

Comparison of Failure Dependent Protection Strategies in Optical Networks

Srinivasan Ramasubramanian

Department of Electrical and Computer Engineering

University of Arizona, Tucson, AZ 85721

srini@ece.arizona.edu

Phone: (520) 621 4521

Fax: (520) 621 8076

Avinash S. Harjani

National Instruments

11500 N Mopac Expwy

Austin, TX 78759-3504

Avinash.Harjani@ni.com

Corresponding Author: Srinivasan Ramasubramanian (srini@ece.arizona.edu)

Abstract

The criticality of survivable network design and operation increases with increasing transmission speed. Path protection strategies achieve better network utilization compared to link protection strategies; however the recovery time of connections in path protection strategies are higher than that in link protection strategies. This paper evaluates and compares the performance of three failure dependent strategies: (1) failure dependent path protection; (2) link protection; and (3) Diversion – a variant of the segmented path protection approach. In addition, A framework for evaluating the connection recovery time is also developed. The protection strategies are compared for their recovery time and blocking performance using extensive simulations.

1 Introduction

Optical networks employing wavelength division multiplexing (WDM) and wavelength sharing among multiple low-rate traffic streams provide a scalable backbone network architecture. Present day networks have transmission speeds of up to 40 Gbps (OC-768) with each wavelength shared by connections with much lower capacity like 155 Mbps (OC-3) or 622 Mbps (OC-12). As the optical processing and buffer technologies are not mature currently to achieve routing individual packets in runtime, optical networks of today and those in the near future are expected to employ connection-oriented service paradigm. In such backbone networks, the major network operation is to establish a connection between source-destination pair on demand and release them when the connection is not needed.

Connection establishment in a connection-oriented network consists of two steps: *path selection* and *channel assignment*. Path selection refers to selecting a path from source to destination based on certain criteria. Channel assignment refers to assigning one or more channels depending on the requirement of the call on every link of the chosen path. Path selection can be carried out in several ways. If a source-destination pair has one pre-selected path, then it is referred to as *fixed-path* approach. If a path is selected depending on the network status from a pre-selected set of candidate paths, then it is referred to as *dynamic path selection*. The set of candidate paths remain the same at all times and do not change with the network status. If the candidate paths are chosen based on the network status, the path selection process is referred to as *exhaustive routing*. Channel assignment refers to allocation of specific resources on every link of a chosen path, for example: (a) fiber, wavelength, and time slot assignment on the links in a WDM grooming network; and (b) fiber and wavelength assignment in a multi-fiber wavelength-routed network. Irrespective of the path selection or channel assignment strategy employed in the network, obtaining information along a path to assess the availability of resources to establish the connection becomes the fundamental requirement. Information collection in WDM grooming networks involve identifying availability of resources on the links along with the grooming capability of intermediate nodes on a specific path to identify resource availability on the path.

Dynamic connection establishment has been extensively studied in the context of wavelength-routed WDM networks [1, 2, 3, 4, 5, 6, 7, 8]. However, this issue has received very little attention in the context of WDM grooming networks until recently [9, 10, 11]. Similarly, survivable routing has also received significant interest in the context of wavelength-routed networks [12], however it is in its early stages of research in the context of grooming networks. In [11], a framework for connection establishment in optical networks employing traffic grooming and heterogeneous switching architectures has been developed. The framework, referred to as Methodology for Information Collection and Routing in Optical Networks (MICRON), outlines a representation mechanism for link information as matrices, approaches to combining link information to obtain path information, and dynamic routing in the presence of a combination of wavelength and time slot switching.

In order to protect connections from link failures in the network, often two paths are assigned: a primary path on which a connection is established and *backup* path on which a connection will be setup in case a primary path fails. A combination of links may share resources in a network, a duct or conduit through which they are laid out, which would result in a failure of more than one link at an instant. Such failures are modeled as Shared Risk Link Group (SRLG). Typically, the objective of the network operation is to protect connections against any SRLG failure. In this paper, we consider only protection schemes as they typically have faster recovery times as compared to restoration schemes.

1.1 Taxonomy of protection schemes

Protection schemes proposed in the literature can be broadly classified as link protection and path protection.

Link protection. Link protection schemes route a connection around a failed link. Re-routing is performed by the node connected to the failed link to the neighboring node on the original path. Such a protection may be

achieved in the network in a way that is transparent to the source node, except in cases where a link connected to the source or destination fails.

Fig. 1 shows an example network in which a connection from node 1 to 4 is established along the primary path 1–2–3–4. Upon failure of a link in the primary path, the connection is re-routed around the failed link as shown in Fig. 2. The channel assignment for the connection in the remaining links of the primary path remains unchanged. For example, when link 2–3 fails, the channel assignment on the backup path 2–5–3 must be such that node 2 can switch the connection from the original channel assignment on link 1–2 to that of 2–5 and node 3 can switch the connection from link 2–5 to the original channel assignment on link 3–4. Such a consistency between primary and backup paths must be maintained for all link failures that affect the primary path of the connection.




Figure 1. Example network in which a connection from node 1 to 4 is established along the primary path 1–2–3–4.

Link protection may be performed at either the granularity of a fiber or connection. Link protection at fiber granularity assumes that every link has primary and spare fibers. Primary fiber is used for routing working connections, while the secondary fiber is used only when a failure occurs. As the protection scheme operates at the granularity of a fiber, the consistency in channel assignment between primary and backup path is automatically satisfied for all link failures. Link protection at the fiber level offers fast recovery time requiring lesser signaling compared to path protection approaches. However, the drawback of switching at the fiber level is that the network cannot take advantage of those connections that may not require protection, as every connection is treated as protected traffic. Link protection at the connection level, referred to as connection switched link protection (CSLP), offers significant improvement when traffic requires different levels of protection [13]. As CSLP operates at the granularity of a connection, the consistency in channel assignment among primary and backup paths must be explicitly guaranteed by the connection establishment procedure.

Path protection. Path protection schemes recover from a failure by re-routing the connections at the source. Path protection schemes may be classified into two categories based on their knowledge of the failure location. Assignment of a backup path that does not require precise knowledge of the link failure is referred to as *failure-independent path protection* (FIPP). Alternatively, if a connection may be assigned more than one backup path depending on the failure, then it is referred to as *failure-dependent path protection* (FDPP). Under FDPP, if the path (and channels) assigned to the connection under no failure is the same as the path (and channels) assigned under any other failure that does not affect the path, then it is referred to as *Strict FDPP* [14], otherwise, it is referred to as *Flexible FDPP*. The L+1 sub-graph routing developed in [15] is an example of flexible FDPP strategy.

This paper focuses on strict FDPP for path protection. For the example in Fig. 1, the backup paths obtained using strict FDPP is shown in Fig. 3. Note that the backup paths may not necessarily be the shortest path in the failed network as the shortest path may not have sufficient capacity to accommodate the connection. As the connections are configured on an end-to-end basis, the backup connection may be treated as a new connection in the failed network, hence no specific constraints are placed on the channel assignment for backup paths. The end-to-end reconfiguration of the paths results in better utilization of network resources compared to link protection, however the connection recovery times are higher in path protection as the failure notification has to be sent to the




Figure 2. Backup paths using link protection strategy.




Figure 3. Backup paths using failure dependent path protection strategy.

source node to re-route the traffic.

Segmented path protection. The segmented path protection strategy divides a primary path into multiple segments, where each segment denotes a set of contiguous links on the primary path [12]. Each segment is then protected using path protection strategy. If a link is present in only one segment, then the backup path corresponding to that segment will protect the connection upon failure. However, if a link is present in more than one segment, then a backup path must be chosen from the set of backup path segments. While there may be several possible ways in which segmented path protection may be achieved [16], this paper considers one variant of the segmented path protection, referred to as *Diversion*, that diverts the connection from the last node before the failed link directly to the destination. For the example considered in Fig. 1, the backup paths obtained using Diversion are shown in Fig. 4. On failure of link 2–3, the connection is diverted at node 2 along the path 2–5–6–4. The connection after the link 2–3 failure is routed along 1–2–5–6–4, with the channel assignment on link 1–2 (or any link before the failed link) remaining unchanged even after the failure. The connection establishment procedure must therefore ensure that the last node in the primary path before the failure can divert the connection successfully to the destination. It is worth noting that the Diversion strategy behaves similar to link protection strategy when the failed link is closer to the destination and like path protection strategy when the failure is closer to the source. This scenario is also depicted in Fig. 4(a) where the backup path under Diversion is the same as that of path protection when the first link in the primary path fails. The backup path when the last link in the primary path fails, see Fig. 4(c), is the same as the link protection strategy.




Figure 4. Backup paths using Diversion strategy.

It is worth noting that the failure independent path protection and link protection are extreme cases of segment protection where the former assumes the entire path as a segment while the latter assumes each link as a segment.

Backup multiplexing. In order to achieve efficient utilization of network resources, multiplexing of resources across multiple backup paths and/or a primary path may be employed. More than one backup path may share a resource as long as any failure in the network will cause at most one of the corresponding working connections to fail. If a resource is shared only among backup paths, then it is referred to as *backup-backup* multiplexing. If a resource is occupied by a working connection and is also assigned to one or more backup paths, then it is referred to as *primary-backup* multiplexing. Any failure that would require the shared resource for establishing a backup connection must lead to the failure of the existing primary connection occupying that resource.

Contribution. This paper develops a connection establishment procedure to route connections that are protected using the three path protection strategies. The computation of the primary path is performed dynamically using the Available Shortest Path (ASP) algorithm that accounts for the availability of a diversion path from any node to the destination. Hence, a successful computation of primary path guarantees the existence of a diversion path (and a consistent channel assignment) for every link failure in the network. The recovery times of different path protection strategies are analyzed using the failure recovery time computation model developed in [17]. The effectiveness of the Diversion approach is studied and compared to CSLP and strict FDPP strategies through extensive simulations on three different networks, where the networks specifically bring out some finer aspects of the protection strategies.

1.2 Organization

The remainder of the paper is organized as follows: Section 2 explains the assumption on the network model, node architecture, and notations employed. Section 3 describes the connection establishment procedure involved in the different protection strategies considered in this paper. The computation of failure recovery time is described in Section 4. The results of the performance study on various protection strategies are described in Section 5. Section 6 concludes the paper.

2 Network model

This paper considers a generic optical network, where links may have multiple fibers, multiple wavelengths per fiber, and multiple time slots per wavelength, and nodes employing heterogeneous switching architectures. Let \mathcal{N} denote the set of nodes and \mathcal{L} denote the set of physical links in the network. The links are assumed to be bi-directional with dedicated resources (fibers) for each direction. Let Ψ denote the set of Shared Risk Link Groups (SRLG) in the network. An element $\psi \in \Psi$ is a subset of \mathcal{L} that denotes the set of links that may fail due to a failure in one or more shared resources.

The developed protection strategies protect the connections (that require protection) against any single SRLG failure at a given instant of time. The following subsections describe the information stored in each link, combining the link information to obtain path information, path selection and capacity allocation and release that are a part of the MICRON framework. For a detailed description of the framework, the readers are referred to [11]. The notations employed in this paper are shown in Table 1 for easy reference.

2.1 Link information

A link ℓ is represented by one or more variables (for each direction). These variables are assumed to be matrices in this paper, which can be reduced to a vector (diagonal matrix) or simply a scalar (1×1 matrix) if needed. For example, if the network converts optical signals to electronic domain at every node, then such a network does not need any wavelength specific information. Hence, the bandwidth available in a link may be represented by

Table 1. Comprehensive list of notations with comments.

Variables	Comments
\mathcal{N}	Set of nodes.
\mathcal{L}	Set of links.
ψ	An SRLG failure. ($\psi \subset \mathcal{L}$)
Ψ	Set of SRLG failures in the network.
α	Time to detect a single link failure.
β	Switch reconfiguration time at a node.
γ	Electronic overhead time in transmitting and receiving a failure notification message.
$\Delta^\psi(n_1, n_2; \{x_\ell\})$	Cost of the least-cost path from n_1 to n_2 with $\{x_\ell\}$ as the cost metric for links under failure ψ .
$\Delta_{\mathcal{P}}(n_1, n_2; \{x_\ell\})$	Cost of the segment from n_1 to n_2 on path \mathcal{P} with $\{x_\ell\}$ as the cost metric for links.
$Pred(n, \mathcal{P})$	Immediate predecessor of node n on path \mathcal{P} .
$\mathcal{P}(s, d, \psi, \{x_\ell\})$	Path computed dynamically from s to d by removing links in the set ψ and $\{x_\ell\}$ as link capacities.
ϕ	Null set.
T_n^ψ	Time taken for node n to receive the failure notification of failure ψ since the instant of failure.
R_n^ψ	Time to finish reconfiguration at node n from the instant of failure ψ .
S_ℓ	Maximum capacity on link ℓ .
P_ℓ	Capacity occupied by primary connections on link ℓ .
G_ℓ^ψ	Capacity gained on link ℓ upon failure ψ .
B_ℓ^ψ	Capacity reserved for backup on link ℓ for failure ψ .
A_ℓ	Available capacity on link ℓ to route primary connection.
A_ℓ^ψ	Available capacity on link ℓ to route backup connection under failure $\psi \in \Psi$.
τ_ℓ	Propagation delay on link ℓ .
s_ℓ	Source node of link ℓ .
d_ℓ	Destination node of link ℓ .
Ψ_ℓ	Set of failures under which link ℓ fails.
Z_ℓ^ψ	Backup path for link ℓ under failure ψ .
X_ℓ	Capacity available to route primary connection.
$\mathcal{Y}_\ell^\psi(\mathcal{R})$	Diversion path for link ℓ under failure ψ for request \mathcal{R} ; from node s_ℓ to destination node of a request $d_{\mathcal{R}}$.
$\mathcal{P}_{\mathcal{R}}$	Set of nodes and links through which the primary path of request \mathcal{R} traverses.
$\mathcal{P}_{\mathcal{R}}^\psi$	Set of nodes and links through which the backup path of request \mathcal{R} traverses under failure ψ .
$\xi_{\mathcal{R}}(\ell)$	Channel assignment for request \mathcal{R} on link ℓ under no failure.
$\xi_{\mathcal{R}}^\psi(\ell)$	Channel assignment for request \mathcal{R} on link ℓ under failure ψ .
$\Psi_{\mathcal{R}}$	Set of failures under which the request \mathcal{R} will be reconfigured to its backup connection.
$s_{\mathcal{R}}$	Source of request \mathcal{R} .
$d_{\mathcal{R}}$	Destination of request \mathcal{R} .
$c_{\mathcal{R}}$	Capacity requirement of request \mathcal{R} .
$x_{\mathcal{R}}^\psi$	Node at which the request \mathcal{R} is re-routed for failure ψ .
$y_{\mathcal{R}}^\psi$	Node at which the reconfigured segment of request \mathcal{R} for failure ψ joins the primary path.
$z_{\mathcal{R}}^\psi$	First node in the primary path such that no link in the path segment $z_{\mathcal{R}}(\psi)$ to the destination is affected by failure ψ .
$L_{\mathcal{R}}^\psi(n)$	Latest time by which connection \mathcal{R} can cross node n in its primary path after failure ψ .
$F_{\mathcal{R}}^\psi(n)$	Earliest time by which connection \mathcal{R} can cross node n in its backup path after failure ψ .
$T_{\mathcal{R}}^\psi$	Recovery time for request \mathcal{R} under failure ψ .

a scalar. If the network employs wavelength routing and the nodes do not employ wavelength conversion, then the bandwidth available on a link is represented by a vector (or a W -tuple, where W represents the number of

wavelengths). If the nodes employ heterogeneous switching architectures (such as a combination of wavelength conversion and time slot interchange), then every link is represented as a matrix.

Let S_ℓ and P_ℓ denote the total number of channels and number of channels occupied by primary connections on link ℓ , respectively. Let G_ℓ^ψ denote the number of channels on link ℓ that are currently occupied by working connections that would fail in case of the SRLG failure ψ . The channels occupied by the failed primary connections will become available, therefore may be assigned for backup connections. Releasing of capacity occupied by a failed primary is referred to as “stub-release” in the literature and has been shown to improve capacity utilization in the networks. Let B_ℓ^ψ denote the number of backup channels required on link ℓ in case of an SRLG failure ψ . In the rest of this paper, the notation ℓ will refer to a link connecting two nodes in general. When used as a suffix for the above matrices, it refers to a specific direction that will depend on (and will be obvious from) the context in which it is used.

2.2 Path information

A path in the network, represented as a set of directional links (ordered in the sequence by which it appears in the path), is denoted by \mathcal{P} . The information on a path is obtained by combining the information of the links in the path as:

$$X_{\mathcal{P}} = \prod_{\ell \in \mathcal{P}} X_\ell \quad (1)$$

where \prod denotes generalized matrix multiplication. Note that the direction of the link is evident from the path. If X , Y , and Z are three matrices and $Z = XY$, then an element z_{ij} of the matrix Z is obtained as:

$$z_{ij} = (x_{i1} \otimes y_{1j}) \oplus (x_{i2} \otimes y_{2j}) \oplus \dots \oplus (x_{iC} \otimes y_{Cj}) \quad (2)$$

where C denotes the number of columns of matrix X , or equivalently number of rows of matrix Y . The operators \otimes and \oplus , denoted as a tuple (\otimes, \oplus) , can be defined in different combinations so that several meaningful results are obtained. It can be observed that when \otimes is integer multiplication and \oplus is integer addition operation, the above equation denotes the traditional matrix multiplication. Some examples of matrix representation of optical grooming networks may be found in [11]. When combining available capacity in particular, the operator tuple (\otimes, \oplus) is set as (\min, \max) .

2.3 Path selection and channel assignment

A path from a node s to d selected by disabling a set of links in a failure set ψ with the available capacity information on links as $\{x_\ell\}$ is denoted by $\mathcal{P}(s, d, \psi, \{x_\ell\})$. This work employs the available shortest path routing algorithm [18] that selects the shortest path (based on hop-count) among the available paths (paths that have sufficient capacity for routing the connection).

A request \mathcal{R} that arrives in a network is provided a connection on a primary path and, if required, one or more backup paths. Let $\mathcal{P}_{\mathcal{R}}$ denote the set of links in the primary path of the connection. The connection, in general, will be reconfigured under a set of failures denoted by $\Psi_{\mathcal{R}}$. For every failure $\psi \in \Psi_{\mathcal{R}}$, a backup path is provided for the request. Let $\mathcal{P}_{\mathcal{R}}^\psi$ denote the links in the backup path of the connection corresponding to the failure ψ .

The network is also assumed to have a channel assignment strategy given the path with sufficient capacity is identified. The work presented here employs first-fit channel assignment strategy. Let $\xi_{\mathcal{R}}(\ell)$ denote the channel assignment on link ℓ for the connection under no failure and $\xi_{\mathcal{R}}^\psi(\ell)$ the channel assignment on link ℓ under failure ψ .

3 Connection Establishment

A network operating with no failures satisfies the following equation:

$$P_\ell \leq S_\ell \quad (3)$$

On a failure ψ , the network will be able to re-assign requests to their backup connections if there is no resource contention, or equivalently, if the following condition is satisfied:

$$P_\ell - G_\ell^\psi + B_\ell^\psi \leq S_\ell \quad (4)$$

Channel allocation under any protection strategy, either for working or backup path, must not violate the above inequality for any SRLG failure in order for the network to be resilient to any SRLG failure.

Based on the above necessary conditions, the available capacity on a link may be computed. When the network has no failures, the available capacity on a link is computed as:

$$A_\ell = S_\ell - P_\ell \quad (5)$$

The available capacity on a link under a failure ψ is computed as:

$$A_\ell^\psi = S_\ell - P_\ell - B_\ell^\psi + G_\ell^\psi \quad (6)$$

It is worth noting that the above assignment only takes into account those connections that would be re-assigned in case of failure ψ . Hence, backup multiplexing (backup-backup and primary-backup) is inherent to the above computation for available capacity under a failure.

Connection establishment in general requires assignment of a primary path and, if required, one or more backup paths for each request. Let X_ℓ denote the available capacity on link ℓ to route a primary connection. Depending on the protection strategy, the computation of resources in the network to route primary and backup connections will differ.

The following subsections describe FDPP, CSLP, and Diversion protection strategies in detail. In order to present the information in a concise manner, a comprehensive list of available capacity and path computation under various protection strategies are shown in Table 2. The readers are recommended to refer to this table for the corresponding protection strategy under consideration.

3.1 Failure Dependent Path Protection (FDPP)

FDPP attempts to provide multiple backup paths, one for each failure under which the primary path may no longer be available. While the backup paths for some failures may be the same, it is not guaranteed to be the same for all failures that affect the primary path. The primary path for the connection may be chosen from a set of candidate paths or computed dynamically. In case of selecting a path from a set of candidate paths, the computation of available capacity for routing primary connection is the same as that described in FIPP approach.

Under a network failure, the primary path of a connection may not be available for two reasons: (1) a failure in the network involves one or more links in the primary path; or (2) a failure in the network does not involve any link on the primary path, however the channel(s) assigned to the connection on the primary path is also assigned to the backup paths of some other connections that is affected by the failure. The latter results in a *domino effect* that might lead to a network-wide reconfiguration.

Strict FDPP. For dynamic path selection schemes which ensure that a connection will not be reconfigured unless a failure affects its primary path, the available capacity must not violate the necessary and sufficient condition under any SRLG failure. The failure dependent protection developed in [14] is an example of the strict FDPP strategy.

Table 2. Comprehensive list of capacity and path computation for various protection strategies.

Metrics	Strict FDPP	CSLP	Diversion
X_ℓ	$\min(A_\ell, \min_{\psi \in \Psi} A_\ell^\psi)$	$\min(A_\ell, \min_{\psi \in \Psi} A_\ell^\psi, \min_{\psi \in \Psi_\ell} R_\ell^\psi)$	
$\mathcal{P}_\mathcal{R}$	$\mathcal{P}(s_\ell, d_\ell, \phi, \{X_\ell\})$		
$\Psi_\mathcal{R}$	$\{\psi \mid \mathcal{P}_\mathcal{R} \cap \psi \neq \phi\}$		
$\mathcal{Y}_\ell^\psi(\mathcal{R})$	–	–	$\mathcal{P}(s_\ell, d_\mathcal{R}, \psi, \{A_\ell^\psi\})$
R_ℓ^ψ	–	$\prod_{\ell' \in \mathcal{Z}_\ell^\psi} A_{\ell'}^\psi$	$\prod_{\ell' \in \mathcal{Y}_\ell^\psi(\mathcal{R})} A_{\ell'}^\psi$
$\mathcal{P}_\mathcal{R}^\psi$	$\mathcal{P}(s_\ell, d_\ell, \psi, \{A_\ell^\psi\})$	Replace links affected by failure with backup paths.	Replace primary path from the failed link with diversion path.

Dynamic routing algorithms require available capacity computation in order to identify a feasible path. Therefore, dynamic routing algorithms will have to estimate the available capacity for routing a primary connection (X_ℓ) in a conservative manner (see Table 2) by ensuring that the capacity is available under all failure scenarios. Based on such a computation of available capacity, the path selection strategy in the network selects an appropriate primary path for the connection. As the above computation guarantees that the capacity assigned for primary path is available under any failure, the connection needs to be reconfigured only for those failures that affect the primary path. Hence, the set of failures that leads to a reconfiguration ($\Psi_{\mathcal{R}}$) is computed as those failures that affect the primary path.

A backup path needs to be computed for every failure $\psi \in \Psi_{\mathcal{R}}$. The backup path under a failure ψ is computed dynamically by removing the links that are affected by the failure. As the backup paths may be computed independently for each failure scenario, the available capacity on a link to route backup path under failure ψ is simply A_ℓ^ψ .

3.2 Connection Switched Link Protection (CSLP)

For link protection at the connection level, the connection is re-routed around the failed link. The channel assignment on the links of the primary path not affected by the failure remains the same. Hence, the capacity to route primary connection on a link is assumed to be available only when the backup path also has the required capacity under the failure of the link. Let Z_ℓ^ψ denote the backup path for link ℓ under failure ψ . The backup path for a link under a failure ψ is computed as the shortest path between the nodes connected by the link after removing all the links affected by the failure. The capacity available on the backup path Z_ℓ^ψ upon failure ψ , denoted by R_ℓ^ψ , is computed as shown in Table 2.

Note that a link may have several backup paths, one for each failure that affects the link. A primary connection routed along this link may be re-routed to any of its backup path depending on the failure. Hence, the capacity that is assigned for primary connection on link ℓ must also be available along all of its backup paths under the corresponding failure. Therefore, the capacity available for routing primary connection is computed by considering the capacity available when the network does not have any failures (A_ℓ), available capacity on a link under all failure scenarios (A_ℓ^ψ), and available capacity on the backup path of the link under the failure scenarios affecting the link (R_ℓ^ψ).

The backup path for a connection under each failure scenario is obtained by simply removing the failed links in the primary path and appending the backup links in their place. A successful connection has one primary path and $|\Psi_{\mathcal{R}}|$ backup paths. The channel assignment on the backup path must be consistent with the channel assignment of the primary path still intact. As the routing of primary connection has taken into account the availability of backup paths, a consistent channel assignment on the backup path is guaranteed to exist¹.

3.3 Diversion

Recall that Diversion protection is similar to that of link protection, except that the connection is re-routed from the node attached to the failed link directly to the destination. In link protection in which the link in the primary path of a connection that are not affected by the failure remain unaffected. In Diversion, the primary path from the source to the node before the failed link remains unaffected, the primary path segment after the failed link is not valid after the failure.

As the primary path will be diverted directly to the destination, a path from every node to the destination is required. The diversion path between every node pair may be fixed or computed dynamically. The computation of diversion path requires the knowledge of request destination, hence may be performed only after request arrival.

¹There may be some cases where path pruning may be necessary. The readers are referred to [17] for a discussion on path pruning.

Let ℓ denote a directed link in the graph and s_ℓ and d_ℓ denote the source and destination of the directed link. Such a notation means a primary path routed through the directional link ℓ traverses from s_ℓ to d_ℓ . Let \mathcal{R} be a request requiring a connection from source $s_{\mathcal{R}}$ to destination $d_{\mathcal{R}}$. Let $\mathcal{Y}_\ell^\psi(\mathcal{R})$ denote the diversion path for failure ψ from link ℓ , which is a path from node s_ℓ to $d_{\mathcal{R}}$. As the diversion path will be used only under a particular failure scenario, the available capacity on a diversion path (R_ℓ^ψ) is computed by only considering that failure scenario as shown in Table 2. The available capacity on a link to route primary connection (X_ℓ) is then computed similar to the CSLP strategy².

The backup path for a connection under a particular failure scenario is obtained by removing the segment of the path from the first failed link to the destination and appending the diversion path from the node before the failed link to the destination. As the primary path computation takes into account the availability of capacity on backup path, a backup path and channel assignment is guaranteed to exist. The channel assignment on the backup (diversion) path must be consistent with that of the primary path segment still intact.

3.4 Generic connection establishment/release procedure

The generic procedure for connection establishment and release is shown in Fig. 5. The connection establishment procedure involves five major steps. At the end of Step 4, the connection is assigned a primary path and a set of backup paths depending on the protection requirement. Once the primary and backup paths are obtained, the capacities on the links are updated. It is worth noting that the way in which the link capacities are maintained allows the different protection strategies to be employed in the same network. The network provides protection at the granularity of a connection (including link protection). The connection release procedure is similar to the Step 5 of connection establishment, except that the capacities are released instead of being assigned.

4 Failure Recovery Time Computation

The failure recovery time for different protection strategies depend on where the reconfiguration takes place. This section computes in detail the failure recovery time required under a single link failure. The failure recovery time computation was originally developed in [17] and the analysis is valid for Diversion strategy as well. The discussion in this section is presented for this paper to be self-contained.

A single-link failure implies a link, on both directions, between two nodes fails. As backup paths and channel assignment for different failures are computed during the connection establishment phase for “protection” strategies, the exact switch configurations at different nodes are known prior to the failure. The switch configurations under different failures may be stored at the individual nodes to reduce reconfiguration time. On a failure, nodes attached to the failed link detect the failure and broadcast a failure notification message. Every node forwards the failure notification further upon receipt and reconfigures its switches corresponding to the failure indicated in the notification message.

On a link failure, nodes connected to the failed link detect the failure first. The failure is assumed to be detected due to a loss of periodic “Hello” packets transmitted over the control channel for a pre-specified duration, hence the time required to detect a failure is assumed to be a constant, denoted by α . On failure detection, the nodes broadcast a failure notification message. The sum of the time required for the node to prepare and transmit the packet on a link and the time to process the packet by the node on the other end of a link is referred to as the electronic overhead time (γ). If τ_ℓ denotes the propagation delay on link ℓ , then $\tau_\ell + \gamma$ denotes the hop delay seen by the failure notification message on link ℓ . Note that as the failure notification message will be converted

²Note that there may be some cases where the backup path (from source to destination) under diversion may traverse a link in the same direction twice, similar to looping in any link protection strategy. In some cases, pruning the resultant path may be necessary to successfully route the backup. Studies conducted by the authors have shown that such scenarios are very rare, a probability of 10^{-5} that a connection encounters such a scenario.

Connection establishment procedure**Input:** Request \mathcal{R} with a specific protection requirement.**Output:**

1. Primary path $\mathcal{P}_{\mathcal{R}}$ and channel assignment on primary path $\xi_{\mathcal{R}}(\ell), \forall \ell \in \mathcal{P}_{\mathcal{R}}$.
2. Failure set $\Psi_{\mathcal{R}}$.
3. A set of backup path for each failure in the failure set $\mathcal{P}_{\mathcal{R}}^{\psi}$ and channel assignment $\xi_{\mathcal{R}}^{\psi}(\ell), \forall \ell \in \mathcal{P}_{\mathcal{R}}^{\psi}$ and $\psi \in \Psi_{\mathcal{R}}$.

Steps:

1. Update the available capacity on each link to route primary connection (X_{ℓ}).
Note: If backup path for a link (for CSLP) or diversion path (for Diversion) needs to be computed dynamically, the paths are computed in this step.
2. Obtain a primary path employing Available Shortest Path (ASP) algorithm. Obtain a sub-trunk assignment on the path employing first-fit strategy. If a path or sub-trunk assignment cannot be obtained the request is rejected. Go to Step 6.
3. Obtain the failure set under which a reconfiguration is required ($\Psi_{\mathcal{R}}$).
4. For every $\psi \in \Psi_{\mathcal{R}}$, obtain a backup path and channel assignment. Update the available capacity on each link to route a backup connection under failure ψ as A_{ℓ}^{ψ} .
 - **Strict FDPP:** Compute the backup path $P_{\mathcal{R}}^{\psi} = \mathcal{P}(s_{\mathcal{R}}, d_{\mathcal{R}}, \psi, \{A_{\ell}^{\psi}\})$.
 - **CSLP:** Construct the backup path $P_{\mathcal{R}}^{\psi}$ by replacing the links affected by the failure ψ with their corresponding backup paths.
 - **Diversion:** Construct the backup path $P_{\mathcal{R}}^{\psi}$ by replacing the primary path from the failed link with the diversion path.
5. Update link capacities.
Note: At this juncture, every request has been assigned: (1) a primary path $\mathcal{P}_{\mathcal{R}}$ with channel assignment $\xi_{\mathcal{R}}$; (2) failure set $\Psi_{\mathcal{R}}$; and (3) a set of backup paths, $\mathcal{P}_{\mathcal{R}}^{\psi}, \forall \psi \in \Psi_{\mathcal{R}}$. In order to update the link capacities, it is not necessary to distinguish which failure scheme is employed.

$$P_{\ell}[\xi_{\mathcal{R}}(\ell)] \leftarrow P_{\ell}[\xi_{\mathcal{R}}(\ell)] + c_{\mathcal{R}} \quad \forall \ell \in \mathcal{P}_{\mathcal{R}}$$

$$G_{\ell}^{\psi}[\xi_{\mathcal{R}}^{\psi}(\ell)] \leftarrow G_{\ell}^{\psi}[\xi_{\mathcal{R}}^{\psi}(\ell)] + c_{\mathcal{R}} \quad \forall \ell \in \mathcal{P}_{\mathcal{R}} \text{ and } \psi \in \Psi_{\mathcal{R}}$$

$$B_{\ell}^{\psi}[\xi_{\mathcal{R}}^{\psi}(\ell)] \leftarrow B_{\ell}^{\psi}[\xi_{\mathcal{R}}^{\psi}(\ell)] + c_{\mathcal{R}} \quad \forall \ell \in \mathcal{P}_{\mathcal{R}}^{\psi} \text{ and } \psi \in \Psi_{\mathcal{R}}$$

6. Exit.

Connection release procedure**Input:** Request \mathcal{R} which has already been accepted.**Steps:**

1. Update link capacities.

$$P_{\ell}[\xi_{\mathcal{R}}(\ell)] \leftarrow P_{\ell}[\xi_{\mathcal{R}}(\ell)] - c_{\mathcal{R}} \quad \forall \ell \in \mathcal{P}_{\mathcal{R}}$$

$$G_{\ell}^{\psi}[\xi_{\mathcal{R}}^{\psi}(\ell)] \leftarrow G_{\ell}^{\psi}[\xi_{\mathcal{R}}^{\psi}(\ell)] - c_{\mathcal{R}} \quad \forall \ell \in \mathcal{P}_{\mathcal{R}} \text{ and } \psi \in \Psi_{\mathcal{R}}$$

$$B_{\ell}^{\psi}[\xi_{\mathcal{R}}^{\psi}(\ell)] \leftarrow B_{\ell}^{\psi}[\xi_{\mathcal{R}}^{\psi}(\ell)] - c_{\mathcal{R}} \quad \forall \ell \in \mathcal{P}_{\mathcal{R}}^{\psi} \text{ and } \psi \in \Psi_{\mathcal{R}}$$

Figure 5. Generic connection establishment/release procedure.




Figure 6. Illustration of $x_{\mathcal{R}}^{\psi}$, $y_{\mathcal{R}}^{\psi}$, and $z_{\mathcal{R}}^{\psi}$ for various protection strategies.

from the optical to electronic domain for processing at every node, this delay may be a significant factor in some networks.

Let $n_{\ell}(\psi)$ and $n'_{\ell}(\psi)$ denote two nodes connected to the failed link corresponding to failure ψ . Both these nodes broadcast the failure notification independently, hence a node n on the network will be aware of the failure from the message that arrives first. The time to get the notification of failure ψ at node n from the instant at which failure occurred in the network, denoted by T_n^{ψ} , is computed as shown in Equation 7. The nodes start to reconfigure their switches when they receive the notification. The time at which the reconfiguration will be completed at a node n for failure ψ , denoted by R_n^{ψ} , is computed as shown in Equation 8. Based on these computations, the recovery time seen by an individual connection is computed.

The recovery time of a connection is the time duration from the last information received on the primary path to the first information received on the backup path.

$$T_n^{\psi} = \alpha + \min[\Delta^{\psi}(n_{\ell}(\psi), n; \{\tau_{\ell} + \gamma\}), \Delta^{\psi}(n'_{\ell}(\psi), n; \{\tau_{\ell} + \gamma\})] \quad (7)$$

$$R_n^{\psi} = T_n^{\psi} + \beta \quad (8)$$

Consider an example connection for a request \mathcal{R} and a failure ψ . The timing calculations are performed with the failure instant as the reference. If t_1 denotes the time instant at which the destination receives the last valid information along the primary path and t_2 denotes the time at which the destination receives the first valid information along the backup path, then the difference $t_2 - t_1$ is treated as the recovery time (or equivalently, outage time). Note that when a link fails, information to the destination is still in transit on the links after the failed link. In order to compute the recovery time, a few more notations are introduced. Let $z_{\mathcal{R}}^{\psi}$ denote the first node in the primary path such that no link in the path segment from $z_{\mathcal{R}}^{\psi}$ to the destination is affected by failure ψ . The segment from $z_{\mathcal{R}}^{\psi}$ to the destination is defined as the *last surviving segment of the primary path*. Similarly, the longest segment of the primary path starting from the source that does not have any failed links is referred to as the *first surviving segment of the primary path*. If a failure does not affect the primary path, then $z_{\mathcal{R}}^{\psi}$ is the source node. The nodes at which the reconfiguration start and end for a connection depends on the protection strategy. Let $x_{\mathcal{R}}^{\psi}$ and $y_{\mathcal{R}}^{\psi}$ denote the nodes at which the reconfiguration starts and ends, respectively. The illustration of $x_{\mathcal{R}}^{\psi}$, $y_{\mathcal{R}}^{\psi}$, and $z_{\mathcal{R}}^{\psi}$ for various protection schemes is shown in Fig. 6. It is worth noting that the above nomenclature is valid irrespective of the protection strategy (link protection, path protection, *segmented* protection).

Let $L_{\mathcal{R}}^{\psi}(n)$ denote the latest time by which connection \mathcal{R} crosses node n , where n is a node in the last surviving segment of the primary path. Let $F_{\mathcal{R}}^{\psi}(n)$ denote the earliest time by which the connection crosses node n in its backup path for failure ψ . The failure recovery time of the connection under failure ψ , denoted by $T_{\mathcal{R}}^{\psi}$, is computed as in Equation 9.

A node n starts to reconfigure its switch as soon as it receives the failure notification message. If the latest time by which the connection crosses the immediate predecessor node of n on the primary path [$Pred(n, \mathcal{P}_{\mathcal{R}})$] is time t and the propagation delay on the link connecting the nodes n and $Pred(n, \mathcal{P}_{\mathcal{R}})$ is τ_{ℓ} , then the latest time by which the connection crosses n at $L_{\mathcal{R}}^{\psi}(n) = \min(T_n^{\psi}, t + \tau_{\ell})$. The recursive way of computing this metric on a node belonging to the last surviving segment of the primary path is shown in Equation 10. When a failure affects the primary path of a connection, the information that just crossed over the failure point is the last bit of information

$$T_{\mathcal{R}}^{\psi} = F_{\mathcal{R}}^{\psi}(y_{\mathcal{R}}^{\psi}) - L_{\mathcal{R}}^{\psi}(y_{\mathcal{R}}^{\psi}) \quad (9)$$

$$L_{\mathcal{R}}^{\psi}(n) = \begin{cases} \min[T_n^{\psi}, L_{\mathcal{R}}^{\psi}(\text{Pred}(n, \mathcal{P}_{\mathcal{R}})) + \Delta_{\mathcal{P}_{\mathcal{R}}}(n, \text{Pred}(n, \mathcal{P}_{\mathcal{R}}), n; \{\tau_{\ell}\})] & \text{if } n \neq z_{\mathcal{R}}^{\psi} \\ T_n(\psi) & \text{if } n = z_{\mathcal{R}}^{\psi} \text{ and } \psi \cap \mathcal{P}_{\mathcal{R}} = \phi \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

$$F_{\mathcal{R}}^{\psi}(n) = \begin{cases} \max[R_n^{\psi}, F_{\mathcal{R}}^{\psi}(\text{Pred}(n, \mathcal{P}_{\mathcal{R}}^{\psi})) + \Delta_{\mathcal{P}_{\mathcal{R}}^{\psi}}(n, \text{Pred}(n, \mathcal{P}_{\mathcal{R}}^{\psi}), n; \{\tau_{\ell}\})] & \text{if } n \neq x_{\mathcal{R}}^{\psi} \\ R_n^{\psi} & \text{if } n = x_{\mathcal{R}}^{\psi} \end{cases} \quad (11)$$

that has the potential to reach the destination. Hence, the starting point for computing the above time is taken as 0 (the failure instant) if the failure affects the primary path³.

After receiving a failure notification, a node n starts to reconfigure its switches and completes it by time R_n^{ψ} . A backup connection routed along the node n cannot cross the node until the reconfiguration is complete. If t denotes the earliest time at which the connection crosses the immediate predecessor node of n on the backup path, then the earliest time at which the backup connection will cross node n is given by $F_{\mathcal{R}}^{\psi}(n) = \max(R_n^{\psi}, t + \tau_{\ell})$, where τ_{ℓ} is the propagation delay of the link connecting the node and its predecessor on the backup path. The reconfiguration begins at node $x_{\mathcal{R}}^{\psi}$ and the earliest time at which the connection will cross this node is R_n^{ψ} , where $n = x_{\mathcal{R}}^{\psi}$. The recursive way of computing the earliest crossover time of the connection through n in its backup path is shown in Equation 11.

5 Performance Evaluation

The performance of the Diversion protection developed in this paper is compared against the (strict) FDPP and CSLP protection strategies on the NSFNET, ARPANET, and 8×2 networks. The topology of the networks are shown in Figure 7. The 8×2 network is chosen specifically because the fixed backup path for every link has a path length of 2, which is the best possible choice for a link protection strategy.

Every link in the network employs two uni-directional fibers, each consisting of sixteen wavelengths and eight time slots per wavelength. Thus, a total of 128 channels constitute a link in each direction. All the nodes in the network are assumed to groom traffic on a wavelength and no wavelength conversion is available at any nodes⁴. This paper considers only single-link failure scenarios in the network; and are modeled SRLG failures with one (bi-directional) link in each group. Upon an SRLG failure, fibers in both directions fail. The network is assumed to have at most one SRLG failure at any given instant.

The links are assumed to have a propagation delay of 5 ms, with 4 ms delay for electronic processing at each node, 4 ms delay for detecting a failed link, 20 ms for reconfiguring the switches at a node. Although the assumption of uniform propagation delay on NSFNET and ARPANET may not reflect that of the corresponding real-life networks, the generic conclusions that are derived from the performance results provide sufficient insight into the working of the various protection strategies.

Failure recovery time. In order to get an estimate of the recovery time of a connection, the networks are simulated independently for each protection strategy. The worst-case failure recovery time for a connection is

³When a link fails, the information after the failure point in the link may still continue to propagate to the next node. However, the propagation delay from the failure point to the first node of the last surviving segment is not taken into account here.

⁴The capacity information on a link is represented as a vector with 16 entries, each corresponding to a wavelength.




Figure 7. Network topologies considered for performance evaluation.

computed as the maximum recovery time among all the failures under which the connection would be reconfigured. The average and standard deviation of the worst-case recovery time of a connection for various protection strategies are shown in Table 3.

Table 3. Average and standard deviation of recovery times (in milliseconds).

Networks	FDPP		CSLP		Diversion	
	Average	Std. Dev.	Average	Std. Dev.	Average	Std. Dev.
NSFNET	53.7	7.2	51.3	5.5	50.6	5.3
ARPANET	61.6	14.3	53.9	8.5	54.6	9.0
8×2	62.3	23.6	38.0	0.0	48.9	12.2

It is observed that the recovery time for the Diversion strategy is closer to the CSLP than to FDPP. In the case of NSFNET network, the average recovery time of a connection under Diversion is smaller than that of CSLP because the average backup path length of a link is significant in comparison to the diversion path for a link to a particular destination. It is also observed that the standard deviation of the recovery time under Diversion is also significantly reduced as compared to FDPP, thus eliminating the dependence on primary path length significantly. Note that the standard deviation of the recovery time for 8×2 network under CSLP is 0 ms, because every link has a backup path length of two, hence all the recovery times are identical.

Backup path length. Table 4 shows the average backup path length of a connection under a failure. It is observed that the performance of the Diversion achieves a trade-off between the performance of FDPP and CSLP.

Table 4. Average backup path length of connections.

Networks	Average backup path length		
	FDPP	CSLP	Diversion
NSFNET	3.25	4.67	3.98
ARPANET	3.77	5.19	4.52
8×2	3.44	4.03	3.56

Blocking performance. Fig. 8 shows the blocking performance for the different protection strategies in the three networks. It is observed that the FDPP approach performs better than CSLP approach in NSFNET and ARPANET networks with Diversion achieving a trade-off between the two strategies. However, in 8×2 network, CSLP performs better than FDPP and Diversion strategies. The reason for such a behavior is that a link is as a backup only for a few other links in CSLP as opposed to FDPP and Diversion. In order to route a primary

connection through a link, the capacity must be available under no failure scenario and for every failure scenario for which the link is used as a backup. As the number of such failure scenarios under which a link is used as a backup is lesser in CSLP, more connections are accepted. This effect is prominent specifically when more wavelengths are employed without wavelength conversion capability. The performance results presented in [17] for a single wavelength system with 128 channels does not show this behavior for the same network. Hence, it is possible that link protection strategies working at the granularity of a connection may out-perform even certain path protection strategies.

Effective network utilization. We compute the resources utilized in the network through *effective network utilization* [18]. A request \mathcal{R} for capacity $c_{\mathcal{R}}$ that is routed along a path with a hop length of H utilizes $c_{\mathcal{R}} \times H$ capacity in the network. However, its effective utilization is only $c_{\mathcal{R}} \times H_s$, where H_s is the shortest path length between the source and destination of the connection. The effective network utilization at any given instant of time is then computed as the sum of the effective utilization of all requests running in the network at that time normalized to the total network capacity. It is to be noted that the effective utilization is computed over only the accepted requests, while the offered load is computed over all the requests.




Figure 8. Blocking performance.

Fig. 9 shows the effective network utilization for the three networks. It is observed that Diversion achieves the desired trade-off in the NSFNET and ARPANET networks, however, performs the worst in 8×2 network. The performance of CSLP and FDPP approaches are almost identical with respect to this metric, although their blocking performance at the corresponding arrival rate is significantly different as the blocking probability values are too low to have significant effect on the network utilization.




Figure 9. Effective network utilization.

Average primary path length. The average primary path length metric, shown in Fig. 10, depicts the effectiveness of the routing algorithms in finding longer routes to accommodate connections. As the network load increases, the average path length of accepted connections is expected to increase for any dynamic path selection strategy. However, beyond a certain network load, longer hop connections tend to get blocked more than the shorter hop

connections. Hence, the average path length of a connection reduces. Such a trend is seen for all the networks under all protection strategies. In particular, such a trend is not very prominent in 8×2 network under CSLP as the network efficiently accommodates the connections using shorter backup paths for links. It is interesting to note that CSLP (a link protection strategy) has a better blocking performance than FDPP (a path protection strategy) while maintaining a lower value for average primary path length of accepted connections. While one could argue that a link protection scheme might reject lesser number of connections by rejecting those connections that require longer hop length, the following metric, average shortest path length, demonstrates that is not the case.

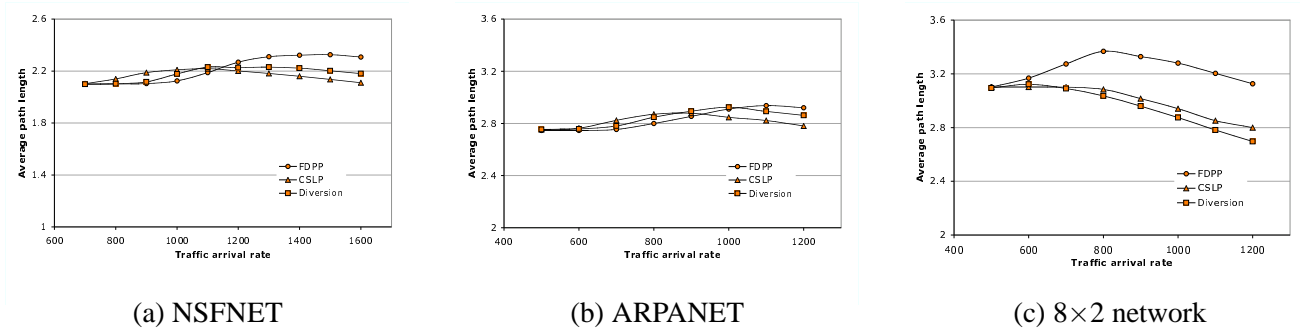


Figure 10. Average primary path length of accepted connections.

Average shortest path length. The average shortest path length metric, shown in Fig. 11, depicts the average shortest path length between the source and destination of accepted connections. Note that, the shortest path length between two nodes in a network does not change with arrival rate. Under low arrival rates, the blocking probability is very low, hence the value of this metric reflects the characteristics of the network for a particular traffic arrival rate. If the routing algorithms treat the connections in a fair manner, then this metric must remain a constant as the network load increases. However, connections that require longer path lengths are likely to get rejected more with increasing network load, the generic trend for any routing algorithm is that the average shortest path length decreases with increasing load. It is observed from Fig. 11 that the performance of Diversion is in between that of FDPP and CSLP for NSFNET and ARPANET network. For the 8×2 network, FDPP and CSLP have similar performance indicating that both approaches are equally fair in treating requests of varying connection lengths. Hence, for the 8×2 network with 16 wavelengths and 8 time slots per wavelength, CSLP performs better than FDPP. However, such a behavior may not emerge prominently when full-wavelength conversion is employed (or when one considers an electronic network where a concept similar to wavelength division multiplexing is not considered on a link.).

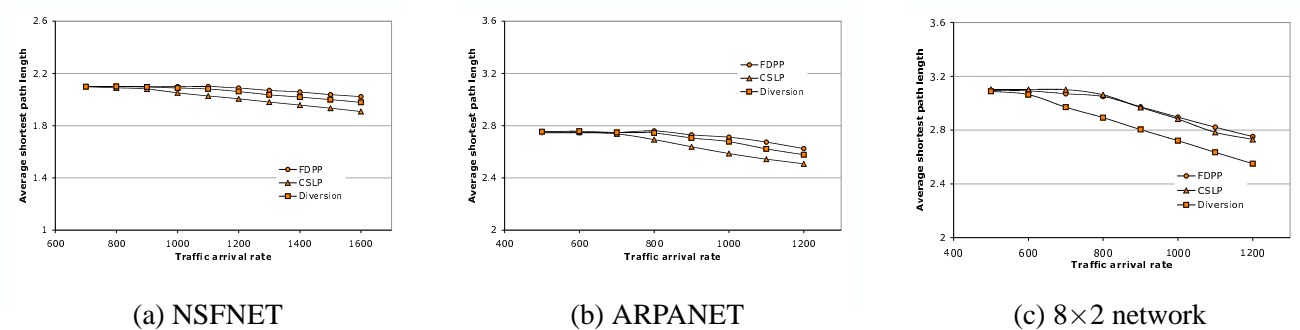


Figure 11. Average shortest path length of accepted connections.

6 Conclusion

This paper evaluates and compares three path protection strategies that are aimed at achieving a trade-off between network utilization and recovery time. A connection establishment procedure is developed to route connections using all the three protection strategies and their effectiveness is evaluated in three networks.

Based on the simulation studies, it may be summarized that the Diversion protection strategy performs better than FDPP strategy in terms of connection recovery times, while compromising on the blocking performance and network utilization. The Diversion technique performs better than CSLP, with respect to blocking performance and utilization, in those networks where FDPP also performs better than CSLP.

Acknowledgment

The research presented in this paper is supported in part by National Science Foundation, under grants 0325979, 0435490, and EES-0333046.

References

- [1] J. Jue and G. Xiao, "An adaptive routing algorithm for wavelength-routed optical networks with a distributed control scheme," in *Proceedings of the Ninth International Conference on Computer Communications and Networks*, Las Vegas, Nevada, USA, October 2000, pp. 192–197.
- [2] H. Zang, L. Shasrabuddhe, J. P. Jue, S. Ramamurthy, and B. Mukherjee, "Connection management for wavelength-routed WDM networks," in *Global Telecommunications Conference, GLOBECOM'99*, Rio de Janeiro, Brazil, 1999, vol. 2, pp. 1428–1432.
- [3] A. Mokhtar and M. Azizoglu, "Adaptive wavelength routing all-optical networks," *IEEE Transactions on Networking*, vol. 6, no. 2, pp. 197–206, April 1998.
- [4] H. Zang, J. Jue, L. Sahasrabuddhe, R. Ramamurthy, and B. Mukherjee, "Dynamic lightpath establishment in wavelength-routed WDM networks," *IEEE Communications*, pp. 100–108, September 2001.
- [5] S. Ramamurthy and B. Mukherjee, "Fixed alternate routing and wavelength conversion in wavelength-routed optical networks," in *Proceedings of the Global Telecommunications Conference, GLOBECOM'98*, Sydney, Australia, November 1998, pp. 2295–2303.
- [6] E. D. Lowe and D. K. Hunter, "Performance of dynamic path optical networks," in *IEE-Proceedings of Optoelectronics*, August 1997, pp. 235–239.
- [7] L. Li and A. K. Somani, "Dynamic wavelength routing using congestion and neighborhood information," *IEEE Transactions on Networking*, vol. 7, no. 5, pp. 779–786, October 1999.
- [8] X. Zhang and C. Qiao, "Wavelength assignment for dynamic traffic in multi-fiber WDM networks," in *7th International Conference on Computer Communication and Networks*, Lafayette, LA, USA, October 1998, pp. 479–485.
- [9] B. Wen and K. M. Sivalingam, "Routing, wavelength and time-slot assignment in time division multiplexed wavelength-routed optical WDM networks," in *Proceedings of IEEE INFOCOM'02*, New York, NY, USA, June 2002, pp. 1442–1450.
- [10] K. Zhu and B. Mukherjee, "Traffic grooming in an optical WDM network," in *IEEE International Conference on Communications*, June 2001, pp. 721–725.

- [11] R. Srinivasan, "MICRON: A framework for connection establishment in optical networks," in *Proceedings of OPTICOMM*, Dallas, TX, USA, October 2003, pp. 139–150.
- [12] W. D. Grover, *Mesh-based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking*, Prentice Hall Publishers, New Jersey, USA, 2003.
- [13] K. Sathyamurthy and S. Ramasubramanian, "Benefits of link protection at connection granularity," in *Proceedings of IEEE International Conference on Broadband Networks (BROADNETS)*, October 2004, pp. 300–309.
- [14] S. Ramasubramanian, "On failure dependent protection in optical grooming networks," in *IEEE International Conference on Dependable Systems and Networks (DSN)*, Florence, Italy, June–July 2004, pp. 475–484.
- [15] M. T. Fredrick and A. K. Somani, "A single-fault recovery strategy for optical networks using subgraph routing," in *Proceedings of the 7th IFIP Working Conference on Optical Network Design and Modelling*, Budapest, Hungary, February 2003, pp. 327–346.
- [16] M. Patel, R. Chandrasekaran, and S. Venkatesan, "A comparative study of restoration schemes and spare capacity assignments in mesh networks," in *Proceedings of 12th International Conference on Computer Communications and Networks*, October 2003, pp. 399–404.
- [17] S. Ramasubramanian and K. Sathyamurthy, "Supporting multiple protection strategies in optical networks," *Technical Report, Department of Electrical and Computer Engineering, University of Arizona*, November 2004.
- [18] R. Srinivasan and A. K. Somani, "Request-specific routing in WDM grooming networks," in *Proceedings of IEEE International Conference on Communications (ICC 2002)*, New York, NY, USA, April 2002, pp. 2876–2880.