Supporting Multiple Protection Strategies in Optical Networks

Srinivasan Ramasubramanian
Department of Electrical and Computer Engineering
University of Arizona, Tucson, AZ 85721
srini@ece.arizona.edu

***Abstract*—This paper develops a framework to support multiple protection strategies in optical networks, which is in general applicable to any connection-oriented network. The capacity available on a link for routing primary and backup connections are computed depending on the protection strategy. The paper also develops a model for computing service outage and failure recovery times for a connection where notifications of failure location are broadcast in the network. The effectiveness of employing multiple protection strategies is established by studying the performance of three networks for traffic with four types of protection requirement.**

Keywords: Optical networks, Dynamic routing, Path protection, Link protection, Multiple protection strategies

I. 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), where each wavelength is shared by connections with much lower capacity like 155 Mbps (OC-3) or 622 Mbps (OC-12). As optical processing and buffer technologies are not mature enough 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 connections between source-destination pairs on-demand and release them when a connection is no longer needed.

Resiliency against link and node failures is critical in optical networks due to the high data rates. *Protection* schemes guarantee 100% recovery by dedicating resources to connections in case of failure. *Restoration* schemes do not dedicate resources, hence providing guarantees on successful or timely reconfiguration of a connection after failure is difficult. Hence, protection schemes are preferred for high-priority traffic; while restoration schemes are employed for low-priority traffic.

Connection establishment in an optical network consists of two steps: *path selection* and *wavelength assignment*. Path selection refers to selecting a path from source to destination based on certain criteria. Wavelength assignment refers to assigning one or more channels depending on the requirement of the call on every link of the chosen path. 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 the primary path fails. We refer to the time duration for which a destination does not receive data from the source due to the failure as the *service outage time*. We refer to the time difference between the instant of failure and the instant at which a destination starts to receive data along the backup path as the *failure recovery time*.

A set of links may share resources, a duct or conduit through which they are laid out, whose failure would result in the failure of multiple links. Such failures are modeled as Shared Risk Link Group (SRLG) failures[1]. Typically, the objective of the network operation is to protect connections against any SRLG failure.

While there have been several protection strategies developed in the literature (see [1], [2], and references therein), it is often the case that a work considers only one kind of protection strategy to be employed in the network. However, the protection requirements of traffic may vary significantly – some may require fast and guaranteed recovery, some require guaranteed recovery but can tolerate higher outage time, while others may require no protection at all. Clearly, no single protection strategy can satisfy such a wide range of protection requirements, as all of them trade-off capacity efficiency with service outage (or failure recovery) time. Hence, a network must support multiple protection strategies for effective operation.

There are two major issues involved in employing multiple protection strategies in a network. First, it is essential to maintain link availability information in a consistent manner such that the available capacity under any kind of protection scheme may be computed easily for routing purposes. Second, every connection in the network must be recovered in exactly one way, thereby avoiding any possible contention between two protection strategies to recover a connection. In order to overcome the above two issues, a framework to support multiple protection strategies is required, where an individual connection is protected using an approach that is best-suited for its requirements.

In this paper, we develop a framework to support multiple protection strategies, specifically identify a consistent mechanism of maintaining the available capacity information across the network. In order to achieve efficient utilization of network

[1]Single link failures are special cases of SRLG failures where each group contains one link.

resources, all protection strategies are employed at the granularity of a connection. Upon a failure, the nodes attached to the failed link send out a failure notification message indicating the failure location. The nodes in the network independently reconfigure their switches corresponding to the failure scenario. We develop a method to compute the service outage time and failure recovery time for individual connections based on the said failure notification and recovery model. We establish the significance of supporting multiple protection strategies by comparing the performance to supporting any one protection strategy using extensive simulations.

The remainder of the paper is organized as follows: Section II provides a taxonomy of the existing protection strategies. Section III describes the network model, information management of links and paths, path selection and wavelength assignment, traffic characteristics, and failure recovery procedure. Section IV describes the path selection process under different protection strategy. The computation of service outage time and failure recovery time of a connection is described in Section V. The results of the performance study on various protection strategies are described in Section VI. Section VII concludes the paper.

II. Taxonomy of protection schemes

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

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 whether they require failure location information or not. If a backup path is assigned without the precise knowledge of the link failure, then it is referred to as *failure-independent path protection* (FIPP). A connection is reconfigured to the same backup path under any failure that affects the primary path. Hence, the backup path must be link-disjoint with the primary path. Fig. 1 shows an example network where the primary connection from node 1 to 4 is established along the path 1–2–3–4. Under FIPP, the backup path for the connection is 1–5–6–4.

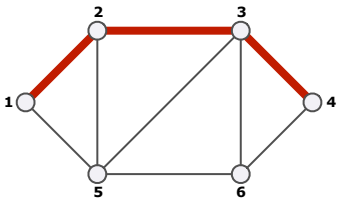

Fig. 1. Example network where a primary connection is established from node 1 to 4 along the path 1–2–3–4.

FIPP strategy has two major drawbacks. First, the primary and/or backup path lengths are usually longer compared to failure dependent schemes as the paths have to be link-disjoint. Second, FIPP strategy does not work under certain multiple failures as SRLG-disjoint paths may not exist between two nodes, i.e. no single SRLG affects the primary and backup paths together, even though no single SRLG failure may disconnect the network.

If a connection is assigned more than one backup path depending on the failure, then it is referred to as *failure-dependent path protection* (FDPP). For the example considered in Fig. 1, the backup path for every link failure is shown in Fig. 2. If the primary path (and wavelength assignment) of the connection is valid under any failure scenario that does not affect the primary path, then the connection need not be reconfigured under such scenarios. A path protection strategy in which connections are reconfigured only when a failure affects the primary path is referred to as a *strict FDPP* strategy [3].

In order to improve the blocking performance, connections may be reconfigured even if a failure does not affect its primary path. One such approach is the L+1 protection strategy [4]. Under L+1 strategy, a network is decomposed into L+1 sub-graphs, where L denotes the number of links in the network. One sub-graph corresponds to normal operation while the others correspond to the network under a distinct single-link failure scenario. A connection that requires protection is accepted only if it can be accommodated in each of the L+1 networks. The connection is attempted independently on each network, hence the path (and wavelength assignment) assigned to a connection under a failure scenario that does not affect its primary path (the path in the network with no failures) may not be identical, thus requiring a reconfiguration under such failures. The above path protection strategy is referred to as *flexible FDPP* in this paper, and may be extended to SRLG failures as well.

Link protection. Link protection schemes route a connection around the 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. 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. The primary fiber is used for routing working connections while the spare fiber is used only when a failure occurs. 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 routed along a link is automatically protected. Link protection at the connection level offers significant improvement when traffic requires different levels of protection [5].

If link protection is employed independently for every connection, then the responsibility is on the connection establishment procedure to provide a consistent wavelength assignment across the primary and backup paths. Note that under link protection strategy, the backup path is obtained by simply replacing the failed link with the backup path of the link. The wavelength assignment of the connection on the links not affected by the failure remain the same. We refer to the link protection at the granularity of a connection as *Connection Switched Link Protection* (CSLP).

Segmented protection. A trade-off between the recovery time and network utilization (or blocking performance) may be achieved by employing segmented path protection strategies.

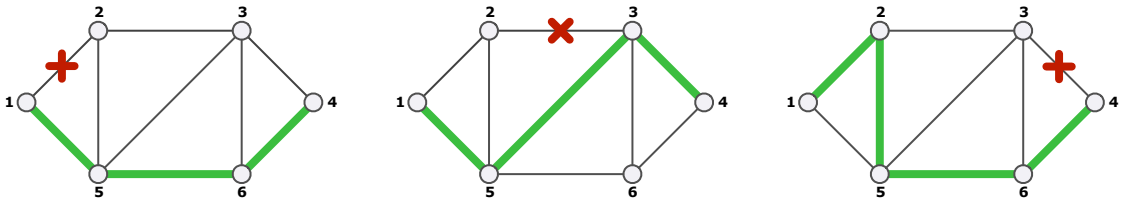

(a) After link 1–2 failure (b) After link 2–3 failure (c) After link 3–4 failure

Fig. 2. Backup paths using failure dependent path protection (FDPP) strategy. For strict FDPP, the primary path is valid under all other failure scenarios. For flexible FDPP, a backup path is provided even for failures that do not affect the primary path.

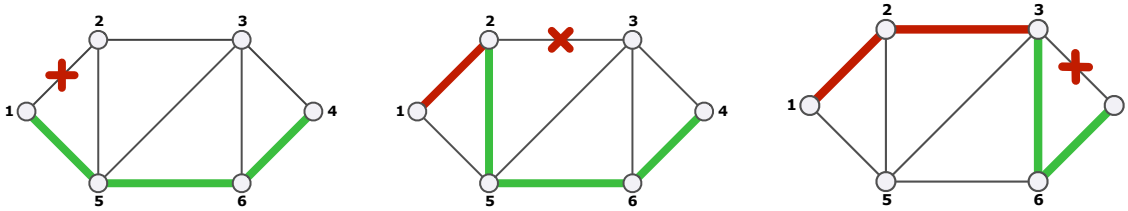

(a) After link 1–2 failure (b) After link 2–3 failure (c) After link 3–4 failure

Fig. 3. Backup paths using the Diversion strategy.

In this approach, the primary path is divided into multiple segments and the segments are protected individually. It is worth noting that the link protection and path protection are extreme cases of segmented protection where the former treats each link as a segment while the latter treats the entire path as a segment.

Segmented protection may be implemented in several ways [6], one of which is to divert the connection from the node before the failed link directly to the destination, referred to as Diversion [7]. Unlike CSLP that employs a backup path for a failed link, Diversion finds a path from the node attached to the failed link to the destination of a connection. Such a path may be identified only after the arrival of the request. For the example considered in Fig. 1, the backup paths obtained under Diversion are shown in Fig. 3. It may be observed that the Diversion strategy has the characteristics similar to that of a path (link) protection when the failures are close to the source (destination).

Backup multiplexing. In order to achieve efficient utilization, multiplexing of resources across multiple backup paths and 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 primary paths 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 already existing primary connection occupying that resource.

III. Framework to Support Multiple Protection Strategies

We consider an optical network where links have multiple fibers, multiple wavelengths per fiber, and multiple time slots per wavelength. Let $\mathcal{N}$ denote the set of nodes and $\mathcal{L}$ denote the set of 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 Group (SRLG) failures 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.

We assume in this paper, for sake of clarity, that all the nodes in the network have full grooming capability. Hence, the capacity information of a link is presented as a scalar. If the nodes do not employ wavelength conversion, then the available capacity information may be represented as a vector. If the nodes employ heterogeneous switching architectures, then the capacity on a link may be presented as a matrix [8].

The notations employed to compute paths using multiple protection strategies are shown in Table I.

A. Available capacity on a link and path

Let S_ℓ and P_ℓ denote the total capacity and capacity occupied by primary connections on link ℓ, respectively. Let G_ℓ^ψ denote the capacity on link ℓ that is currently occupied by working connections that would fail in case of the SRLG failure ψ. The capacity used by the connections affected by the failure become available, which may be assigned for backup connections. Incorporating this information into path selection enables primary-backup multiplexing, often referred to as "stub-release" in the literature. Let B_ℓ^ψ denote the number of backup channels required on link ℓ in case of an SRLG failure ψ.

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 available capacity on a path is obtained by combining the information of the links in the path as:

$$A_\mathcal{P} = \min_{\ell \in \mathcal{P}} A_\ell \tag{1}$$

If the variables are vectors, then the available capacity vector of the path is computed as the element-wise minimum of the available capacity of the links in the path.

TABLE I

COMPREHENSIVE LIST OF NOTATIONS WITH COMMENTS EMPLOYED TO COMPUTE PATHS UNDER MULTIPLE PROTECTION STRATEGIES.

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.
$\mathcal{P}(s, d, \psi, \{x_\ell\})$	Path computed dynamically from s to d by removing links in the set ψ and assigning $\{x_\ell\}$ as the capacity of the links.
ϕ	Null set.
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 $\psi \in \Psi$.
$\mathcal{Z}_\ell^\psi$	Backup path for link ℓ under failure ψ.
X_ℓ	Capacity available to route primary connection.
Y_ℓ^ψ	Capacity available to route backup connection under failure ψ.
$\mathcal{P}_\mathcal{R}$	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.
$\xi_\mathcal{R}(\ell)$	Wavelength assignment for request $\mathcal{R}$ on link ℓ under no failure.
$\xi_\mathcal{R}^\psi(\ell)$	Wavelength 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 $\mathcal{R}$.

B. Path selection and wavelength 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 on the links as $\{x_\ell\}$ is denoted by $\mathcal{P}(s, d, \psi, \{x_\ell\})$. We assume that the path selection is based on Dijkstra's algorithm that would select the shortest path (based on hop-count) among the available paths (paths that have sufficient capacity for routing the connection), referred to as Available Shortest Path (ASP) [9].

Let $\mathcal{R}$ denote a request for a connection from source $s_\mathcal{R}$ to destination $d_\mathcal{R}$ for capacity $c_\mathcal{R}$. The request is assigned a primary path and, if required, one or more backup paths. We refer to the request as *connection* if it is accepted. Let $\mathcal{P}_\mathcal{R}$ denote the set of links in the primary path of the connection. The connection 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.

Let $\xi_\mathcal{R}(\ell)$ denote the wavelength assignment on link ℓ for the primary path of a connection and $\xi_\mathcal{R}^\psi(\ell)$ the wavelength assignment on link ℓ under failure ψ.

C. Traffic requirements

We assume that requests that arrive in the network have certain Quality-of-Protection (QoP) requirements based on which we classify the requests into one of the following four types.

- *Type 1* – Guaranteed protection with less than 50 ms service outage time.
- *Type 2* – Guaranteed protection with no outage time requirement; connection reconfiguration probability less than 0.3.
- *Type 3* – Guaranteed protection with no outage time and reconfiguration requirements.
- *Type 4* – No protection (from failures that affect the primary path).

Type-1 traffic requires stringent requirement on outage time. Type-2 traffic can at times tolerate the overhead involved in having a higher outage time. A high outage time in the optical layer may lead to several undesired reactions in the electronic layers above, hence it is of importance to minimize this effect. Assume that a link in a network is unavailable for a few hours in a month for maintenance[2], and the connections will have to be reconfigured to backup paths before maintenance. A connection may have a requirement of a certain number of instances it can tolerate over a certain fixed period of time (say a month). This value is translated into the number of failures under which the connection will be reconfigured. A probability of 0.3 implies that the number of failures for which the connection will be reconfigured must be less than 0.3 times the total number of failures in the network. Type 4 traffic is the lowest priority and do not require any protection against link failures that affect its primary path. However, the connections may not be dropped if the failure does not affect its primary path.

A request is assigned a protection strategy based on its QoP requirements. Based on our earlier study of the individual schemes [3], [5], we assign the protection type to the traffic in a static manner: Type 1 connections are protected through

[2]Maintenance schedule may be treated as failures with known failure time and duration.

link protection; Type 2 connections by strict FDPP, Type 3 connections by flexible FDPP.

D. Failure recovery

A link failure is identified by the nodes attached to the link. The nodes send out broadcast messages indicating the location of the failure. Upon reception of the failure notification, every node reconfigures its switches corresponding to the failure scenario. Note that as the network employs protection, the paths on which the connections will be routed upon a failure are known apriori. Hence, it is not necessary for the source node to initiate a request for reconfiguring the connection to its backup path. In addition, we assume that the failure notification messages are given highest preference, hence the delay they encounter is primarily the propagation delay (and some additional delay to process the message).

IV. Path Computation Under Multiple Protection Strategies

In this section, we describe the computation of primary and backup paths for under different protection strategies. When the network does not have failures, the capacity assigned to primary connections on a link is upper bounded by the available capacity on the link. Therefore,

$$P_\ell \leq S_\ell \tag{2}$$

On a failure ψ, the network will be able to re-assign requests to their backup connections if the following condition is satisfied:

$$P_\ell - G_\ell^\psi + B_\ell^\psi \leq S_\ell \tag{3}$$

The above inequality must be obeyed by any routing and wavelength assignment scheme for all SRLG failures if the network must be resilient to any single SRLG failure. From the above, the available capacity on a link when there are no failures is computed as:

$$A_\ell = S_\ell - P_\ell \tag{4}$$

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

$$A_\ell^\psi = S_\ell - P_\ell - B_\ell^\psi + G_\ell^\psi \tag{5}$$

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

The following subsections describe the path computation under each protection strategy. 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 Fig. 4. The readers are recommended to refer to this table for the corresponding protection strategy under consideration.

A. Failure Independent Path Protection (FIPP)

FIPP provides fixed primary and backup paths that are link-disjoint. The connection established along the primary path will be reconfigured to the backup path only when a failure affects a link on the primary path. Under any other failure, the connection will remain unaffected. Hence, the available capacity to route primary connection (X_ℓ) must be available when the network has no failures and under those failures that do not affect the primary path. As the primary path ($\mathcal{P}_\mathcal{R}$) is known apriori, the set of failures that affect the primary path ($\Psi_\mathcal{R}$) is also known.

The capacity allocated for the backup path must be available only in those failures that will affect the primary path. Note that the backup path is the same under all failures that affect the primary path. In addition, if the wavelength assignment must also be the same on all the links, then the available capacity to route the backup connection (Y_ℓ^ψ) is computed by considering the available capacity under all the failure scenarios in which the connection will be reconfigured. Hence, the wavelength assignment on the backup path need to be computed only once.

A successful connection would be assigned one primary path and $|\Psi_\mathcal{R}|$ backup paths such that the backup path (along with the wavelength assignment) for all failures in $\Psi_\mathcal{R}$ are identical. Hence, the maximum number of distinct backup paths (along with wavelength assignment) is one.

B. 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 of the 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 on the primary path; or (2) a failure in the network does not involve any link on the primary path, however the capacity assigned to the connection on the primary path is also assigned to the backup paths of some other connections that are affected by the failure. The latter requires reconfiguration of a connection to its backup path even if a failure does not affect the primary path of the connection. Such a reconfiguration can lead to a *domino effect* requiring network-wide reconfiguration.

Strict FDPP. Under strict FDPP, the connections may be reconfigured only under those failures that affect the primary path. Hence, the capacity assigned to a connection on the primary path must be available under all failure scenarios that does not affect the primary path. The computation of the primary path requires the knowledge of the available capacity on links (to compute the shortest available path). The computation of available capacity on a link in turn requires the knowledge of the primary path. To avoid such a mutual dependence, the computation of the available capacity on a link to route a primary

Metrics	FIPP	Strict FDPP	Flexible FDPP	CSLP	No protection
X_ℓ	$\min\left(A_\ell, \min_{\psi\in\Psi\setminus\Psi_\mathcal{R}} A_\ell^\psi\right)$	$\min\left(A_\ell,\ \min_{\psi\in\Psi} A_\ell^\psi\right)$	A_ℓ	$\min\left\{A_\ell,\ \min_{\psi\in\Psi} A_\ell^\psi,\ \min_{\psi\in\Psi_\ell} R_\ell^\psi\right\}$	$\min\left(A_\ell,\ \min_{\psi\in\Psi} A_\ell^\psi\right)$
$\mathcal{P}_\mathcal{R}$	Fixed	$\mathcal{P}(s_\ell, d_\ell, \phi, \{X_\ell\})$			
$\Psi_\mathcal{R}$	$\{\psi \mid \mathcal{P}_\mathcal{R} \cap \psi \neq \phi\}$	$\{\psi \mid \mathcal{P}_\mathcal{R} \cap \psi \neq \phi\}$	Ψ	$\{\psi \mid \mathcal{P}_\mathcal{R} \cap \psi \neq \phi\}$	
Y_ℓ^ψ	$\min_{\psi\in\Psi_\mathcal{R}} A_\ell^\psi$	A_ℓ^ψ			
$\mathcal{Z}_\ell^\psi$	–	–	–	Fixed or $\mathcal{P}(s_\ell, d_\ell, \psi, \{Y_\ell^\psi\})$	–
$\mathcal{Y}_\ell^\psi(\mathcal{R})$	–	–	–	–	–
R_ℓ^ψ	–	–	–	$\prod_{\ell'\in\mathcal{Z}_\ell} A_{\ell'}^\psi$	–
$\mathcal{P}_\mathcal{R}^\psi$	Fixed	$\mathcal{P}(s_\ell, d_\ell, \psi, \{Y_\ell^\psi\})$	$\mathcal{P}(s_\ell, d_\ell, \psi, \{Y_\ell^\psi\})$	Replace links affected by failure with backup paths.	ϕ
Maximum distinct backup paths	1	$\lvert\Psi_\mathcal{R}\rvert$	$\lvert\Psi\rvert$	$\lvert\Psi_\mathcal{R}\rvert$	0

Fig. 4. Comprehensive list of capacity and path computation for various protection strategies.

connection (X_ℓ) is computed in a conservative manner (see Fig. 4) by ensuring that the capacity is available under all failure scenarios.

The path selection strategy in the network then 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.

Flexible FDPP. Recall that the flexible FDPP approach is equivalent of treating the network as L+1 sub-graphs (one with no failure and others with one distinct link failed) [4] and accepting the request only when a successful path and wavelength assignment is available in all the sub-graphs. This approach is shown to be capacity efficient, but requires a connection to be reconfigured under a very high number of failures [3].

A connection may be rerouted even when a failure does not affect its primary path. Hence, the capacity available to route a primary connection is computed by simply taking the capacity available when the network does not have any failures. Clearly, the capacity assigned for the primary path may not be available under a failure scenario, hence the connection must be provided a backup path for every SRLG failure in the network. Hence, the set of failures for which a backup path computed is the same as the set of failure scenarios in the network.

Irrespective of the way in which the primary path is computed, the backup path has 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.

The path selection approach described above takes into account only 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 of available capacity under a failure.

C. Connection Switched Link Protection (CSLP)

For link protection at the connection level, a connection is re-routed around the failed link. A link ℓ is assumed to have a backup path for every failure ψ that affects it, denoted by $\mathcal{Z}_\ell^\psi$. The wavelength 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 $\mathcal{Z}_\ell^\psi$ denote the backup path for link ℓ under failure ψ. The capacity available on the backup path $\mathcal{Z}_\ell^\psi$ upon failure ψ, denoted by R_ℓ^ψ is computed as shown in Fig. 4.

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 one of its backup paths depending on the failure. 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 a 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 availabile capacity on the backup path of the link under the failure scenarios affecting the link (R_ℓ^ψ).

We note that the backup path for a link ℓ under a failure ψ Z_ℓ^ψ may be computed dynamically based on the network status, $Z_\ell^\psi = \mathcal{P}(s_\ell, d_\ell, \psi, \{A_\ell^\psi\})$, by simply not considering the links that are affected by failure ψ and computing the available shortest path between the nodes connected by link ℓ. The computation of the backup path for a link is independent of the request, hence may be performed before request arrival. In this paper, we assume that the backup path for a link under a failure is computed as the shortest path under the link failure; hence the backup path for a link under a failure does not change with network traffic.

The backup path for a connection under each failure scenario is obtained by simply replacing the failed link by the links in the backup backup path. A successful connection has one primary path and $|\Psi_\mathcal{R}|$ backup paths.

The wavelength assignment on the backup path must be consistent with the wavelength assignment of the primary path still intact. As the routing of primary connection has taken into account the availability of backup paths, a consistent wavelength assignment on the backup path is guaranteed to exist[3].

D. No Protection

Requests that do not require any protection need to be provided with only a primary path. Although the request does not require protection, a connection is provided that may not be removed unless a link in its primary path fails. The available capacity on a link to route the primary connection is computed similar to the strict FDPP strategy. The capacity occupied by the primary connection would be released whenever a failure affects the primary path. For that reason, the set of failures under which a reconfiguration is necessary ($\Psi_\mathcal{R}$) is computed as those failures that affect the primary path. As the connections do not need protection, no backup paths are assigned for any of the failures ($\mathcal{P}_\mathcal{R}^\psi = \phi, \ \forall \psi \in \Psi_\mathcal{R}$). When a failure $\psi \in \Psi_\mathcal{R}$ occurs, the capacity assigned to the connection will be released.

E. Connection establishment and release

We assume that the network is managed through a centralized control, or equivalently, the network employs link state protocol where every node has up-to-date network state information. The procedures for connection establishment and release is shown in Fig. 5 when the network supports multiple protection strategies. The connection establishment procedure takes as input the current network status and request. The protection scheme for the request is statically determined based on its requirements. The output of the connection establishment procedure is to provide a primary path and backup path(s), if necessary, along with wavelength assignment.

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 connection release procedure is similar to the Step 5 of connection establishment, except that the capacities are released instead of being assigned.

[3]Path pruning may be necessary to avoid looping in certain cases.

V. Computation of Service Outage and Failure Recovery Times

In this section, we compute the service outage and failure recovery times for a connection under single link failure scenario. The timing calculations are performed with the failure instant as the reference. The notations employed in computing the service outage time is shown in Table II.

Upon a link failure, the nodes connected to the failed link detect the failure. The failure is assumed to be detected due to the loss of a periodic "Hello" packets exchanged 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. Upon detecting the failure, the nodes broadcast a failure notification message. The sum of the time required for the node to prepare and transmit a 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, denoted by γ. If τ_ℓ denotes the propagation delay on link ℓ, then $\tau_\ell + \gamma$ denotes the delay experienced by a failure notification message on the link. Note that as the failure notification message will be converted 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. The nodes attached to the failed link broadcast failure notification independently. A node n 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 6. This notification time includes the time to detect the failure (α) and the minimum of the propagation delay from the nodes attached to the failed link to node n in the network after removing the failed link.

The nodes start to reconfigure their switches when they receive the notification. The worst-case time to reconfigure the switches at a node is assumed to be β. The time at which the reconfiguration will be completed at node n for failure ψ, denoted by R_n^ψ, is computed as shown in Equation 7.

Consider a connection established for a request $\mathcal{R}$ and a single link failure ψ. Let $z_\mathcal{R}^\psi$ denote the first node in the primary path such that no link in the primary 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 starts and ends 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. The above nomenclature is

Connection establishment procedure
Input:
1) Current network state.
2) Request $\mathcal{R}$ with a specific protection requirement.

Output:
1) Primary path $\mathcal{P}_\mathcal{R}$ and wavelength 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 wavelength 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).
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 wavelength assignment. Update the available capacity on each link to route a backup connection under failure ψ as Y_ℓ^ψ.
   - **FIPP:** It is sufficient to compute for one failure scenario as the computation of Y_ℓ^ψ takes into account the capacity availability under all failure scenarios in which the connection needs to be reconfigured.
   - **Strict/Flexible 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.
   - **No protection:** $P_\mathcal{R}^\psi = \phi$.
5) Update link capacities.
   Note: At this juncture, every request has been assigned: (1) a primary path $\mathcal{P}_\mathcal{R}$ with wavelength 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.

$$
\begin{aligned}
P_\ell\left[\xi_\mathcal{R}(\ell)\right] &\leftarrow P_\ell\left[\xi_\mathcal{R}(\ell)\right] + c_\mathcal{R} && \forall \ell \in \mathcal{P}_\mathcal{R} \\
G_\ell^\psi\left[\xi_\mathcal{R}^\psi(\ell)\right] &\leftarrow G_\ell^\psi\left[\xi_\mathcal{R}^\psi(\ell)\right] + c_\mathcal{R} && \forall \ell \in \mathcal{P}_\mathcal{R} \text{ and } \psi \in \Psi_\mathcal{R} \\
B_\ell^\psi\left[\xi_\mathcal{R}^\psi(\ell)\right] &\leftarrow B_\ell^\psi\left[\xi_\mathcal{R}^\psi(\ell)\right] + c_\mathcal{R} && \forall \ell \in \mathcal{P}_\mathcal{R}^\psi \text{ and } \psi \in \Psi_\mathcal{R}
\end{aligned}
$$

6) Exit.

Connection release procedure
Input: Request $\mathcal{R}$ which has already been accepted.
Steps:
1) Update link capacities.

$$
\begin{aligned}
P_\ell\left[\xi_\mathcal{R}(\ell)\right] &\leftarrow P_\ell\left[\xi_\mathcal{R}(\ell)\right] - c_\mathcal{R} && \forall \ell \in \mathcal{P}_\mathcal{R} \\
G_\ell^\psi\left[\xi_\mathcal{R}^\psi(\ell)\right] &\leftarrow G_\ell^\psi\left[\xi_\mathcal{R}^\psi(\ell)\right] - c_\mathcal{R} && \forall \ell \in \mathcal{P}_\mathcal{R} \text{ and } \psi \in \Psi_\mathcal{R} \\
B_\ell^\psi\left[\xi_\mathcal{R}^\psi(\ell)\right] &\leftarrow B_\ell^\psi\left[\xi_\mathcal{R}^\psi(\ell)\right] - c_\mathcal{R} && \forall \ell \in \mathcal{P}_\mathcal{R}^\psi \text{ and } \psi \in \Psi_\mathcal{R}
\end{aligned}
$$

Fig. 5. Generic connection establishment/release procedure.

applicable to a variety of protection strategies (link protection, path protection, *segmented* protection [10]).

When a link fails in the primary path, the destination continues to receive the information that is still in propagation on the last surviving segment. We first compute the latest time until when the destination continues to receive information sent on the primary path. Let $L_\mathcal{R}^\psi(n)$ denote the *latest crossing time* defined as the latest when the information of the connection $\mathcal{R}$ crosses node n, where n is a node in the last surviving segment of the primary path. Similarly, let $F_\mathcal{R}^\psi(n)$ denote the *earliest crossing time* defined as the earliest time when the information of 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 shown in Equation 10.

A node n starts to reconfigure its switch as soon as it receives the failure notification message. A node in the last surviving segment may receive the failure notification message either along the last surviving segment itself (if that is shortest path) or along a different path. If the latest crossing time at 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 crossing time at node n is given by $t + \tau_\ell$ if the node n has not already started its switch reconfiguration. Otherwise, it is given by T_n^ψ. Therefore, the minimum of the above two times defines the latest crossing time at node n, i.e. $L_\mathcal{R}^\psi(n) = \min(T_n^\psi, t + \tau_\ell)$. The recursive way of computing the latest crossing time at a node in the last surviving segment of the primary path is shown in Equation 8.

When a failure affects the primary path of a connection, the information that just crossed over failure point is the last bit of information that has the potential to reach the destination.

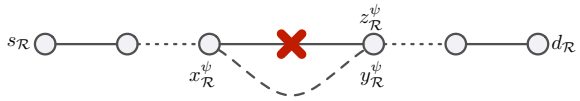

(a) CSLP

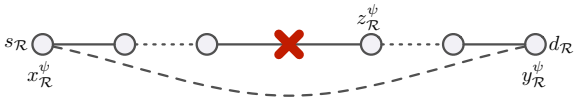

(c) FIPP, Strict and Flexible FDPP
(failure affects primary path)

(d) Flexible FDPP
(failure does not affect primary path)

Fig. 6. Illustration of $x^{\psi}_{\mathcal{R}}$, $y^{\psi}_{\mathcal{R}}$, and $z^{\psi}_{\mathcal{R}}$ for various protection strategies.

TABLE II

COMPREHENSIVE LIST OF NOTATIONS EMPLOYED IN SERVICE OUTAGE TIME COMPUTATION WITH COMMENTS.

Variables	Comments
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 $\mathcal{P}$.
T^{ψ}_n	Time taken for node n to receive the failure notification of failure ψ since the instant of failure.
R^{ψ}_n	Time to finish reconfiguration at node n from the instant of failure ψ.
τ_ℓ	Propagation delay on link ℓ.
$x^{\psi}_{\mathcal{R}}$	Node at which the request $\mathcal{R}$ is re-routed for failure ψ.
$y^{\psi}_{\mathcal{R}}$	Node at which the reconfigured segment of request $\mathcal{R}$ for failure ψ joins the primary path.
$z^{\psi}_{\mathcal{R}}$	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.
$L^{\psi}_{\mathcal{R}}(n)$	Latest time by which connection $\mathcal{R}$ can cross node n in its primary path after failure ψ.
$F^{\psi}_{\mathcal{R}}(n)$	Earliest time by which connection $\mathcal{R}$ can cross node n in its backup path after failure ψ.
$O^{\psi}_{\mathcal{R}}$	Outage time for request $\mathcal{R}$ under failure ψ.

$$T^{\psi}_n = \alpha + \min[\Delta^{\psi}(n_\ell(\psi), n; \{\tau_\ell + \gamma\}),\ \Delta^{\psi}(n'_\ell(\psi), n; \{\tau_\ell + \gamma\})] \tag{6}$$

$$R^{\psi}_n = T^{\psi}_n + \beta \tag{7}$$

$$L^{\psi}_{\mathcal{R}}(n) = \begin{cases} \min[T^{\psi}_n,\ L^{\psi}_{\mathcal{R}}(Pred(n, \mathcal{P}_{\mathcal{R}})) + \Delta_{\mathcal{P}_{\mathcal{R}}}(Pred(n, \mathcal{P}_{\mathcal{R}}), n; \{\tau_\ell\})] & \text{if } n \neq z^{\psi}_{\mathcal{R}} \\ T_n(\psi) & \text{if } n = z^{\psi}_{\mathcal{R}} \text{ and } \psi \cap \mathcal{P}_{\mathcal{R}} = \phi \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

$$F^{\psi}_{\mathcal{R}}(n) = \begin{cases} \max[R^{\psi}_n,\ F^{\psi}_{\mathcal{R}}(Pred(n, \mathcal{P}^{\psi}_{\mathcal{R}})) + \Delta_{\mathcal{P}^{\psi}_{\mathcal{R}}}(Pred(n, \mathcal{P}^{\psi}_{\mathcal{R}}), n; \{\tau_\ell\})] & \text{if } n \neq x^{\psi}_{\mathcal{R}} \\ R^{\psi}_n & \text{if } n = x^{\psi}_{\mathcal{R}} \end{cases} \tag{9}$$

$$O^{\psi}_{\mathcal{R}} = F^{\psi}_{\mathcal{R}}(y^{\psi}_{\mathcal{R}}) - L^{\psi}_{\mathcal{R}}(y^{\psi}_{\mathcal{R}}) \tag{10}$$

Hence, the starting point for computing the above time is taken as 0 (the failure instant) if the failure affects the primary path[4].

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

However, in flexible FDPP protection, a connection may be reconfigured even if a failure does not affect the primary path. In such cases, the source continues to transmit on the primary path until it receives a failure notification message. Therefore, the latest crossing time at the source is T^{ψ}_s.

After receiving a failure notification, a node n starts to reconfigure its switches and completes it by time R_n^ψ. The information of the connection sent along the backup path cannot cross the node n until the reconfiguration is complete. If t denotes the earliest crossing time of the immediate predecessor node of n on the backup path and τ_ℓ denotes the propagation delay of the link connecting the node and its predecessor on the backup path, then the earliest crossing time at node n is given by $t + \tau_\ell$ if node n has already completed its reconfiguration by that time. Otherwise, the earliest crossing time at node n is given by R_n^ψ. The maximum of the above two times defines the earliest crossing time at node n, i.e. $F_{\mathcal{R}}^{\psi}(n) = \max(R_n^\psi,\ t+\tau_\ell)$, . The reconfiguration begins at node $x_{\mathcal{R}}^{\psi}$ and the earliest crossing time at this node is is R_n^ψ, where $n = x_{\mathcal{R}}^{\psi}$. The recursive way of computing the earliest crossing time of the connection through n in its backup path is shown in Equation 9.

The difference between the first crossing time and the last crossing time at the node where the reconfiguration ends gives the service outage time, while the first crossing time at the node where the reconfiguration ends gives the failure recovery time for the connection.

VI. Performance Evaluation

The performance evaluation of the CSLP, Strict FDPP, Flexible FDPP, and MIXED protection strategies are carried out on three networks as shown in Fig. 7. The links in these networks 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 are still valid.

Every link in the network employs two unidirectional fibers each 128 channels per link (in each direction). All the nodes in the network are assumed to be full-grooming nodes. The network is assumed to have only single link failures. Due to the single link failure assumption, FIPP is not considered for performance evaluation as they are known to perform worse than FDPP schemes. Every link is provided with a fixed backup path over which all the connections routed through the link will be reconfigured upon a failure.

The networks employ available shortest path algorithm (ASP) [9] and first-fit wavelength assignment for all the considered protection methodologies. The available shortest path algorithm computes the shortest path among those paths that have sufficient resources for connection establishment.

Request arrival follows a Poisson process with rate λ and have an exponential holding time with unit mean. Every request has one channel capacity requirement. The source and destination of a request is assumed to be equally likely among all node pair combinations[5]. In addition, a connection has certain failure outage time requirements as mentioned in Section III-C and a connection is assigned a protection strategy based on the requirement. The simulations are performed in five rounds with 10,000 requests per round. The average of the values obtained over these five rounds are presented in this paper.

[5]Under non-uniform traffic, it is often difficult to quantify if a certain observation in the network is the effect of the given protection strategy or non-uniform traffic. Hence, the uniform traffic assumption.

A. Service outage time estimation for different networks

In order to get an estimate of the service outage time of a connection under each protection strategy, the networks are first simulated independently for each protection strategy. The worst-case service outage time for a connection is computed as the maximum outage time among all the failures under which the connection would be reconfigured. The average and standard deviation of the worst-case outage time of a connection for various protection strategies are shown in Table III.

The outage time under flexible FDPP is the highest. Strict and flexible FDPP schemes are path protection strategies, hence their outage times are higher than that of CSLP. The outage times for flexible FDPP is higher than strict FDPP as the connection will have to be reconfigured even under failures that does not affect the primary path. Flexible FDPP scheme is known to enforce reconfiguration in more failures than Strict FDPP. Table IV shows the average length of primary and backup paths along with the number of reconfigurations for network load at which the blocking probability is of the order of 10^{-3}. The total number of failures is the same as the number of links in the network. The number of reconfigurations for a connection under flexible FDPP is the highest, despite its low primary path length. For example, a connection established in an NSFNET network will require reconfiguration in an average of 15.9 link failures among the 22 link failures. For CSLP and Strict FDPP, a connection is not reconfigured unless a link in its primary path fails, hence the average number of reconfigurations is the same as the average primary path length.

The average outage times of Strict FDPP and Flexible FDPP are higher than CSLP by 4% and 6% for NSFNET network; 14.71% and 16.76% for ARPANET network; 63.42% and 68.16% for 8×2 network. The NSFNET and ARPANET have low connectivity, hence the outage times under CSLP is higher as the average backup path length for a link is significantly higher. This property can be observed in Table IV where the average backup path length with CSLP is higher than that of FDPP schemes. Hence, for these two networks, CSLP does not offer significant advantages in terms of outage times. In the 8×2 network, every link has a backup path length of 2. Hence, the outage time of CSLP is significantly lower compared to FDPP schemes. This is also the reason for a 0.0 ms standard deviation in the outage time observed for CSLP in Table III.

From earlier studies [3], [5], it is observed that Flexible FDPP has the best capacity utilization followed by Strict FDPP, while CSLP performs the worst. Hence, in networks where there is no significant difference in outage times between CSLP and FDPP approaches, FDPP (strict or flexible) is expected to perform better than CSLP.

B. Performance under mixed traffic

The performance of employing any one particular type of protection strategy (CSLP, strict FDPP, or flexible FDPP) for

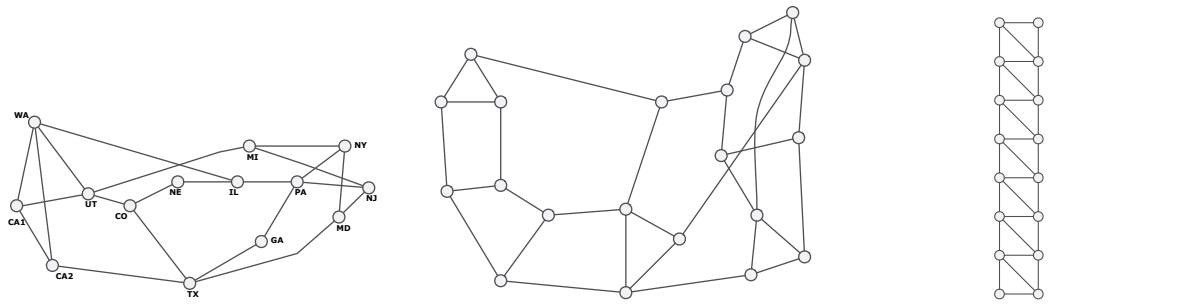

(a) NSFNET network (14 nodes, 22 links)

(b) ARPANET network (20 nodes, 32 links)

(c) 8×2 (hypothetical) network (16 nodes, 29 links)

Fig. 7. Network topologies considered for performance evaluation.

TABLE III

AVERAGE AND STANDARD DEVIATION OF OUTAGE TIMES.

Networks	CSLP		Strict FDPP		Flexible FDPP	
	Average	Std. Dev.	Average	Std. Dev.	Average	Std. Dev
NSFNET	**51.3 ms**	5.5 ms	**53.8 ms**	7.2 ms	**54.4 ms**	7.3 ms
ARPANET	**53.3 ms**	8.6 ms	**61.7 ms**	14.4 ms	**62.5 ms**	14.4 ms
8x2	**38.0 ms**	0.0 ms	**62.4 ms**	23.6 ms	**64.1 ms**	24.0 ms

TABLE IV

AVERAGE PRIMARY AND BACKUP PATH LENGTHS WITH AVERAGE NUMBER OF RECONFIGURATIONS FOR A CONNECTION.

Networks	CSLP		Strict FDPP		Flexible FDPP		
	Primary/ Reconfig.	Backup	Primary/ Reconfig.	Backup	Primary	Backup	Reconfigs.
NSFNET	2.10	4.68	2.09	3.37	2.11	2.21	15.9
ARPANET	2.83	5.25	2.76	3.93	2.76	2.85	22.4
8x2	3.09	4.10	3.28	3.60	3.09	3.17	21.5

all protected traffic versus employing mixed protection strategies depending on the connection requirement is studied. In all the cases, the networks still support unprotected traffic and does not provide any backup paths for them.

Networks that employ only one kind of protection strategy attempt to satisfy all the requests using that strategy. The connections are rejected if the outage time requirement or number of reconfiguration requirement cannot be satisfied. Recall that the network employing multiple protection strategies (referred to as MIXED in the graphs) is assumed to follow a static fixed selection policy for protecting connections depending on their requirements; Type-1 calls are protected using CSLP, Type-2 with Strict FDPP, and Type-3 with Flexible FDPP.

Performance metrics. The performance of the protection strategies are studied through two metrics: (1) blocking probability; and (2) effective network utilization. The blocking probability is computed for individual traffic types, in addition to the overall, as the ratio of the number of requests rejected (of a particular type) to the number of requests received (of that particular type). The effective network utilization is computed as follows. 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, $|\mathcal{L}| \times C$, where C denotes the number of channels in a link. It is to be noted that the effective utilization is computed over the accepted requests only, while the offered load is computed as the effective network utilization over all requests.

Results and discussion. Figures 8, 9, and 10 show the blocking probability (overall and individual types) and effective network utilization for NSFNET, ARPANET, and 8×2 networks, respectively.

For the NSFNET and ARPANET networks, the overall blocking probability [Figs. 8(a) and 9(a)] is high as no scheme can guarantee the requirement of most of the Type-1 calls [Figs. 8(c) and 9(c)]. Such calls account for approximately 20% of the traffic, hence is rejected in both cases. The 8×2 network [Figs. 10(a) and (c)] shows a different trend. While strict FDPP rejects most of the Type-1 requests, flexible FDPP rejects most of the Type-2 and Type-1 requests. This is expected due to the average outage times for strict and flexible FDPP being higher than 50 ms, and the number of reconfigurations under flexible FDPP is more than 0.3 times the total number of failures. It is observed that the overall blocking performance of MIXED

strategy is better than CSLP as the former uses a combination of strategies to optimize utilization.

The blocking performance of Type-2 traffic [Figs. 8(d), 9(d), and 10(d)] shows a similar trend across all networks: (1) flexible FDPP performs the worst due to its inherent characteristic of high reconfigurations; and (2) CSLP performing worse than strict FDPP and MIXED strategies due to excessive resource utilization that is common to any link protection scheme. The blocking performance of Type-3 and Type-4 requests [Figs. 8(e)-(f), 9(e)-(f), and 10(e)-(f)] depend on the performance of a given strategy for Type-1 and Type-2 traffic. In general, if a particular strategy rejected a higher percentage of traffic with more requirements, it accepts more traffic with lesser requirements. These trends are more prominent in the 8×2 network as there is a clear distinction in the outage times of link and path protection strategies.

The effective network utilization under MIXED strategy is the highest for all the networks [Figs. 8(b), 9(b), and 10(b)]. Note that the effective utilization does not account for the priority of the requests, giving equal credits to all types, hence the difference in utilization are less (specifically for NSFNET and ARPANET). A weighted utilization (that reflects the revenue obtained) would show significant improvement.

From the experimental results, we can thus conclude supporting multiple protection strategies is advantageous in the network that it allows for suitable assignment of protection strategies to connections based on their recovery requirements.

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.

VII. CONCLUSION

This paper develops a framework to support multiple protection strategies in optical networks, which is in general applicable to any connection-oriented network. The capacity available on a link for routing primary and backup connections are computed depending on the protection strategy. The paper also develops a model for computing the service outage and failure recovery times for a connection where notifications of failure location are broadcast in the network. The effectiveness of employing multiple protection strategies is established by considering four kinds of traffic on three networks.

In networks where the service outage time for a connection does not vary significantly under path protection and link protection strategies, path protection strategies perform better due to better resource utilization. The advantages of employing multiple protection strategies is significant in networks that have a large difference in the outage times offered by link and path protection strategies. The study conducted in this paper selects a protection strategy in a static manner based on the requirement of the connection. As path protection strategies are known to be more capacity efficient than link protection strategies, connection establishment strategy that first attempts path protection and then attempting link protection strategy, if the outage time requirements cannot be tolerated, may offer better performance compared to static schemes.

REFERENCES

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

[2] S. Ramamurthy, L. Sahasrabuddhe, and B. Mukherjee, "Survivable WDM mesh networks," *Journal of Lightwave Technology*, vol. 21, no. 4, pp. 870–883, April 2003.

[3] 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.

[4] 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.

[5] K. Sathyamurthy and S. Ramasubramanian, "Performance study of link protection schemes under varying quality-of-protection requirements," in *Proceedings of IEEE International Conference on Broadband Networks (BROADNETS)*, October 2004, pp. 300–309.

[6] M. Patel, R. Chandrasekaran, and S. Venkatesan, "A comparative study of restoration schemes and spare capacity assignments in mesh networks," in *in Proceedings of 12th International Conference on Computer Communications and Networks*, October 2003, pp. 399–404.

[7] S. Ramasubramanian and A. Harjani, "DIVERSION: A trade-off between link and path protection strategies," in *Proceedings of the 9th IFIP Working Conference on Optical Network Design and Modelling*, Milan, Italy, February 2005, pp. 321–334.

[8] R. Srinivasan, "MICRON: A framework for connection establishment in optical networks," in *Proceedings of OPTICOMM*, Dallas, TX, USA, October 2003, pp. 139–150.

[9] 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.

[10] C.V. Saradhi and C.S.R. Murthy, "Dynamic establishment of segmented protection paths in single and multi-fiber wdm mesh networks," in *in Proceedings of OPTICOMM*, July-August 2002, pp. 211–222.

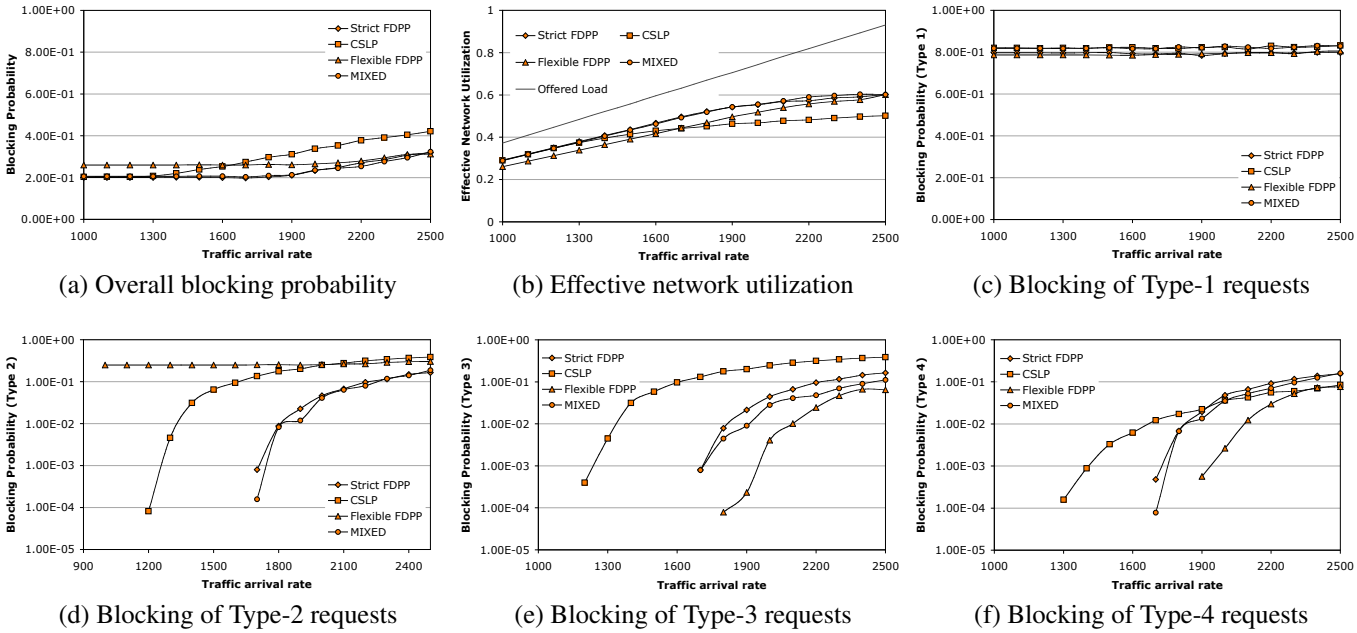

(a) Overall blocking probability (b) Effective network utilization (c) Blocking of Type-1 requests

(d) Blocking of Type-2 requests (e) Blocking of Type-3 requests (f) Blocking of Type-4 requests

Fig. 8. Performance results for NSFNET network.

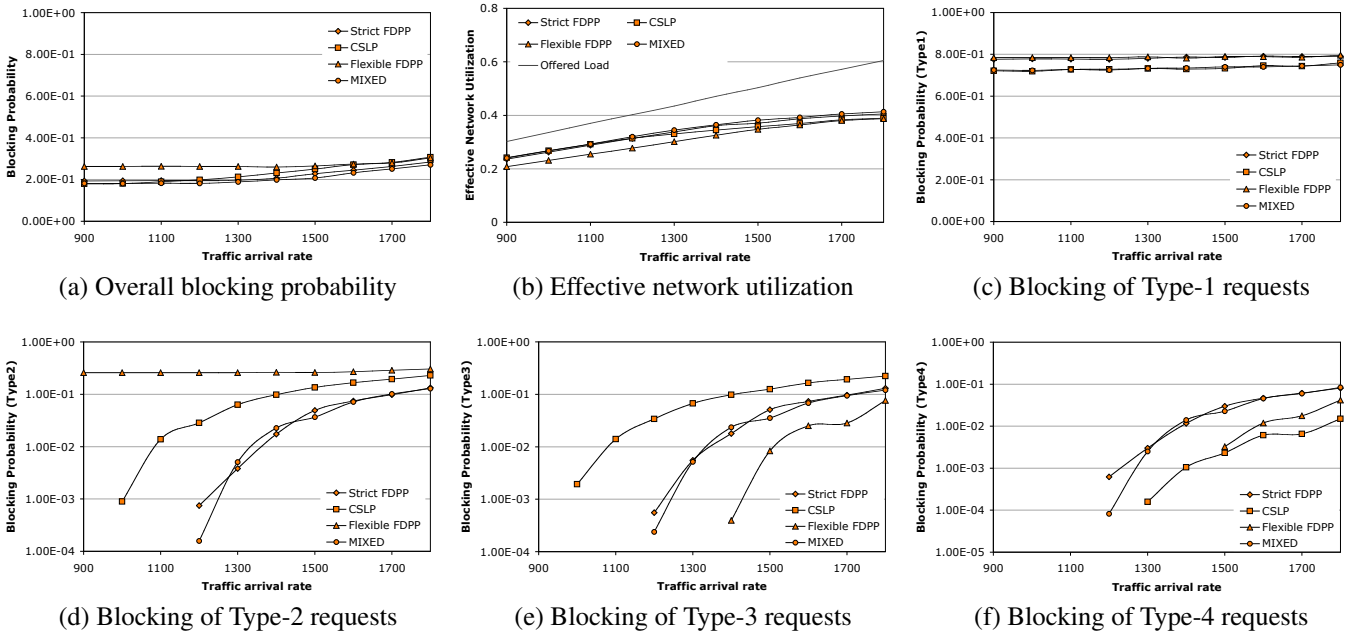

(a) Overall blocking probability (b) Effective network utilization (c) Blocking of Type-1 requests

(d) Blocking of Type-2 requests (e) Blocking of Type-3 requests (f) Blocking of Type-4 requests

Fig. 9. Performance results for ARPANET network.

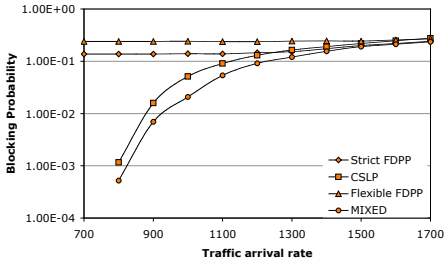

(a) Overall blocking probability

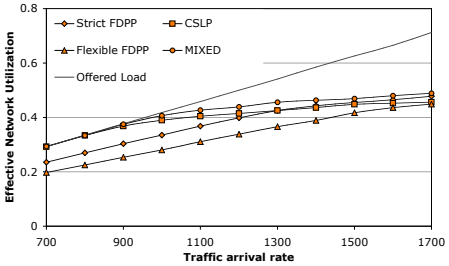

(b) Effective network utilization

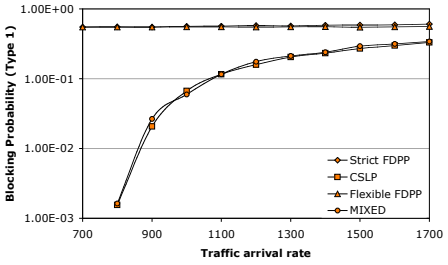

(c)Blocking of Type-1 requests

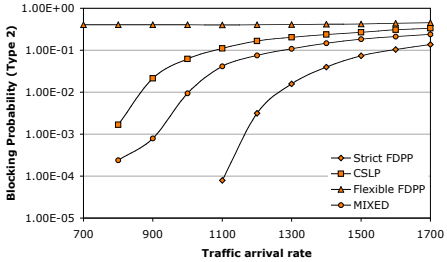

(d) Blocking of Type-2 requests

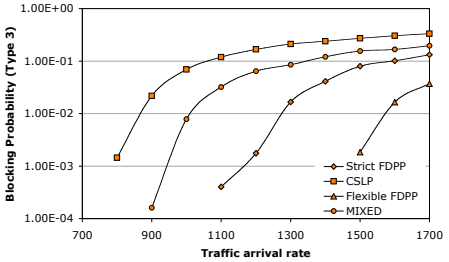

(e) Blocking of Type-3 requests

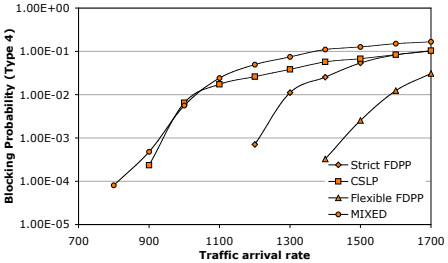

(f) Blocking of Type-4 requests

Fig. 10. Performance results for 8×2 network.