Dual-Link Failure Resiliency Through Backup Link Mutual Exclusion

Amit Chandak and Srinivasan Ramasubramanian
Department of Electrical and Computer Engineering
University of Arizona, Tucson, AZ 85721
E-mail: {chandak, srini}@ece.arizona.edu

***Abstract*— One of the strategies to recover from dual-link failures is to employ link protection for the two failed links independently which requires that two links may not use each other in their backup paths if they may fail simultaneously. Such a requirement is referred to as Backup Link Mutual Exclusion (BLME) constraint and the problem of identifying a backup path for every link that satisfies the above requirement is referred to as the BLME problem.**

This paper explores the BLME problem in depth by: (1) formulating the backup path selection as an integer linear program; and (2) developing a pseudo-polynomial time approximation algorithm based on minimum cost path routing. The ILP formulation and heuristic are applied to six networks and their performance is compared to approaches that assume precise knowledge of dual-link failure. The heuristic approach is shown to obtain feasible solutions that are resilient to most dual-link failures, although the backup path lengths may be significantly higher than optimal. In addition, the paper demonstrates the significance of the knowledge of failure location by illustrating that network with higher connectivity may require lesser capacity than one with a lower connectivity to recover from arbitrary dual-link failures.

I. Introduction

Optical networks of today operate in a circuit-switched manner as optical header processing and buffering technologies are still in the early stages of research for wide-scale commercial deployment. Protecting the circuits or connections established in such networks against single-link failures may be achieved in two ways: *path protection* or *link protection*. Path protection attempts to restore a connection on an end-to-end basis by providing a backup path in case the primary (or working) path fails. The backup path assignment may be either independent or dependent on the link failure. For example, a backup path that is link-disjoint with the primary path allows recovery from single-link failures without the precise knowledge of failure location. On the other hand, more than one backup path may be assigned for a primary path and the connection is reconfigured on the backup path corresponding to the failure scenario that resulted in the primary path failure. The former is referred to as failure independent path protection (FIPP) while the latter is referred to as failure dependent path protection (FDPP).

Link protection recovers from a single-link failure by rerouting connections around the failed link. Such a recovery may be achieved transparent to the source and destination of the connections passing through the failed link. Link protection at the granularity of a fiber switches all the connections on a fiber to a separate (spare) fiber on the backup path. The time needed to detect the fault, communicate to the end-nodes, re-initiate connection requests on the backup paths, and reconfigure the switches at the intermediate nodes could sometimes cause the layers above the optical layer to resort to their own restoration solutions. Link protection reduces the communication requirement compared to path protection, thus provides faster recovery. However, the downside of link protection is that the capacity requirement is higher than that of path protection.

Algorithms for protection against link failures have traditionally considered single-link failures [1], [2], [3] (for a detailed description on protection approaches, refer to [4]). However, dual-link failures are becoming increasingly important due to two reasons. Firstly, links in the networks share resources such as conduits or ducts and the failure of such shared resources result in the failure of multiple links. Secondly, the average repair time for a failed link is in the order of few hours to few days [4] and this repair time is sufficiently long for a second failure to occur. Although algorithms developed for single-link failure resiliency is shown to cover a good percentage of dual-link failures [5], [6], [7], [8], these cases often include links that are far away from each other. Considering the fact that these algorithms are not developed for dual-link failures, they may serve as an alternative to recover from independent dual-link failures. However, reliance on such approaches may not be preferable when the links close to one another in the network share resources, leading to correlated link failures.

Dual-link failures may be modeled as shared risk link group (SRLG) failures. A connection established in the network may be given a backup path under every possible SRLG failure. This approach assumes a precise knowledge of failure locations to reconfigure the failed connections on their backup paths. An alternative is to protect the connections using link protection, where only the nodes adjacent to the failed link (and those involved in the backup path of the link) will perform the recovery. The focus of this paper is to protect end-to-end connections from dual-link failures using link protection.

A. Dual-link failure resiliency using link protection

Assume that two links, ℓ_1 and ℓ_2, failed one after the other (even if they occur together, assume that one failed first followed by the other) in a network $\mathcal{G}$. The backup path of the first failed link is analogous to a connection (at the granularity of a fiber) established between two non-adjacent nodes in the network $\mathcal{G} - \{\ell_1\}$, which is required to be protected against a single-link failure in $\mathcal{G} - \{\ell_1\}$. Therefore, strategies developed for protecting connections against single-link failures may be directly applied for dual-link failures that employ link protection to recover from the first failure.

Dual-link failure resiliency strategies may be classified based on the nature in which the connections are recovered from first and second failures. The recovery from the first link failure is assumed to employ link protection strategy. Figure 1 shows an example network where link 1–2 is protected by the backup path 1–3–4–2 when failed. The second protection strategy will refer to the manner in which the backup path of the first failed link is recovered.

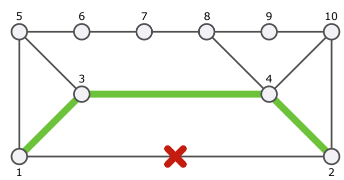

Fig. 1. Link 1–2 protected by backup path 1–3–4–2 when failed.

Link Protection – Failure Independent Protection (LP-FIP). One approach to dual-link failure resiliency using link protection is to compute two link-disjoint backup paths for every link. Given a three-edge-connected[1] network, there exists three link-disjoint paths between any two nodes [9]. Thus, for any two adjacent nodes, there exists two link-disjoint backup paths apart from the link connecting the two nodes. Let B_ℓ and B'_ℓ denote the two link-disjoint backups for link ℓ. Without loss of generality, assume that on failure of link ℓ, a connection routed along link ℓ will be re-routed on B_ℓ. If the first failure in the network happens to be one of the links belonging to B_ℓ, a subsequent failure of link ℓ will necessitate switching of connections to backup path B'_ℓ. Hence, the node connected to link ℓ must have some knowledge of the first failure (not necessarily the location).

The backup path under LP-FIP may be computed as 1–5–6–7–8–9–10–2. The backup path is identical under any second failure that affects the path 1–3–4–2. When the second failure occurs, a failure notification must be sent to node 1, although this notification need not explicitly mention which link failed in the path 1–3–4–2. Note that every link has two backup paths assigned, one for single-link failure scenario and another for dual-link failure scenario.

Link Protection - Failure Dependent Protection (LP-FDP). For every second failure that affects the backup path, a backup path under dual-link failure is provided. This backup path is obtained by eliminating the two failed links from the network and computing shortest path between the specific node pairs. Figure 2 shows the backup path assigned for link 1–2 under dual-link failures. The backup path assignment is different for different failures that affect the path. When a second link failure occurs, a failure notification must be sent to node 1, indicating the precise failure location on the path 1–3–4–2. It is fairly straight forward to see that the average backup path length under dual-link failures using LP-FDP will be lesser than that using LP-FIP. Every link is assigned one backup path for single-link failure and multiple backup paths (depending on the number of links in the backup path for the single-link failure) under dual-link failures.

Link Protection - Link Protection (LP-LP). Notification of the second failure location to different nodes for them to reconfigure their backup paths may result in a high recovery time. In order to avoid notification to other nodes and reconfiguring at the end of the paths, link protection may be adopted to recover from the second link failure as well. Figure 3 shows how the backup path 1–3–4–2 is reconfigured after possible second link failures. Under this strategy, every link will have only one backup path (for all failure scenarios). In order for this strategy to work, the backup path under the second failure must not pass through the first failed link! This constraint is referred to as the *Backup Link Mutual Exclusion* (BLME) constraint. Note that the above approach would fail if the backup path for link 3–4 is 3–1–2–4 for the example in Figure 3.

Backup Link Mutual Exclusion (BLME) Problem. Given a network $\mathcal{G}$ and a set of dual-link failures $\mathcal{F}$, where each element $f \in \mathcal{F}$ contains at most two links that can fail and that the failure does not disconnect the network. Then, for every link ℓ in the network, identify a backup path $\mathcal{P}_\ell$, such that if a link ℓ' is in $\mathcal{P}_\ell$ and both links ℓ and ℓ' can fail together, then the backup path of link ℓ' does not use link ℓ.

The problem is concisely described as: Given $\mathcal{G}$ and $\mathcal{F}$, identify $\mathcal{P}_\ell \;\; \forall \ell \in \mathcal{G}$ such that $\forall \ell' \neq \ell$:

$$[\ell' \in \mathcal{P}_\ell] \wedge [\exists f \in \mathcal{F} : (\ell \in f) \wedge (\ell' \in f)] \quad \Rightarrow \quad \ell \notin \mathcal{P}_{\ell'}$$

Background and prior work. A network must be three-connected for it to be resilient to any two arbitrary link failures, irrespective of the protection strategy employed. The only related work in this context is [10] that attempts to provide a heuristic solution to the BLME problem. The heuristic solution, referred to as the Maximum Arbitrary Double-Link Protection Algorithm (MADPA), is shown to find backup paths satisfying the BLME constraint for all links in two networks out of the three networks considered in their paper, however it is not guaranteed to find a solution even if one exists. It may be shown that three edge-connectivity is a sufficient condition for the existence of a solution to the BLME problem, the proof of which is omitted in this paper due to space constraints. The readers are referred to the technical report [11], which to the best of our knowledge is the first and only known proof that three-edge-connectivity is a sufficient condition for the existence of a solution.

[1] A k-edge-connected network is simply referred to as a k-connected network in the rest of this paper.

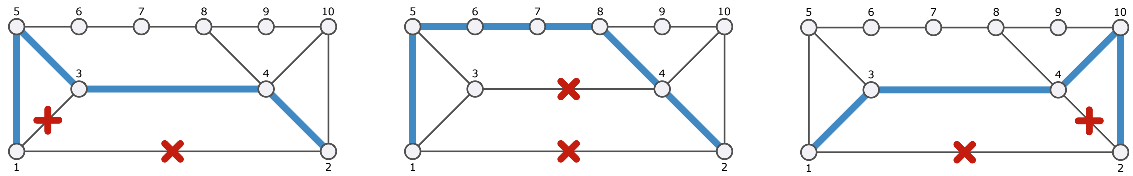

(a) (b) (c)

Fig. 2. Dual-link failure resiliency using LP-FDP.

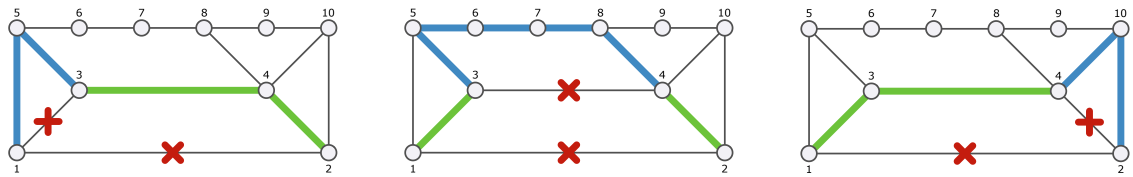

(a) (b) (c)

Fig. 3. Dual-link failure resiliency using LP-LP.

B. Scope, contribution, and organization

This paper assumes that link protection is performed at the granularity of a fiber, where every link is equipped with one primary fiber and one or two spare fibers in each direction as necessary. Recovery from a failure is performed using fiber switching, hence issues relating to specific connections flowing through a link (such as wavelength continuity constraint) do not arise. As the concept of dual-link failure resiliency using only link protection has not received much attention in the literature, the scope of the paper is restricted to obtaining solution to the BLME problem with backup path length as the only optimization metric. Considerations for actual traffic carried on a link and recovery at the granularity of a connection are beyond the scope of this paper.

Solution methodologies to the BLME problem are developed using two approaches by: (1) formulating the BLME problem as an integer linear program (ILP); and (2) developing a pseudo-polynomial time approximation algorithm. The trade-offs involved in solutions that have and do not have the precise knowledge of the failure location are compared by applying the solution techniques to six networks. In addition, the paper also establishes the potential benefits of a four-connected network over a three-connected network when the knowledge of the failure location is available.

The rest of the paper is organized as follows: Section II describes the formulation of the BLME problem as an ILP with some specific insights into the objective function and formulation for networks that may not be three-connected. Section III develops a pseudo-polynomial time approximation algorithm, called *Iterative Minimum Cost Path* heuristic. Section IV discusses the possibility of loop formation under the BLME approach. Section V describes the results obtained by applying the ILP formulation and heuristic to six networks. Section VI shows the advantages of designing a four-connected network over a three-connected network given that the location of two failures are known. Conclusions and future work are discussed in Section VII.

II. BLME Formulation

Consider a network represented as a graph $\mathcal{G}(\mathcal{N}, \mathcal{L})$, where $\mathcal{N}$ and $\mathcal{L}$ denote a set of nodes and undirected links, respectively. A link $\ell \in \mathcal{L}$ is assumed to be bi-directional. Each link is assumed to employ one primary fiber in each direction.

Let $\mathcal{F}$ denote the set of dual-link failures to be tolerated. An element $f \in \mathcal{F}$ consists of exactly two undirected links.

The BLME problem is formulated as an Integer Linear Program (ILP) using undirected links. The central idea behind this ILP formulation is to view the network as $|\mathcal{L}|$ distinct graphs. Each graph, denoted as $\mathcal{G}_\ell$, will provide a backup path for link ℓ. Equivalently, each graph $\mathcal{G}_\ell$ will have a ring with link ℓ present in it.

Let $F_{\ell\ell'}$ denote the existence of a failure $f \in \mathcal{F}$ such that $\ell \in f$ and $\ell' \in f$; 1 if true, 0 otherwise.

Let $\alpha_{\ell\ell'}$ be a binary variable that indicates whether link ℓ' is present in graph $\mathcal{G}_\ell$: 1 if present, 0 otherwise. Similarly, let $\beta_{\ell i}$ be a binary variable that indicates whether node i is present in graph $\mathcal{G}_\ell$ or not: set to 1 if present, 0 otherwise. Let $C_{\ell i}$ denote whether node i is attached to link ℓ or not: 1 if true, 0 otherwise.

The formulation of backup path selection for all links satisfying BLME constraint is shown in Figure 4. The objective function is set to minimize the sum of the backup path lengths of all links, or equivalently the average backup path length of a link under a single-link failure. The average backup path length under single-link failure, denoted by $\overline{H}$, is computed as:

$$\overline{H} = \left(\frac{1}{|\mathcal{L}|} \sum_{\ell, \ell' \in \mathcal{L}} \alpha_{\ell\ell'} \right) - 1 \tag{1}$$

The constraint *GC* ensures that a graph $\mathcal{G}_\ell$ must contain a ring with link ℓ present in it by forcing the corresponding link variable to take a value of 1. The *BLME* constraint ensures that for two links ℓ and ℓ', if link ℓ' is present in $\mathcal{G}_\ell$, then ℓ is not present in graph $\mathcal{G}_{\ell'}$ provided the two links ℓ and ℓ' may

Objective function.
Minimize $\sum_{\ell \in \mathcal{L}} \sum_{\ell' \in \mathcal{L}} \alpha_{\ell\ell'}$

Graph constraint (GC).
$\alpha_{\ell\ell} = 1 \qquad \forall \ell \in \mathcal{L}$

BLME constraint (BLMEC).
$\alpha_{\ell\ell'} + \alpha_{\ell'\ell} + F_{\ell\ell'} \leq 2 \qquad \forall \ell, \ell' \in \mathcal{L}$ and $\ell \neq \ell'$

Ring constraint (RC).
$2\beta_{\ell i} - \sum_{\ell' \in \mathcal{L}} \alpha_{\ell\ell'} C_{\ell' i} = 0 \qquad \forall \ell \in \mathcal{L}$ and $i \in \mathcal{N}$

Bounds.
$\alpha_{\ell\ell'} = \{0, 1\} \qquad \forall \ell, \ell' \in \mathcal{L}$
$\beta_{\ell i} = \{0, 1\} \qquad \forall \ell \in \mathcal{L}$ and $i \in \mathcal{N}$

Fig. 4. ILP formulation of backup path selection with BLME constraint.

be unavailable at the same time. Otherwise, such a restriction is not imposed.

The constraint *RC* ensures every graph has a ring, every node i that is present in a graph must have exactly two outgoing (or incoming) links. The above constraint introduces $|\mathcal{L}| \times |\mathcal{N}|$ additional variables $\beta_{\ell i}$ to the formulation, however, they are strongly correlated to the link variables $\alpha_{\ell\ell'}$.

The variables employed in the formulation are limited to take binary values using *Bounds*.

For a network to be resilient to arbitrary dual-link failures, the network must be three-connected. However, some real-life networks may not be three-connected. For example, in the NJ-LATA network of Figure 7, three dual-link failures involving link pairs 1 and 2; 11 and 16; and 22 and 23 will disconnect the network. If a dual-link failure involving two links ℓ and ℓ' disconnects the network, it is not possible to find backups for links ℓ and ℓ' satisfying the BLME constraint as they will have to include each other. The BLME constraint is relaxed for these failure scenarios by not considering these dual-link failures. In other words, $F_{\ell\ell'}$ is assigned 0 (assumed not to fail) even though the links can be unavailable at the same time, if their joint failure will disconnect the network.

III. A Heuristic Approach

As ILP solution times for large networks may be prohibitively high, a heuristic approach is also developed. The heuristic solution is based on iterative computation of minimum cost routing. The network is treated as an undirected graph $\mathcal{G}$. A set of auxiliary graphs corresponding to failure of a link $\ell \in \mathcal{G}$ is created: $\mathcal{X}_\ell = \mathcal{G}(\mathcal{N}, \mathcal{L} - \{\ell\})$. In each auxiliary graph $\mathcal{X}_\ell$, the objective is to obtain a path between the nodes that were originally connected by link ℓ. Let $\mathcal{P}_\ell$ denote the path selected in auxiliary graph $\mathcal{X}_\ell$. If a link ℓ' is a part of the path selected on graph $\mathcal{X}_\ell$, then the path in graph $\mathcal{X}_{\ell'}$ must avoid the use of link ℓ. This is accomplished by imposing a cost on the links in the auxiliary graphs and selecting the minimum cost path. Let $w_{\ell,\ell'}$ denote the cost of link ℓ' on graph $\mathcal{X}_\ell$ such that it indicates that graph $\mathcal{X}_{\ell'}$ contains link ℓ and the two links ℓ and ℓ' may be unavailable simultaneously. Hence, the cost values are binary in nature.

The cost of a path in an auxiliary graph is the sum of the cost of links in it. At any given instant during the computation, the total cost of all the paths (T) is the sum of the cost of the paths across all auxiliary graphs. It may be observed that the total cost must be an even number, as every link ℓ' in a path $\mathcal{P}_\ell$ that has a cost of 1 implies that link ℓ in path $\mathcal{P}_{\ell'}$ would also have a cost of 1. For a given network, the minimum value of the total cost would then be two times the number of dual-link failure scenarios that would disconnect the network. If τ denotes the number of dual-link failure scenarios that would disconnect the graph, then the termination condition for the heuristic is given by $T = 2\tau$. Figure 5 shows the steps involved in the Iterative Minimum Cost Path (IMCP) heuristic.

The complexity of an iteration of the IMCP heuristic is dictated by the (backup) path selection step (Step 4.2.), whose complexity is $\mathcal{O}(|\mathcal{L}| + |\mathcal{N}| \ log \ |\mathcal{N}|)$. The complexity of an iteration of IMCP (decided by Step 4) is $\mathcal{O}\left(|\mathcal{L}|^2 + |\mathcal{L}||\mathcal{N}| \ log \ |\mathcal{N}|\right)$. The number of iterations required to obtain a solution depends on the connectivity of the network. Although there is a terminating condition specified for this algorithm, the algorithm is not guaranteed to terminate for certain networks. Hence, one can limit the number of iterations to be performed to a fixed threshold. If the algorithm is restricted to a maximum of K iterations, the IMCP heuristic is pseudo-polynomial with a worst-case complexity of $\mathcal{O}\left(K|\mathcal{L}|^2 + K|\mathcal{L}||\mathcal{N}| \ log \ |\mathcal{N}|\right)$.

IV. Loop Formation

Recall that the backup path of a link ℓ after its failure is analogous to a connection established in a network which is protected using link protection (for the second failure). Hence, all the properties of a link protection strategy for a connection in a regular network is valid in dual-link failure resiliency using link protection. Loop formation is one among them!

Consider a failed link (connecting nodes 1 and 8) whose backup is established along a path where the nodes in the path are numbered from 1 through 8. Figure 6 shows two kinds of loop formation.

In Figure 6(a), link 4–5 is protected by the backup path 4–7–6–5. Upon failure of link 4–5, the backup between nodes 1 and 8 is modified as 1–2–3–4–7–6–5–6–7–8, resulting in the loop 7–6–5–6–7. While this loop could be pruned using signaling, it is only necessary to reduce path delay. The backup path for 4–5 will route both its primary fiber and the secondary fiber of link 1–8 through the path 4–7–6–5, hence two spare fibers will be used in each of these links. The loop formation involves the same links, however the links are traversed in opposite direction. As every link needs to be equipped with two spare fibers in each direction (for bi-directional connectivity), there will not be a resource contention.

Figure 6(b) shows another kind of loop formation where there is a potential resource contention. The backup path of

Iterative Minimum Cost Path (IMCP) Heuristic

Step 1. Obtain auxiliary graphs $\mathcal{X}_\ell$ for every $\ell \in \mathcal{X}$ as $\mathcal{X}_\ell = \mathcal{X} - \{\ell\}$. Note that every link $\ell \in \mathcal{X}$ is assumed to be bi-directional in nature.
Step 2. Initialize the path to be found in every graph $\mathcal{X}_\ell$ as an empty set. $\mathcal{P}_\ell \leftarrow \phi \quad \forall \ell \in \mathcal{X}$.
Step 3. Initialize the cost of all the links in every auxiliary graph to 0. $w_{\ell,\ell'} \leftarrow 0, \quad \forall \ell \in \mathcal{X}, \quad \ell' \in \mathcal{X}_\ell$.
Step 4. For every auxiliary graph $\mathcal{X}_\ell$

1) Erase the old path and update the cost in auxiliary graphs; i.e., for every link $\ell' \in \mathcal{P}_\ell$, update $w_{\ell',\ell} \leftarrow 0$. $\mathcal{P}_\ell \leftarrow \phi$.
2) Re-compute the least cost path $\mathcal{P}_\ell$.
3) If a link ℓ' is present in this graph, then increase the cost of link ℓ in auxiliary graph $\mathcal{X}_{\ell'}$; i.e. For every link $\ell' \in \mathcal{P}_\ell$, update $w_{\ell',\ell} \leftarrow F_{\ell,\ell'}$.

Step 5. Compute the total cost of all paths over all the auxiliary graphs; i.e. $T = \sum_{\ell \in \mathcal{X}} \sum_{\ell' \in \mathcal{P}_\ell} w_{\ell,\ell'}$.
Step 6. If the total cost of all the paths equals the threshold of 2τ, where τ is the number of dual-link failure scenarios that would disconnect the graph, then it indicates the best possible solution has been obtained; i.e., if $T = 2\tau$, go to Step 7. Otherwise, go to Step 4.
Step 7. Stop.

Fig. 5. Steps involved in the Iterative Minimum Cost Path (IMCP) Heuristic solution.

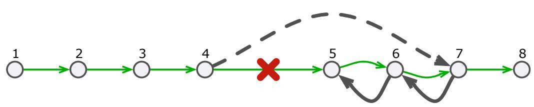

(a) No resource contention

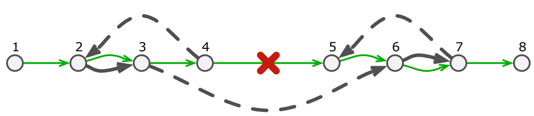

(b) Contention on links traversed in same direction

Fig. 6. Illustration of loop formation in dual-link failure resiliency with BLME. (a) Looping with a link traversed in opposite directions. (b) Looping with links 2–3 and 6–7 are traversed twice in the same direction; thus requiring pruning.

link 4–5 is 4–2–3–6–7–5. Upon failure of link 4–5, the backup between nodes 1 and 8 is modified as: 1–2–3–4–2–3–6–7–5–6–7–8; resulting in two loops 2–3–4–2 and 7–5–6–7. Note that links 2–3 and 6–7 are traversed in the same direction. If such a loop formation is allowed, then links 2–3 and 6–7 must be equipped with three spare fibers as the backup path for link 4–5 would be switching two fibers. As the network is assumed to have at most two link failures, it must be sufficient to equip every link with two spare fibers. Therefore, if the links have only two spare fibers, pruning of the backup paths cannot be avoided. The pruned backup path between nodes 1 and 8 after link 4–5 failure would be 1–2–3–6–7–8; while the backup path of link 4–5 would be 4–2–3–6–7–5.

The looping problem described here is similar to those encountered in any link protection mechanism. The paths can be pruned using a signaling mechanism that would be required to establish the backup paths. Note that the signaling cannot be completely avoided as a link can serve as backup for more than two other links, hence protection switches cannot be configured prior to failure.

V. Performance Evaluation

The performance of the ILP and heuristic algorithm developed in this paper are evaluated by applying them to six networks as shown in Figure 7: (a) ARPANET; (b) NSFNET[2]; (c) Node-16; (d) Node-28; (e) Mesh-4×4; and (f) NJ-LATA. All networks except NJ-LATA are three-connected. The NJ-LATA network is not three-connected as nodes 1, 6, and 11 have degree 2 and is considered "as is" for performance evaluation. The Node-16 and Node-28 networks are hypothetical networks used to test the limits of the ILP. All the nodes in these two networks have exactly three links connected to them, thus these two networks are *minimally three-connected*.

Dual-link failure scenarios occur in network due to two reasons as mentioned in Section I. Firstly, link resources such as conduit or duct are shared by multiple links for ease of layout. Such sharing of resources is typically limited to links that are close to each other, such as adjacent links. Hence, dual-link failure scenarios under such shared resource failure typically affect only near-by links. The second case of dual-link failure scenario is due to the time required to repair a failed link. Before a failed link is repaired, another link in the network could fail. The possibilities of such a link failure are typically rare. If it can be assumed that most of the dual-link failures may be caused because of failure of shared resources, then it is of interest to identify backup paths assignments by considering only failure of near-by links.

Two kinds of dual-link failures are considered for performance evaluation: (1) arbitrary dual-link failures; and (2) adjacent dual-link failures. Two links that are attached to a node are referred to as adjacent links. Note that any dual-link failure that will disconnect the network is not considered in computing the number of failures that can be tolerated.

A. *Performance Metrics*

The performance metrics considered specifically for the ILP solutions are: (1) solution time and (2) optimality bound. The optimality bound is relevant in scenarios where the ILP could not obtain optimal solution, but has a feasible solution with a known bound on optimality. The ILP is solved using the CPLEX 8.1 solver [12] on a single-processor Pentium-4 2.53

[2]The NSFNET network considered here has been modified from the original network with the addition of link numbered 23 to keep the network three-connected.

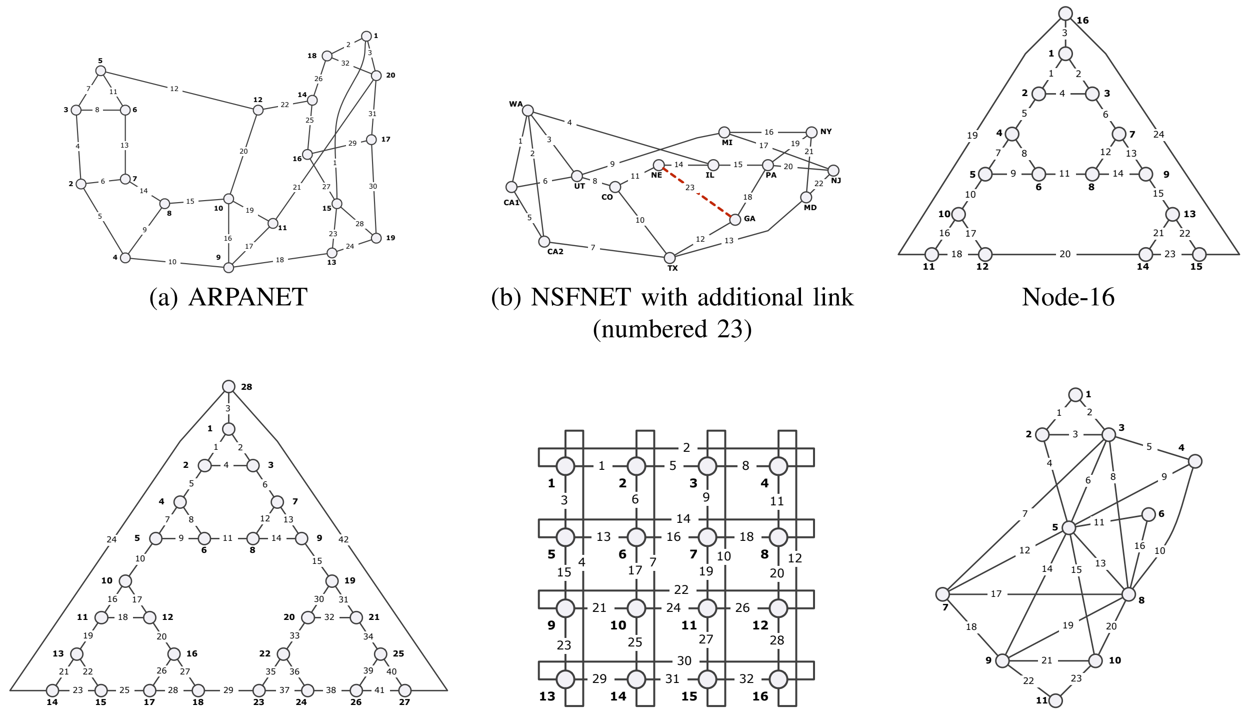

(a) ARPANET (b) NSFNET with additional link (numbered 23) Node-16

(d) Node-28 (e) Mesh-4×4 (f) NJ-LATA

Fig. 7. Networks considered for performance evaluation.

GHz computer with 512 MB of RDRAM. The metric that is considered specifically for heuristic is the number of dual-link failures that can be tolerated, as the heuristic is not guaranteed to recover from all dual-link failures.

In addition to the above metrics that are specific to the ILP and heuristic, certain common metrics for both the approaches are also evaluated: (1) average and maximum backup path length for a single-link failure; (2) average and maximum backup path length under dual-link failure scenario; and (3) total spare capacity required.

The backup path length under a single-link failure scenario is simply the length of the backup path when the only failure in the network is the particular link itself. The backup path of a link under dual-link failure scenario is computed as: (1) the backup path length of the link alone if the backup does not use the second link; (2) it is computed after pruning if the backup employs the second failed link. The average and maximum backup path length under dual-link failures is computed over all dual-link failures in the network.

The spare capacity requirement is computed as discussed in [10]:

- If a link ℓ_1 is not in the backup of any other link, i.e., $\alpha_{\ell_2,\ell_1} = 0 \ \ \forall \ell_2$, then no additional capacity is required.
- A link will require 200% additional capacity under two conditions: (1) link ℓ_1 is in the backup path for two links ℓ_2 and ℓ_3 that could be unavailable simultaneously, mathematically represented as $\sum_{\ell_2,\ell_3 \in \mathcal{L}'; \ell_1 \neq \ell_2 \neq \ell_3} \alpha_{\ell_2,\ell_1} \alpha_{\ell_3,\ell_1} F_{\ell_2,\ell_3} > 0$; or (2) link ℓ_1 is in the backup of ℓ_2, ℓ_2 is in the backup of ℓ_3, and ℓ_2 and ℓ_3 can both be unavailable together, mathematically represented as $\sum_{\ell_2,\ell_3 \in \mathcal{L}'; \ell_1 \neq \ell_2 \neq \ell_3} \alpha_{\ell_3,\ell_2} \alpha_{\ell_2,\ell_1} F_{\ell_2,\ell_3} > 0$.
- In all other cases, the link requires one additional fiber, a 100% additional capacity.

ILP results. Table I shows the results for the six networks to be resilient to arbitrary dual-link failures obtained using ILP with the objective to optimize the average backup path length under single-link failure scenario. The CPLEX program terminated due to insufficient memory for Node-28, and Mesh-4×4 networks. While this is indicative of the complexity of the problem, feasible solutions were obtained as intermediate values. The best value obtained before termination is reported for these networks. It is observed that the solution time increases significantly with increase in the network size. For scenarios where an optimal solution is not found, the value of optimality bound indicates the worst-case deviation of best value from the optimal. It is to be noted that although an optimal solution may not be obtained, a feasible solution is obtained for all the networks considered, confirming the existence of a solution.

It is observed that a 200% additional capacity (two spare fibers) is required in all the links of all the networks except NJ-LATA. Such a requirement can be immediately deduced from the connectivity of the network. For example, whenever a link is necessary[3] to keep the network three-connected, then such a link must have two spare fibers. Thus, the networks Node-16 and Node-28 will require 200% additional capacity even when only adjacent links may fail together. Such a 200%

[3]Removal of the link will result in the network not being three-connected.

requirement in capacity may be reduced only on those links whose removal does not affect the three-connectivity property of the network. For example, the link between WA and UT in the NSFNET may be removed without affecting the three-connectivity property of the network. However, such a solution would have an increased average backup path length. For NJ-LATA network, two of the three link pairs whose failure disconnects the network were not present in more than one backup path, hence a reduction of 4 fibers was obtained. The high connectivity in the NJ-LATA network results in this reduction even when the objective function is not set to minimize capacity, which is purely coincidental. Such a reduction cannot be guaranteed for all networks.

It is also observed that the average backup path length under dual-link failures (for Node-16 and Node-28 networks) may be lower than the average backup path length under single-link failure scenarios due to path pruning.

Table II shows the results obtained from ILP considering only adjacent dual-link failures. The number of adjacent dual-link failures is significantly less than the number of arbitrary dual-link failures, thus leading to a significant improvement in solution time and average hop length. The CPLEX program terminated due to insufficient memory for the Node-28 network. The best integer solution obtained at the time of termination is reported in the table. The backup path length under a single-link failure scenario is considerably reduced when compared to the results for any arbitrary link failures (10.191 to 6.786 for Node-28).

The backup path length under dual-link failures decreases for adjacent failures compared to arbitrary failures, and so does the number of dual-link failures in the network. If the reduction in the backup path length dominates the reduction in the number of dual-link failures, then the average backup path length for adjacent link failures will be lesser than that of arbitrary link failures, otherwise it will be higher. The former is the case for Node-16 and Node-28 networks but latter is the case for NSFNET, ARPANET, NJ-LATA networks. Mesh-4$\times$4 remains unaffected.

Heuristic results. Tables III and IV show the results obtained using the IMCP heuristic approach for tolerating arbitrary dual-link failures and adjacent dual-link failures, respectively. Interestingly, this simple heuristic obtains solution that can tolerate most dual-link failures for the networks considered; the solution cannot recover from four dual-link failures for Node-28 under arbitrary dual-link failures.

The heuristic produces a solution in relatively less number of iterations for five of the six scenarios. A maximum of thirty iterations were performed. The running time of the IMCP heuristic for obtaining a solution under arbitrary dual-link failures was found to be less than 20 ms for all networks except the Node-28 network for which the algorithm ran for 170 ms to complete thirty iterations. The running time for the IMCP heuristic for obtaining a solution under adjacent dual-link failures was observed to be less than 10 ms for all the networks considered.

While the objective of the heuristic is to obtain a feasible solution, it is not guaranteed to find a solution (as seen in the Node-28 network for arbitrary dual-link failures). The results presented in this paper are obtained by iterating over a given ordered set of links. In our study where we considered iterating the auxiliary graphs in the reverse order, the heuristic could not find a feasible solution for Node-16 and Node-28 networks under arbitrary dual-link failures. The number of iterations required to arrive at the solution depends on a lot of parameters, specifically the order in which the auxiliary graphs are considered and the weights employed. Comparing the results of the heuristic to that of the ILP, it is observed that the heuristic can be as worse as 60% above optimal for average backup path lengths.

Comparing to the MADPA heuristic in [10], [13], the IMCP heuristic obtains a solution that recovers from all dual-link failures for ARPANET while MADPA tolerates only 490 (out of 496) unordered dual-link failures. Comparison to average path length obtained by MADPA is not performed as the authors of the work indicated that the results presented in [10], [13] were obtained before pruning.

Comparison to LP-FDP. Tables III and IV also compare the heuristic results (LP-LP) to the LP-FDP approach which assigns failure dependent backup paths for every link. The backup paths for a link ℓ for the LP-FDP scheme under single and dual-link failures are obtained by simply eliminating the failed link(s) and computing the shortest path between the nodes connected by link ℓ. Firstly, LP-FDP recovers from all dual-link failures that do not disconnect the network. It is observed that the knowledge of precise failure locations significantly reduces the backup path length under dual-link failure scenarios. The optimal (or best) values of the backup path lengths obtained by optimizing the average backup path length using ILP are found to be as high as twice the backup path lengths obtained under LP-FDP. The increased backup path length for LP-LP may also result in an increased recovery time from the second failure compared to LP-FDP.

While the advantages of the LP-LP strategy are that every link will have only one backup path and the failure location need not be explicitly notified to nodes that are not present in the backup path, the latter becomes a significant drawback in networks with very high connectivity, specifically when optimizing the total spare capacity.

VI. Capacity Optimization for Dual-link Failure Resiliency in Highly Connected Networks

The ILP developed in this paper has an objective of reducing the average path length while the IMCP heuristic attempts to just find any feasible solution. Often network designers are interested in minimizing spare capacity. The ILP developed in this paper could certainly be modified to minimize total spare capacity under the BLME approach. However, a direct way to minimize the total spare capacity for arbitrary dual-link failures would be to reduce the given network to a *minimally three-connected* network and employ two spare fibers per link. The ILP or IMCP heuristic may be employed on the minimally three-connected network to obtain backup paths.

TABLE I

ILP results for tolerating arbitrary dual-link failures.

Metric	ARPANET	NSFNET	Node-16	Node-28	Mesh-4x4	NJ-LATA
Solution time in seconds (Optimality bound)	1655.55 (Optimal)	63.16 (Optimal)	37529.08 (Optimal)	8740.42 (45.13%)	26450.83 (8.8%)	1.59 (Optimal)
Average backup path length (single failure)	4.719	4.087	6.125	10.191	3.5	2.130
Maximum backup path length	9	6	12	23	5	3
Average backup path length (dual failure)	4.816	4.314	6.069	9.256	3.667	2.192
Maximum backup path length	13	10	13	22	9	4
Total spare capacity	64	46	48	84	64	42

TABLE II

ILP results for tolerating adjacent dual-link failures.

Metric	ARPANET	NSFNET	Node-16	Node-28	Mesh-4x4	NJ-LATA
Solution time in seconds (Optimality bound)	18.88 (Optimal)	1.63 (Optimal)	10.05 (Optimal)	4494.16 (23.4%)	0.17 (Optimal)	1.16 (Optimal)
Average backup path length (single failure)	4.125	3.696	4.917	6.786	3.00	2.130
Maximum backup path length	8	5	7	18	3	3
Average backup path length (dual failure)	5.158	4.532	6.021	7.565	3.667	2.309
Maximum backup path length	13	9	12	21	5	4
Total spare capacity	64	46	48	84	64	42

TABLE III

Performance comparison of IMCP heuristic and LP-FDP for tolerating arbitrary dual-link failures.

Metric	ARPANET		NSFNET		Node-16		Node-28		Mesh-4x4		NJ-LATA	
	IMCP	LP-FDP	IMCP	LP-FDP	IMCP	LP-FDP	IMCP	LP-FDP	IMCP	LP-FDP	IMCP	LP-FDP
Number of dual-link failures tolerated	496	496	253	253	276	276	857	861	496	496	250	250
Average backup path length (single failure)	6.375	3.094	4.870	3.043	7.708	3.125	8.595	3.500	4.750	3.000	3.609	2.000
Maximum backup path length (single failure)	12	6	9	5	13	6	19	8	9	3	7	2
Average backup path length (dual failure)	6.375	3.212	4.986	3.140	7.232	3.406	8.944	3.718	4.815	3.00	3.486	2.016
Maximum backup path length (dual failure)	14	8	11	5	14	7	21	11	11	3	8	3
Total spare capacity	64	64	46	46	48	48	84	84	64	64	46	46

TABLE IV

Performance comparison of IMCP heuristic and LP-FDP for tolerating adjacent dual-link failures.

Metric	ARPANET		NSFNET		Node-16		Node-28		Mesh-4x4		NJ-LATA	
	IMCP	LP-FDP	IMCP	LP-FDP	IMCP	LP-FDP	IMCP	LP-FDP	IMCP	LP-FDP	IMCP	LP-FDP
Number of dual-link failures tolerated	72	72	54	54	48	48	84	84	96	96	92	92
Average backup path length (single failure)	6.094	3.094	4.609	3.043	6.333	3.125	8.119	3.500	4.688	3.000	3.043	2.000
Maximum backup path length (single failure)	11	6	7	5	12	6	16	8	7	3	5	2
Average backup path length (dual failure)	6.413	3.729	4.484	3.398	6.771	4.417	8.625	5.071	4.729	3.00	3.175	2.079
Maximum backup path length (dual failure)	13	8	10	5	12	7	18	9	9	3	7	3
Total spare capacity	64	64	46	46	48	48	84	84	64	64	46	46

The average backup path lengths will increase as the reduced network will have lower connectivity. The reason for the above argument is that the recovery from the second failure takes place without the precise knowledge of the first failure location. Hence, certain links are forced to employ 200% spare capacity. If a network is minimally three-connected, then every link must carry 200% spare capacity, irrespective of the strategy employed for dual-link failure resiliency. However, as the connectivity of the network is increased, it is possible to reduce the total spare capacity significantly given the precise knowledge of the dual-link failure.

Assume that every link is assigned a backup path for every dual-link failure (similar to LP-FDP) by eliminating the failed links and computing . Under a dual-link failure involving links ℓ_1 and ℓ_2, if the backup paths $\mathcal{B}_1$ (for ℓ_1) and $\mathcal{B}_2$ (for ℓ_2) are link-disjoint, then every link on these two backup paths will require only one spare fiber even under dual-link failure scenarios. If such a requirement is satisfied for all dual-link failures, then every link requires only one spare capacity. If the network is only three-connected, a few links could still carry only one spare fiber.

If a network is (even minimally) four-connected, then it is possible to find two link-disjoint backup paths for the two failed links, as shown by the following theorem, hence requiring only one spare fiber on all the links.

Theorem: Given a four-connected network $\mathcal{X}$ with a dual-link failure involving any two arbitrary links ℓ_1 (connecting nodes s_1 and d_1) and ℓ_2 (connecting s_2 and d_2), there exists backup paths $\mathcal{B}_1$ for link ℓ_1 and $\mathcal{B}_2$ for link ℓ_2 such that $\mathcal{B}_1$ and $\mathcal{B}_2$ are link-disjoint.

Proof. After two link failures, the resultant network $\mathcal{G}' = \mathcal{G} - \{\ell_1, \ell_2\}$ is either (1) at least three-connected, or (2) two-connected.

Case 1: If $\mathcal{G}'$ is three-connected, then by a variant of Menger's theorem [9], there exists three link-disjoint paths from a source x to three different destinations y_1, y_2, and y_3. Setting $x = s_1$, $y_1 = d_1$, $y_2 = s_2$, and $y_3 = d_2$, there exists three link-disjoint paths $\mathcal{P}(s_1, d_1)$, $\mathcal{P}(s_1, s_2)$, and $\mathcal{P}(s_1, d_2)$. By combining the paths $\mathcal{P}(s_1, s_2)$ and $\mathcal{P}(s_1, d_2)$, a path $\mathcal{P}(s_2, d_2)$ (through s_1) is obtained that is link-disjoint with $\mathcal{P}(s_1, d_1)$.

Case 2: If $\mathcal{G}'$ is two-connected, there exists a cut-set[4] in the graph as shown in the Figure 8.

Given that the original network is four-connected, there exists four link-disjoint paths between any two nodes. Therefore, there exists four link-disjoint paths from s_1 to d_1. The four link-disjoint paths must be of the form:

$$\begin{aligned} \mathcal{P}_1(s_1, d_1) &= \{\ell_1\} \\ \mathcal{P}_2(s_1, d_1) &= \mathcal{P}(s_1, s_2) \cup \{\ell_2\} \cup \mathcal{P}(d_2, d_1) \\ \mathcal{P}_3(s_1, d_1) &= \mathcal{P}(s_1, s_3) \cup \{\ell_3\} \cup \mathcal{P}(d_3, d_1) \\ \mathcal{P}_4(s_1, d_1) &= \mathcal{P}(s_1, s_4) \cup \{\ell_4\} \cup \mathcal{P}(d_4, d_1) \end{aligned}$$

[4]The figure depicts only the nodes and links that are of interest. Some nodes and links that lie within the sets to keep the network four-connected are not shown.

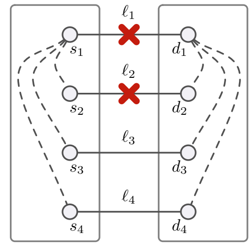

Fig. 8. A cut-set of a four-connected network where removal of links ℓ_1 and ℓ_2 results in a two-connected network. The solid lines denote a link between two nodes; while the dotted lines denote a path.

As the above four paths between s_1 and d_1 are link-disjoint, the paths $\mathcal{P}(s_1, s_2)$, $\mathcal{P}(s_1, s_3)$, and $\mathcal{P}(s_1, s_2)$ must also be mutually link-disjoint. By combining the paths $\mathcal{P}(s_1, s_2)$ and $\mathcal{P}(s_1, s_3)$, a path $\mathcal{P}(s_2, s_3)$ can be formed that is link disjoint with the path $\mathcal{P}(s_1, s_4)$. Using a similar argument, two link disjoint paths $\mathcal{P}(d_3, d_2)$ and $\mathcal{P}(d_4, d_1)$ may be obtained.

Note that the paths $\mathcal{P}(s_2, s_3)$ and $\mathcal{P}(s_1, s_4)$ are obviously link-disjoint with the paths $\mathcal{P}(d_3, d_2)$ and $\mathcal{P}(d_4, d_1)$ as they belong to two disjoint cut-sets. Using the above link-disjoint paths, backup paths for links ℓ_1 and ℓ_2 may be obtained as:

$$\begin{aligned} \mathcal{B}(s_1, d_1) &= \mathcal{P}(s_1, s_4) \cup \{\ell_4\} \cup \mathcal{P}(d_4, d_1) \\ \mathcal{B}(s_2, d_2) &= \mathcal{P}(s_2, s_3) \cup \{\ell_3\} \cup \mathcal{P}(d_3, d_2) \end{aligned}$$

The above two backup paths are link-disjoint as the individual path segments are mutually link-disjoint. **Q.E.D.**

Incorporating the link-disjoint path constraint for obtaining backup paths under dual-link failure scenarios will result in increased backup path lengths. However, the capacity savings obtained with the knowledge of failure locations could be significant. For example, consider a network with N nodes (assume N to be even) that can be either minimally three-connected or minimally four-connected. The minimum number of links required to form a minimally three-connected network is $1.5N$. These links must be equipped with one primary fiber and two spare fibers to recover from any dual-link failure (irrespective of whether failure locations are known or not). Thus a total of $4.5N$ capacity is required. On the other hand, the minimum number of links required to form a minimally four-connected graph is $2N$, where each link will be equipped with one primary and one spare fiber. Thus, a total of $4N$ capacity is required[5]. The MESH-4×4 network (16 nodes and 32 links) considered in this study is a minimally four-connected network that can be reduced to a minimally three-connected network by removing 8 links. The capacity provisioned under minimally three-connected and minimally four-connected scenarios would be 72 (24 primary + 48 spare) and 64 (32 primary + 32 spare) fibers-links, respectively.

The advantages of a minimally four-connected network over a minimally three-connected network are: (1) reduction in total

[5]Note that this computation does not take into account the lengths (or cost) of the link. The cost of increasing the connectivity could be higher if the additional links are longer (or higher cost), resulting in lower capacity savings.

capacity by up to 11% (and possibly associated link costs like line amplifiers, etc.); (2) higher connectivity (resulting in shorter paths, faster recovery notification); and (3) up to 33% increased primary capacity even with the 11% reduction in total capacity. The drawbacks of the four-connected network are: (1) possible increase in network cost due to longer links; (2) increased switching requirement at nodes (due to increase in node degree); and (3) increased link failure rate due to increase in the number of links by up to 33%. Although precise quantification of the trade-offs involved in the design will vary among networks, it is worth noting that designing a network with increased network connectivity may not necessarily be an expensive proposition.

While the capacity savings obtained through the knowledge of precise failure location may not be significant if the network is close to minimally three-connected, significant capacity savings may be obtained in networks with higher connectivity. Hence, solutions for dual-link failure resiliency that do not employ the precise knowledge of failure, such as the one considered in this paper, are typically useful only in networks that are closer to a minimally three-connected networks than four-connected networks.

VII. Conclusion and Future Work

This paper formally classifies the approaches for providing dual-link failure resiliency. Recovery from a dual-link failure using an extension of link protection for single-link failure results in a constraint, referred to as Backup Link Mutual Exclusion (BLME) constraint, whose satisfiability allows the network to recover from dual-link failures without the need for broadcasting the failure location to all nodes. The paper formulates the problem of finding backup paths for links satisfying the BLME constraint as an ILP and further develops a pseudo-polynomial time heuristic algorithm. The formulation and heuristic are applied to six different networks and the results are compared. The heuristic is shown to obtain a solution for most scenarios with a high failure recovery guarantee, although such a solution may have longer average hop lengths compared to the optimal values. The paper also establishes the potential benefits of knowing the precise failure location in a four-connected network that has lower installed capacity than a three-connected network for recovering from dual-link failures. Hence, approaches for dual-link failure resiliency that do not employ precise knowledge of failure location are well-suited for networks that are closer to minimally three-connected networks rather than four-connected networks.

The focus of this paper has been restricted to considering one primary fiber per link (in each direction) and recovering from failures through link protection at the granularity of a fiber. Our future work in this direction involves consideration of cost and capacity considerations on each link. It is also of interest to develop heuristic approaches that are guaranteed to obtain a feasible solution, although it may not be optimal.

Acknowledgment

The research presented in this paper is supported in part by National Science Foundation, under grants ANI-0325979 and ANI-0435490, and Connection One Circuits and Systems Research Center.

References

[1] M. Medard, S. G. Finn, and R. A. Barry, "WDM loop-back recovery in mesh networks," in *in Proceedings of IEEE INFOCOM*, March 1999, pp. 752–759.

[2] S. S. Lumetta, M. Medard, and Y. C. Tseng, "Capacity versus robustness: A tradeoff for link restoration in mesh networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 18, no. 12, pp. 1765–1775, December 2000.

[3] G. Ellinas, G. Halemariam, and T. Stern, "Protection cycles in WDM networks," *IEEE Journal of Selected Areas in Communication*, vol. 8, no. 10, pp. 1924–1937, October 2000.

[4] W. D. Grover, *Mesh-based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking*. New Jersey: Prentice Hall Publishers, 2003.

[5] M. Fredrick, P. Datta, and A. K. Somani, "Sub-graph routing: A novel fault-tolerant architecture for shared-risk link group failures in WDM optical networks," in *4th International Workshop on Design of Reliable Communication Networks*, October 2003.

[6] M. Clouqueur and W.Grover, "Availability analysis of span-restorable mesh networks," *IEEE Journal of Selected Areas in Communications*, vol. 20, no. 4, pp. 810–821, May 2002.

[7] ——, "Mesh-restorable networks with complete dual-failure restorability and with selectively enhanced dual-failure restorability properties," in *in Proceedings of OPTICOMM*, July-August 2002, pp. 1–12.

[8] J. Doucette and W. D. Grover, "Shared-risk logical span groups in span-restorable optical networks: Analysis and capacity planning model," *Photonic Network Communications*, vol. 9, no. 1, pp. 35–53, January 2005.

[9] J. A. Bondy and U. S. R. Murthy, *Graph Theory with Applications*. New York: American Elsevier Publishing, 1976.

[10] H. Choi, S. Subramaniam, and H. Choi, "On double-link failure recovery in WDM optical networks," in *Proceedings of IEEE INFOCOM*, June 2002, pp. 808–816.

[11] S. Ramasubramanian and A. Chandak, "Dual-link failure resiliency through backup link mutual exclusion," *Technical Report, Department of Electrical and Computer Engineering, University of Arizona*, December 2004.

[12] CPLEX Solver, http://www.cplex.com.

[13] H. Choi, S. Subramaniam, and H. Choi, "Loopback recovery from double-link failures in optical mesh networks," *IEEE/ACM Transactions on Networking*, vol. 12, no. 6, pp. 1119–1130, December 2004.