Optical Networking*, vol. 7, pp. 351–364, Mar. 2008.

[4] M. Kurant, H. X. Nguyen, and P. Thiran, “Survey on dependable ip over fiber networks,” *Lecture Notes in Computer Science*, vol. 4028, pp. 55–81, Nov. 2006.

[5] J. Armitage, O. Crochat, and J.-Y. L. Boudec, “Design of a survivable WDM photonic network,” in *IEEE INFOCOM*, Apr. 1997.

[6] O. Crochat and J.-Y. L. Boudec, “Design protection for WDM optical networks,” *IEEE JSAC*, vol. 16, no. 7, Sep. 1998.

[7] Q. Deng, G. Sasaki, and C.-F. Su, “Survivable IP over WDM: a mathematical programming problem formulation,” in *Proc 40th Annual Allerton Conference*, Oct. 2002.

[8] 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.

[9] E. Modiano and A. Narula-Tam, “Survivable lightpath routing - a new approach to the design of wdm-based networks,” *IEEE JSAC*, vol. 20, no. 4, pp. 800–809, May 2002.

[10] A. Sen, B. Hao, and B. H. Shen, “Minimum cost ring survivability in WDM networks,” in *Workshop on High Performance Switching and Routing*, 2003, pp. 183– 188.

[11] A. Todimala and B. Ramamurthy, “Survivable virtual topology routing under Shared Risk Link Groups in WDM networks,” in *BROADNETS ’04*, Washington, DC, USA, 2004.

[12] K. Lee and E. Modiano, “Cross-layer survivability in WDM-based networks,” in *IEEE INFOCOM*, Rio de Janeiro, Brazil, Apr. 2009.

[13] K. Lee, H.-W. Lee, and E. Modiano, “Reliability in layered networks with random link failures,” in *IEEE INFOCOM*, Mar 2010.

[14] M. Chari and C. Colbourn, “Reliability polynomials: A survey,” *Journal of Combinatorics, Information and System Sciences*, vol. 22, pp. 177–193, 1997.

[15] W. Myrvold, “Reliable network synthesis: Some recent developments,” in *International Conference on Graph Theory, Combinatorics, Algorithms, and Applications, volume II*, 1996, pp. 650–660.

[16] H.-W. Lee, K. Lee, and E. Modiano, “Maximizing reliability in wdm networks through lightpath routing,” MIT, Tech. Rep., 2011. [Online]. Available: http://web.mit.edu/hwlee/www/reliability.pdf