Hybrid Survivability Schemes Achieving High Connection Availability With a Reduced Amount of Backup Resources [Invited]

Jawwad Ahmed, Cicek Cavdar, Paolo Monti, and Lena Wosinska

***Abstract*—Maximizing connection availability in wavelength division multiplexing (WDM) networks is critical because even small disruptions can cause huge data losses. However, there is a trade-off between the level of network survivability and the cost related to the backup resources to be provided. One-hundred percent survivability can be achieved by dedicated path protection (DPP) with multiple prereserved protection paths for each provisioned connection, i.e., DPP (1:*N*). Unfortunately, the blocking probability performance of DPP (1:*N*) is negatively affected by the large number of prereserved backup wavelengths standing by unutilized. On the other hand, path restoration (PR)-based solutions ensure good blocking performance at the expense of lower connection availability. The work in this paper aims at finding hybrid network survivability strategies that combine the benefits of both techniques (i.e., high availability with low blocking rate). More specifically, the paper focuses on a double link failure scenario and proposes two strategies. The first one couples DPP (1:1) with path restoration (referred to as DPP + PR) to minimize the number of dropped connections. The second scheme adds the concept of backup reprovisioning (BR), referred to as DPP + BR + PR, in order to further increase the connection availability achieved by DPP + PR. Integer linear programming models for the implementation of the proposed schemes are formulated. Extensive performance evaluation conducted in a path-computation-element-based WDM network scenario shows that DPP + BR + PR and DPP + PR can significantly lower the blocking probability value compared to DPP (1:2) without compromising too much in terms of connection availability.**

***Index Terms*—Availability; Backup reprovisioning; Failure recovery; Path protection; Path restoration; Survivability; WDM.**

I. Introduction

Achieving high connection availability is extremely important in wavelength division multiplexing (WDM) networks, where each connection (i.e., the wavelength channel referred to as the ligthpath) is carrying a huge amount of data.

Manuscript received May 14, 2013; revised August 15, 2013; accepted August 16, 2013; published September 19, 2013 (Doc. ID 190553).

The authors are with the KTH Royal Institute of Technology, School of Information and Communication Technology, Electrum 229, 164 40 Kista, Sweden (e-mail: jawwad@kth.se).

http://dx.doi.org/10.1364/JOCN.5.00A152

One way to provide 100% network survivability against *single failure* is to rely on path-protection-based approaches [1–4]. Dedicated path protection (DPP) schemes are preferred by network operators primarily because of their fast protection switching times compared to shared path protection (SPP)-based approaches [2]. In a DPP-based scheme, each protected connection has one precomputed, prereserved, and dedicated disjoint backup path that is used when the primary path is disrupted. In SPP, similar to DPP, one backup path is used to protect each primary path. SPP additionally allows sharing of backup wavelength resources among multiple protection paths, under the condition that their respective primaries are not sharing a network resource (i.e., link or node, depending on the failure scenario under exam). In other words, SPP allows for a more efficient use of backup resources while providing 100% survivability against single failures. However, these advantages come at the expense of a higher complexity of the path provisioning and recovery phase (i.e., in terms of path computation, switching time, and control message exchange [2,5]).

In the presence of *multiple failures*, path-protection-based approaches are still able to guarantee 100% survivability, but they need to rely on as many protection paths as the number of failures from which they want to be able to recover [i.e., N backup paths to protect against N simultaneous failures, referred to as the DPP (1:N) scheme]. On the other hand, prereserving a high number of backup resources has a nonnegligible impact on the network ability to provision future connection requests, which in turn results in worse blocking probability performance.

Path restoration (PR)-based techniques [2,6–9] represent an alternative to path protection. A path restoration scheme is *reactive* in nature, as opposed to path protection that is *proactive*. This translates into a more efficient use of network resources, where no prereserved backup wavelengths are standing by unutilized waiting to be used only in the case of a failure. Path restoration schemes are very flexible since they can automatically adapt to any number (single and multiple) and any type (link and node) of failures. However, with path restoration, failure recovery times are typically high because they now comprise a backup path computation phase (unless it is preplanned) and signaling phase for reservation of backup resources. Furthermore, path restoration approaches do not provide any guarantee because wavelength resources might not be available when one or more backup paths need to be reserved, which can be the case in medium to high traffic load conditions. This may translate to nonnegligible downtime values for some of the provisioned connections.

It is apparent that there is a trade-off between the network survivability level and the amount of backup resources that need to be prereserved in the network (leading to possibly high connection blocking probability). It is particularly important to address this trade-off in the presence of multiple failures, a scenario whose likelihood cannot be ignored (e.g., double link failure occurrence probability is nonnegligible [10]). Please note that in this work double link failures refer to scenarios in which the second fault occurs while the first failed element is still under reparation. Furthermore, it is assumed that faults occur randomly in the network on any link.

For this reason hybrid survivable schemes protecting against multiple failures have gained interest. They are considered as an option to provide acceptable survivability levels, but without a drastic impact on the network blocking performance. The idea behind a hybrid scheme is to guarantee survivability at least against one failure using a path-protection-based approach, while the disruptive effects of possible additional failures are addressed using less reliable (i.e., not path-protection-based) survivability strategies. For example, one possibility is to augment DPP (1∶1) with backup reprovisioning (BR) [11–19]. The key idea behind a backup reprovisioning approach is to proactively reprotect those connections that are left in an unprotected state because of the primary/backup path failure. This is different from path restoration, which is a reactive approach, i.e., restoration computes a new path and reserves resources only after a connection is disrupted by failure(s). A successful backup reprovisioning attempt results in a connection having one primary path and a backup path still available in case an additional failure would affect this connection.

A detailed study on the benefits of using backup reprovisioning in WDM networks is presented in [11]. An extended and more recent version of this work is presented in [12], where an algorithm called backup reprovisioning after network state updates (BAND) is proposed to trigger a backup reprovisioning procedure when the network state changes, i.e., after a failure, as well as after a connection arrival and/or departure.

Backup reprovisioning, on the other hand, is still a proactive technique in which backup resources are prereserved regardless of whether they are effectively needed or not. The work in [20] presents an approach in which restoration is used instead of backup reprovisioning to improve blocking performance. More specifically, a link-based restoration scheme is used for quick failure recovery (i.e., the fault notifications are not sent to the source and destination nodes of the failed path). However, link-based restoration is not as effective as path-based restoration because the latter offers more choices for finding a feasible alternate path for an affected connection. Also, local restoration tends to provision suboptimal paths, resulting in increased network resource utilization [2].

In this paper, we propose efficient hybrid survivable schemes offering both high connection availability and low blocking probability. More specifically, we focus on a double link failure scenario and propose a strategy in which both backup reprovisioning and path restoration are used on top of DPP (1∶1). The intuition behind this work is that a combination of backup reprovisioning and path restoration can achieve connection availability values comparable to the ones offered by pure path-protection-based approaches, i.e., DPP (1∶2). With this objective in mind, we propose two dynamic failure recovery schemes: the first is referred to as dedicated path protection + backup reprovisioning + path restoration (DPP + BR + PR), and the second is a simplified version of the first in which only path restoration is used, i.e., dedicated path protection + path restoration (DPP + PR). These proposed schemes provide multiple recovery options for connections affected by double link failures and are specifically aimed at maximizing the average connection availability without the need to multiply the backup resource usage. Consequently, no significant impact on blocking performance has been expected.

Time-efficient integer linear programming (ILP) models are also presented for implementing the proposed schemes. These models are particularly suited for a centralized dynamic provisioning framework (e.g., [21]). Results show that the proposed schemes can achieve high connection availability compared to a pure backup reprovisioning-based approach. At the same time the blocking probability performance of the proposed approaches is far better than that of protection strategies based on multiple dedicated backup paths [i.e., DPP (1∶2)]. To the best of our knowledge, this is one of the first works that investigate the benefits of achieving high connection availability via integrating the backup reprovisioning concept with an end-to-end path restoration scheme in a DPP-based provisioning scenario. Note that, although the proposed failure recovery schemes are presented for a DPP-based dynamic provisioning scenario, the main concept is also applicable to SPP or hybrid path-protection-based strategies [20]. The proposed schemes can also be easily extended to include scenarios with any number and any type of failure.

The paper is organized as follows. Section II introduces and describes the proposed hybrid survivability schemes along with two benchmark solutions used for performance evaluation. ILP formulations for dynamic provisioning of connection requests requiring DPP (1∶2), DPP + BR, DPP + PR, and DPP + BR + PR are presented in Section III. Extensive performance evaluation results are shown in Section IV. Finally, Section V draws some concluding remarks.

II. Survivable Provisioning Framework and Failure Recovery Schemes

This section first provides a general description of the survivable connection provisioning framework considered in this paper. Then it describes the intuition behind the two proposed hybrid survivability schemes, i.e., DPP + PR and DPP + BR + PR. Finally, two additional schemes, i.e., DPP (1:2) and DPP + BR, used for benchmarking purposes are also introduced and described.

A. Framework for Survivable Connection Provisioning

The work in this paper assumes a dynamic traffic provisioning scenario, where no more than two failures, regardless of their type, can affect the network simultaneously, i.e., only a double failure scenario in the network is considered.

Figure 1 illustrates the finite state machine describing the various states in which a connection can be when provisioned in the network. Transitions among states are possible in the presence of specific events: i) a failure happening in the network, ii) a failed element being successfully repaired, iii) a backup reprovisioning attempt being successful or unsuccessful, and iv) a path restoration attempt being successfully/unsuccessfully made for a specific connection. Upon the occurrence of a failure, a connection can find itself in a number of different states, depending on its current status, the failure recovery scheme adopted, and the occurrence of other events (i.e., reparation and/or additional failure) in the network. More details are provided next.

According to the proposed framework, a connection is established along with a reserved, dedicated protection path

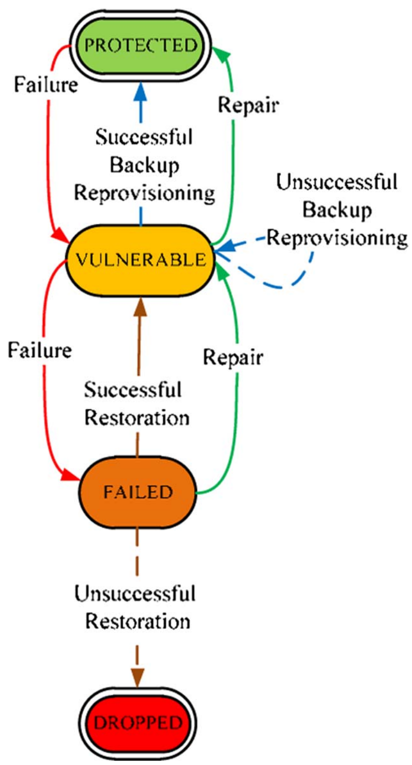

Fig. 1. Finite state machine showing the possible states of a connection and the type of transition a connection can experience while provisioned in the network.

(i.e., DPP). This guarantees the ability to survive at least one failure, i.e., the connection is in a PROTECTED state. If a failure affects a protected connection, then a transition to the VULNERABLE state takes place. This state is used to characterize a connection that is still working normally but is vulnerable to a possible additional failure striking the network, i.e., while another failure is already under reparation. Note that a connection ends up in the VULNERABLE state regardless of which path (i.e., primary or backup) is affected by the failure. One possibility to get back to the PROTECTED state is via a backup reprovisioning attempt. If the attempt is not successful (i.e., lack of wavelength resources), then the connection stays in the VULNERABLE state until the failure is repaired. In the meantime, if another failure affects this vulnerable connection, then a transition to the FAILED state takes place. This state is used to represent a connection that is interrupted, i.e., the services provisioned via this connection are down. At this point, two options are available to get back to the VULNERABLE state. The first one is to do a path restoration attempt. If path restoration is not successful (i.e., lack of wavelength resources) then the connection ends up in the DROPPED state. Once a connection is dropped, it cannot be recovered and its wavelength resources are released. Restoration is chosen in this framework for its reactive nature and for its inherent ability to efficiently use wavelength resources, i.e., backup paths are computed only after a failure occurs, and no backup resources are reserved beforehand. The other option to go back to the VULNERABLE state is to just wait until one of the two failures is repaired. This option might not always be viable, especially for services with strict downtime requirements.

Using the framework described so far it is possible to propose a number of hybrid survivability schemes that combine one or more of the transition options just described. The choice depends on which performance measure (i.e., number of dropped connections versus average connection downtime versus efficient use of network resources) or combination of them is to be addressed. In this paper we propose specifically two survivability schemes: DPP + PR and DPP + BR + PR.

Finally, the last two survivability schemes described in this section, i.e., DPP (1:2) and DPP + BR, will be used in Section IV for benchmarking purposes.

B. Dedicated Path Protection + Path Restoration

The DPP + PR scheme tries to minimize the number of dropped connections while limiting the occupied wavelength resources overall in the network. This is achieved by attempting path restoration each time a connection is in the FAILED state but no backup reprovisioning is attempted in any case. On the other hand, as it was already pointed out, path restoration does not give any guarantee of success, especially at a medium or high network load where wavelength resources might become scarce. In addition, the recovery time of path restoration is relatively long. Both of these aspects are addressed in the scheme described next.

C. Dedicated Path Protection + Backup Reprovisioning + Path Restoration

The DPP + BR + PR scheme uses the same principle and has the same objective as the DPP + PR scheme, i.e., path restoration is attempted for each connection in the FAILED state. This is done to limit the number of dropped connections. However, the DPP + BR + PR scheme has an additional feature. To limit the number of connections that ultimately will require path restoration, each time a connection ends up in the VULNERABLE state backup reprovisioning is attempted. This is done with the intent of reducing the number of connections that, upon the occurrence of a second failure, will transition to the FAILED state and need path restoration. We expect that it would improve the average recovery time since protection-based mechanisms are typically faster than restoration-based techniques. As a result, DPP + BR + PR has the ability to provide lower downtime compared to DPP + PR. On the other hand, backup reprovisioning comes at the cost of additional backup resources being reserved in the network.

D. Dedicated Path Protection (1:2)

The DPP (1:2) [4] is our first benchmarking scheme. It assigns two mutually disjoint protection paths to each primary path. In this way DPP (1:2) guarantees that connections are always in the PROTECTED state as long as a single or double failure scenario is considered, i.e., there is always a protection path available. DPP (1:2) is also rather fast because switches on the intermediate nodes can be preconfigured for the backup paths ahead of time [2]. All these benefits come, on the other hand, at the expense of very high wavelength resource usage.

E. Dedicated Path Protection + Backup Reprovisioning

The DPP + BR [11,12] is our second benchmarking scheme. With DPP + BR, each time a connection becomes vulnerable, backup reprovisioning is attempted once in order to bring the connection back to the PROTECTED state. If backup reprovisioning fails the connection stays in the VULNERABLE state with the risk that, if affected by another failure, the connection becomes DROPPED, i.e., transition back from the FAILED state to VULNERABLE state may not be possible for this survivability scheme.

It is important to note that, upon the occurrence of a failure, not just one, but a considerable number of connections can be potentially disrupted simultaneously. It will then be up to the selected survivability scheme to find an alternate backup route for each affected connection. In such a scenario concurrent optimization schemes have already been proven to be very efficient in finding routing solutions that optimize the resource usage in the network [21]. In this respect, the next section presents a number of ILP formulations that can be used to implement the survivability schemes proposed in this section while exploiting the benefits of concurrent optimization. Although the proposed framework is general enough to handle any type of failure, for the sake of simplicity in the remainder of the paper only (double) link failure scenarios will be considered.

III. ILP Formulations

The purpose of this section is to introduce and explain how to solve the survivable provisioning problem presented in Section II, i.e., DPP + BR + PR and DPP + PR. Remember that these schemes utilize both dedicated path protection along with the backup reprovisioning and the path restoration concepts. For this reason, two ILP models, i.e., ILP_DPP_BR and ILP_DPP_PR, are presented here. A solution for the DPP + BR + PR problem, on the other hand, can then be easily obtained by dynamically solving, for each disrupted connection, the ILP_DPP_BR and the ILP_DPP_PR formulation instances in the correct order (i.e., ILP_DPP_BR first, after a first failure hits a connection, then ILP_DPP_PR is used in the event that the same connection is struck by a second failure). Finally, an ILP formulation for the DPP (1:2) problem is also presented and described.

Before going into the details of each ILP formulation, a description of the inputs and the variables used in the solution of each problem is provided.

Given:

- $G(N,E)$: a physical topology consisting of a set of N nodes and E links.
- W: the maximum number of wavelengths supported on each link.
- W_{xy}: the number of free wavelengths on a link (x,y).
- D: the set of connections disrupted by a failure for which a specific survivable provisioning problem instance needs to be solved.
- D': the set of connections that need to be provisioned into the network. This input is used only in the modeling of the DPP (1:2) case.[1]
- λ_c: a request c from source s to destination d, where $\lambda_c \in D$. Note that a *request* may be to provision a new connection or to perform a failure recovery (i.e., restoration or backup reprovisioning) for an existing connection in the network.
- λ_c^p: equal to 1 if the primary path of connection c from source s to destination d has not failed.
- λ_c^{b1}: equal to 1 if the first backup path of connection c from source s to destination d has not failed.
- λ_c^{b2}: equal to 1 if the second backup path of connection c from source s to destination d has not failed in the case of DPP (1:2).

[1]Note that the cardinality of set D' may vary depending on whether connections need to be provisioned concurrently or one by one [21].

Variables:

- p_{xy}: the total number of wavelengths used by primary paths on link (x, y).
- $b1_{mn}$: the total number of wavelengths used by the first backup paths on a link (m, n).
- $b2_{ij}$: the total number of wavelengths used by the second backup paths on a link (i, j). This variable is used in DPP (1:2) case only.
- p_{xy}^c: equal to 1 if the primary path of request c from source s to destination d passes through link (x, y).
- $b1_{mn}^c$: equal to 1 if the first backup path of request c from source s to destination d passes through a link (m, n).
- $b2_{ij}^c$: equal to 1 if a second backup path of a request c from source s to destination d passes through link (i, j). This variable is used in the DPP (1:2) case only.
- A_c: equal to 1 if a request c is successfully (re)provisioned/restored. Note that the variable A_c is incorporated to make it possible for the presented ILP formulations to correctly handle multiple input requests (typically for concurrent optimization) in scenarios where only a subset of input requests can be successfully satisfied.

A. ILP DPP Backup Reprovisioning

ILP_DPP_BR is used to model and optimally solve the backup reprovisioning problem for the set of connections disrupted by a fault (i.e., D). As mentioned earlier, the ILP_DPP_BR model is used as a part of the DPP + BR and DPP + BR + PR strategy solutions. Note that backup reprovisioning is triggered when either the primary or the backup of a given connection fails (i.e., $\lambda_c^p + \lambda_c^{b1} = 1, \forall \lambda_c \in D$).

Objective 1:

$$\text{Minimize } \alpha \cdot (|D| - \Sigma_{\forall c} A_c) + \beta \cdot \Sigma_{\forall (x,y)} p_{xy} + \gamma 1 \cdot \Sigma_{\forall (m,n)} b1_{mn}.$$

Constraints:

$$\sum_{\forall x} p_{xk}^c - \sum_{\forall x} p_{kx}^c = \begin{cases} A_c, & k = d \\ -A_c, & k = s \\ 0, & k \neq s, d \end{cases}, \quad \forall\, k, c, \exists \lambda_c^p = 0, \quad (1.1)$$

$$\sum_{\forall x} b1_{xk}^c - \sum_{\forall x} b1_{kx}^c = \begin{cases} A_c, & k = d \\ -A_c, & k = s \\ 0, & k \neq s, d \end{cases}, \quad \forall\, k, c, \exists \lambda_c^{b1} = 0, \quad (1.2)$$

$$p_{xy} = \sum_{\forall c, \exists \lambda_c^p = 0} p_{xy}^c, \quad \forall\, (x, y), \quad (1.3)$$

$$b1_{mn} = \sum_{\forall c, \exists \lambda_c^{b1} = 0} b1_{mn}^c, \quad \forall\, (m, n), \quad (1.4)$$

$$p_{xy}^c = 0, \quad \forall\, (x, y), \quad \forall\, c, \exists (\lambda_c^{b1} = 1 \wedge b1_{xy}^c = 1), \quad (1.5)$$

$$b1_{xy}^c = 0, \quad \forall\, (x, y), \quad \forall\, c, \exists (\lambda_c^p = 1 \wedge p_{xy}^c = 1), \quad (1.6)$$

$$p_{mn} + b1_{mn} \le W_{mn}, \quad \forall\, (m, n) \in E. \quad (1.7)$$

The first term in Objective 1 aims at minimizing the number of failed backup reprovisioning attempts, while the second and the third terms aim at minimizing the wavelength resource usage for the primary and backup paths, respectively. Constraints (1.1) and (1.2) are flow conservation constraints for the primary and backup paths, respectively. Constraints (1.3) and (1.4) compute the primary and backup link load, respectively (in terms of the total number of wavelengths used on each link). Constraints (1.5) and (1.6) ensure that a reprovisioned path is link disjoint from the path it is supposed to protect. Constraint (1.7) ensures that the value of the load of any link does not exceed the number of free wavelength channels.

B. ILP DPP Path Restoration

This ILP formulation is used to model and optimally solve the path restoration problem for the set of connections disrupted by a fault (i.e., D). As mentioned earlier, the ILP_DPP_PR model is used as a part of the DPP + PR and DPP + BR + PR solutions.

Objective 2:

$$\text{Minimize } \alpha \cdot (|D| - \Sigma_{\forall c} A_c) + \beta \cdot \Sigma_{\forall (x,y)} p_{xy}.$$

Constraints:

$$\sum_{\forall x} p_{xk}^c - \sum_{\forall x} p_{kx}^c = \begin{cases} A_c, & k = d \\ -A_c, & k = s \\ 0, & k \neq s, d \end{cases}, \quad \forall\, k, c, \quad (2.1)$$

$$p_{xy} = \sum_{\forall c} p_{xy}^c, \quad \forall\, (x, y), \quad (2.2)$$

$$p_{mn} \le W_{mn}, \quad \forall\, (m, n) \in E. \quad (2.3)$$

The first term of Objective 2 aims to minimize the number of failed restoration attempts, while the second term is used to minimize the wavelength resources consumed by the restored paths. Constraint (2.1) is used for flow conservation of each primary path. Constraint (2.2) calculates the load of each link of the primary path. Constraint (2.3) ensures that the value of the load of any link does not exceed the number of free wavelength channels.

C. ILP Dedicated Path Protection (1:2)

This ILP formulation is used to model and optimally solve the DPP (1:2) provisioning problem for the set of connection requests belonging to set D'.

Objective 3:

Minimize $\alpha\cdot(|D^{'}| - \Sigma_{\forall c}A_c) + \beta\cdot\Sigma_{\forall(x,y)}p_{xy} + \gamma 1\cdot\Sigma_{\forall(m,n)}b1_{mn} + \gamma 2\cdot\Sigma_{\forall(i,j)}b2_{ij}$.

Constraints:

Constraints (2.1), (2.2), and

$$\sum_{\forall x} b1^c_{xk} - \sum_{\forall x} b1^c_{kx} = \begin{cases} A_c, & k = d \\ -A_c, & k = s \\ 0, & k \neq s,d \end{cases}, \quad \forall\, k, c, \tag{3.1}$$

$$\sum_{\forall x} b2^c_{xk} - \sum_{\forall x} b2^c_{kx} = \begin{cases} A_c, & k = d \\ -A_c, & k = s \\ 0, & k \neq s,d \end{cases}, \quad \forall\, k, c, \tag{3.2}$$

$$b1_{mn} = \sum_{\forall c} b1^c_{mn}, \quad \forall\, (m,n), \tag{3.3}$$

$$b2_{ij} = \sum_{\forall c} b2^c_{ij}, \quad \forall\, (i,j), \tag{3.4}$$

$$p^c_{xy} + b1^c_{xy} \leq A_c, \quad \forall\, (x,y), \quad \forall\, c, \tag{3.5}$$

$$p^c_{xy} + b2^c_{xy} \leq A_c, \quad \forall\, (x,y), \quad \forall\, c, \tag{3.6}$$

$$b1^c_{xy} + b2^c_{xy} \leq A_c, \quad \forall\, (x,y), \quad \forall\, c, \tag{3.7}$$

$$p_{xy} + b1_{xy} + b2_{xy} \leq W_{xy}, \quad \forall\, (x,y) \in E. \tag{3.8}$$

The first term in Objective 3 focuses on the minimization of the number of failed provisioning attempts, i.e., the number of connections for which a working path or either one of the two mutually link-disjoint protection paths cannot be computed. The second term in the objective function minimizes the resources used by the provisioned primary paths. The third and fourth terms are used to minimize the wavelength resources used to provision the first and second backup path, respectively. Constraints (3.1) and (3.2) are flow conservation constraints used for the first and the second backup paths, respectively. Constraints (3.3) and (3.4) compute the link loads for the first and second backup paths, respectively. Constraints (3.5), (3.6), and (3.7) ensure that the provisioned primary and backup paths are mutually link disjoint. Constraint (3.8) ensures that the value of the load of any link does not exceed the number of free wavelength channels.

In all objective functions, factor α is assigned the highest value in order to maximize the number of provisioned, reprovisioned, and/or restored connections (depending on the specific objective function). Parameters β, $\gamma 1$, and $\gamma 2$ are assigned lower values, and they minimize the wavelength resource usage for the primary and backup paths, respectively.

IV. Simulation Setup and Numerical Results

This section presents some performance assessment results for the proposed hybrid survivable strategies and their respective benchmark techniques described in Sections II and III. A number of simulation experiments are carried out using a Java-based discrete event-driven simulator [22] running on a Red Hat Enterprise Linux (RHEL) workstation. The simulation platform comprises dual Intel Xeon CPUs (4 cores per CPU) clocked at 2.0 GHz and with 12 GB of memory. The Cost239 network topology [23] (Fig. 2) is used for performance evaluation. All fiber links in the network are assumed to be bidirectional with one fiber in each direction, and with 16 wavelengths per fiber. Each lightpath is assumed to carry a bandwidth equivalent to an entire wavelength capacity. The presented results are averaged over 50 replications. The confidence interval for all the plotted results is 5% or less, with a 95% confidence level except for very low values of the blocking probability. The connection holding time is exponentially distributed with average equal to 1 time unit. Moreover, connection request arrivals are assumed to follow a Poisson distribution, while sources and destinations are randomly and uniformly picked throughout the network. For dynamic (DPP 1:1) connection provisioning, the heuristic presented in [24] is used. The ILP models presented in Section III are solved dynamically using the Gurobi Optimizer 4.51 [25].

The time between occurrence of failures and reparation time of a broken link are assumed to be exponentially distributed. The mean time between arrivals of failures in the network is assumed to be 5.0 time units (corresponding to a failure rate equal to 0.2/time unit). Mean time to repair (MTTR) of a failed link is considered to be equal to 0.5 time unit. Link failures are uniformly distributed in the network. Note that the value chosen for the failure rate is on the aggressive side with respect to MTTR. However, this setting makes it possible to complete the simulation-based performance assessment work in an acceptable time frame

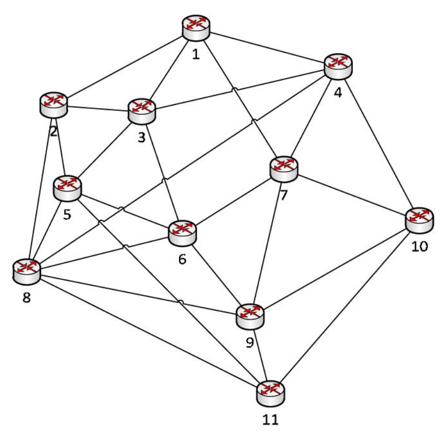

Fig. 2. Cost239 topology.

while still achieving a good statistical accuracy [26,27]. Furthermore, a high failure rate allows us to test the robustness of the proposed schemes considering adverse network conditions (i.e., high double link failure rate), which are not that uncommon in some Asian regions (e.g., India [28]). At the end of the failure recovery phase, the original primary path is restored back (i.e., path reversion) after a link is repaired [3], and a *no-stub release* approach (resources for the original primary are kept reserved for the lifetime of a connection to ease path reversion) is used to implement the path reversion mechanism. The values for α, β, $\gamma 1$, and $\gamma 2$ used in each ILP objective function are assumed to be 10,000, 1.0, 0.5, and 0.25, respectively, for the reasons explained in Section III. The remainder of the section is devoted to analyzing the performance of the proposed hybrid DPP + PR and DPP + BR + PR schemes in terms of a number of performance metrics, i.e., blocking probability, connection unavailability, overall number of dropped connections, restoration success rate, and usage of primary and backup resources.

Note that a preliminary version of this work was presented in [29], where the performance was evaluated in the NSF network topology. However, in this work, we present extended simulation results considering a Cost239 network with additional performance metrics. Furthermore, the considered failure recovery framework and proposed survivability schemes are described in much more detail than in [29].

A. Blocking Probability

Figure 3 presents blocking probability values as a function of the network load. It can be noticed that DPP (1:2) shows a substantially higher blocking probability compared to all the other schemes because two mutually link-disjoint backup paths need to be reserved for each primary path during the connection provisioning phase. This scheme results in a high number of wavelength resources reserved for backup purposes, with a negative effect in terms of successful provisioning of possible future connection arrivals.

As expected, the DPP + PR scheme shows the best blocking performance, particularly at lower loads. This is due to the fact that path restoration does not need any backup resources to be reserved beforehand. Both schemes relying on backup reprovisioning (i.e., DPP + BR and DPP + BR+ PR) show similar levels of blocking probability performance, which is slightly worse than DPP + PR, but still far better than DPP (1:2). This is because backup reprovisioning is performed only after the occurrence of a failure and additional backup resources are reserved only for those connections in the VULNERABLE state for which the backup reprovisioning attempt was successful.

B. Connection Unavailability

Connection availability is defined as the ratio between the total uptime and the sum of the uptime and the downtime of a connection. Connection unavailability is defined as the inverse of connection availability. Figure 4 shows that the DPP + BR scheme experiences the worst connection unavailability performance. This is because once a connection is in the VULNERABLE state, only one backup reprovisioning attempt is made and, if not successful, the connection might be dropped if affected again by another failure. Adding a restoration attempt after an unsuccessful backup reprovisioning greatly improves the unavailability performance. This intuition is confirmed in the figure where the DPP + PR and the DPP + BR + PR schemes show significantly lower unavailability values. Moreover, the DPP + BR + PR scheme is able to achieve connection unavailability that is almost half that of the DPP + PR scheme. This is because, first, backup reprovisioning helps in minimizing the downtime for active connections by maintaining as many connections as possible in a protected state (i.e., to allow quick recovery from a future failure via protection switching) [12]. The second reason is that backup reprovisioning significantly reduces the number of connections that have to resort to path restoration.

As expected, DPP (1:2) presents the lowest connection unavailability values because with this scheme there are always two backup paths available for each newly provisioned connection, essentially protecting the connection against any possible double link failure scenario. Short interruptions may still be incurred in the case of DPP (1:2)

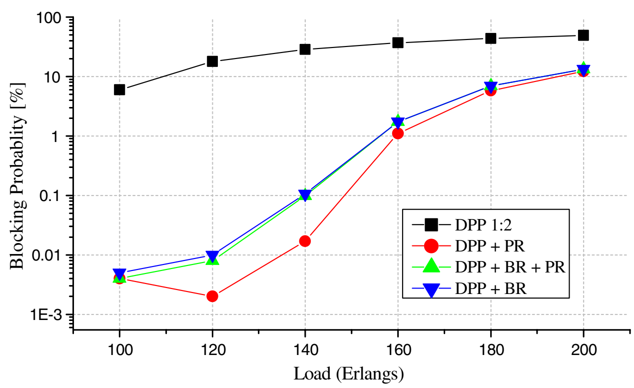

Fig. 3. Blocking probability.

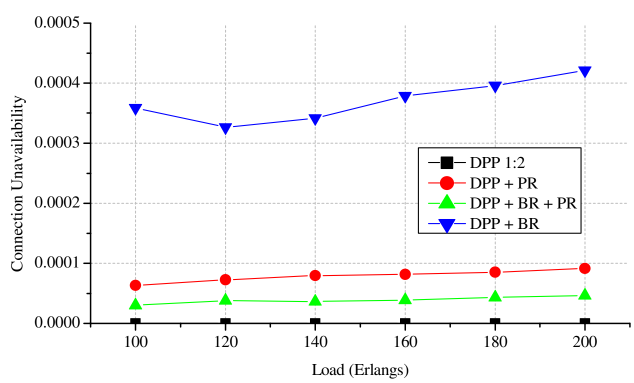

Fig. 4. Connection unavailability.

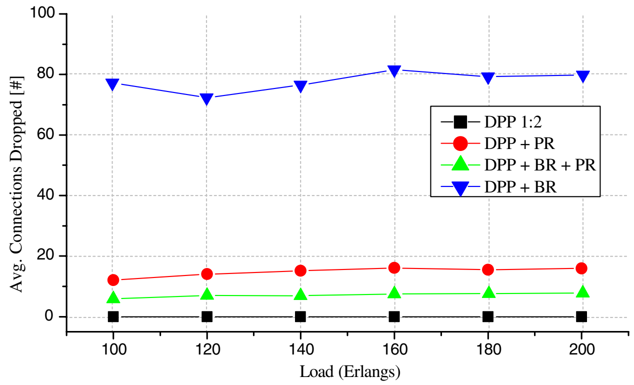

Fig. 5. Average connections dropped.

because of the protection switching mechanism triggered in the event of a link failure on the working path. Nevertheless, the connection unavailability performance of both proposed schemes is reasonably close to DPP (1:2).

C. Number of Dropped Connections

Figure 5 shows the average number of dropped connections as a function of network load. This is an important performance parameter for service providers. It refers to the number of connections that, with the primary and the backup paths both disrupted by failures, could not be successfully restored via a path restoration attempt (i.e., the DPP + PR and/or the DPP + BR + PR scheme). A large number of dropped connections can significantly increase the average connection unavailability. As expected, the number of dropped connections is rather high for the DPP + BR scheme. On the other hand, using the DPP + PR and the DPP + BR + PR schemes can substantially reduce the number of dropped connections. This is related to the fact that with DPP + BR once a connection is in the FAILED state the only way to transition back to the VULNERABLE state is to wait for the failure to be completely repaired. During this time if the connection is affected by another failure it will be dropped (i.e., transition to a DROPPED state). In the case of DPP + PR (and DPP+ BR + PR), on the other hand, a connection in the FAILED state is allowed to attempt restoration that, if successful, will allow the connection to transition back to the VULNERABLE state. Hence, DPP + PR and DPP + BR + PR result in a much lower number of connections being dropped, and consequently in a much lower connection unavailability, as shown in the previous subsection. DPP (1:2) does not experience any dropped connections at all, but this comes at the expense of a much higher blocking probability.

D. Double Link Failure Restorability and Number of Restoration Attempts

The double link failure restorability (DLFR) measures the percentage of connections that experienced a double

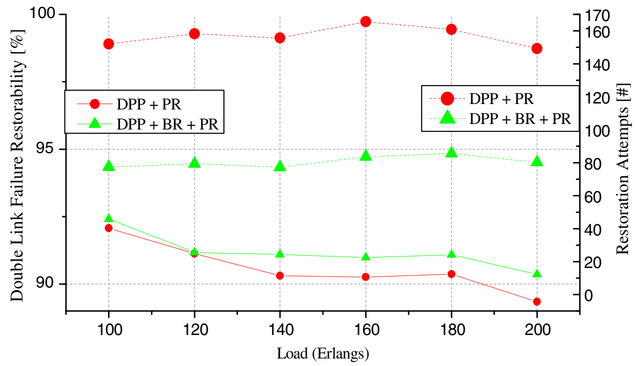

Fig. 6. DLFR and restoration attempts.

link failure but were able to survive due to a successful path restoration attempt. This performance metric applies only to the DPP + PR and the DPP + BR + PR schemes. As shown in Fig. 6, DLFR is around 92% for both the DPP + PR and DPP + BR + PR schemes. This means that path restoration was effective in avoiding drop of 92% of the connections that found themselves in the FAILED state because of a second failure striking them. The value of DLFR tends to decrease with increasing value of the load, as the network resources start to become saturated.

A closely related and equally important performance metric is the total number of connections that have to

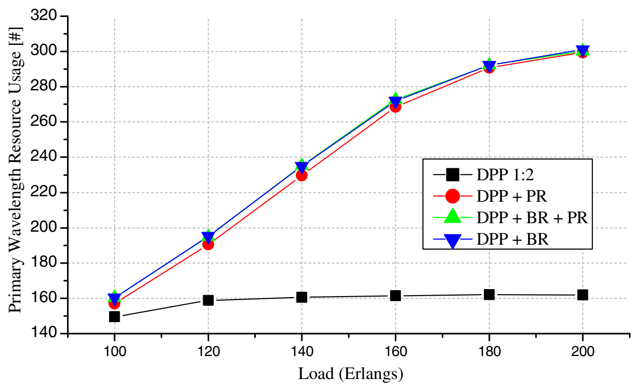

Fig. 7. Primary wavelength resource usage.

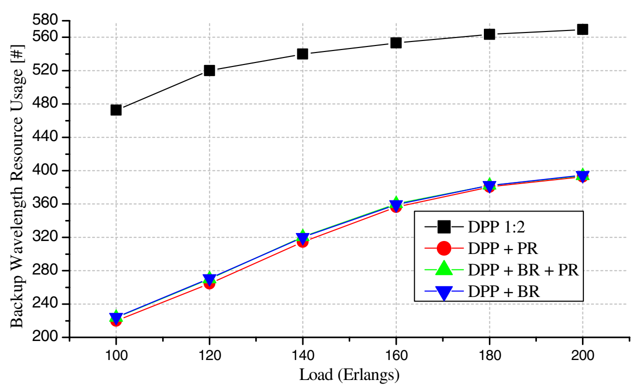

Fig. 8. Backup wavelength resource usage.

TABLE I
TOTAL CONNECTION DOWNTIME AND AVERAGE COMPUTATION TIME FOR THE PATH RESTORATION PROCEDURE

	DPP + PR			DPP + BR + PR			DPP + BR		
Load (Erlangs)	100	140	180	100	140	180	100	140	180
Downtime (time units)	0.415	5.31	17.56	0.345	3.34	8.68	273	284	309
Computation time ILP_DPP_PR (ms)	30.986	31.239	31.277	30.993	31.263	31.24	—	—	—

resort to path restoration. Avoiding a large number of restoration attempts is important because a path restoration procedure takes significant time as it involves the calculation (on the fly) of a new path between the affected source–destination pair as well as the signaling time for the setup of a new path in the network. In contrast, if a backup path is available at the time instance in which a connection fails, then just a fast protection switching operation is involved [2]. The backup reprovisioning can be helpful here, in the sense that it proactively minimizes the number of vulnerable connections in the network. As expected, Fig. 6 shows that with the DPP + BR + PR scheme 50% fewer connections have to resort to path restoration, compared to the DPP + PR scheme.

E. Wavelength Resource Usage

Figures 7 and 8 present, respectively, some performance results in terms of primary and backup wavelength link resource usage, as a function of network load. All the schemes involving backup reprovisioning consume slightly higher resources (i.e., DPP + BR and DPP + BR + PR) as compared to the schemes that utilize only path restoration (i.e., DPP + PR). DPP (1∶2) has the lowest primary wavelength resource usage, which happens due to high blocking probability compared to other schemes.

The figure showing the backup wavelength resource usage (Fig. 8) shows similar behavior for the DPP + BR, the DPP + BR + PR, and the DPP + PR scheme. DPP (1∶2), on the other hand, shows higher backup resource utilization. This happens because DPP (1∶2) provides two backup paths for each provisioned connection as compared to the other schemes that only provision one at a time.

F. Connection Downtime and ILP_DPP_PR Computation Time

Table I shows the connection downtime values experienced when using different failure recovery schemes. The connection downtime is defined as the time in which a connection is not in normal working conditions because of a protection switching event, a restoration attempt being under way, or because of the connection being dropped. As can be expected from the connection unavailability results presented in Fig. 4, the DPP + BR scheme shows the worst downtime values because of the lack of any dynamic restoration procedure to recover from a failure after an unsuccessful backup reprovisioning attempt. On the other hand, the DPP + PR and the DPP + BR + PR schemes show much lower downtime values. Note that the downtime results for DPP (1∶2) are not shown because of their negligible values since the only time a connection is down is during protection switching.

Table I also shows the value of the average computation time required to solve the ILP_DPP_PR formulation as part of the process of finding a solution for the DPP + PR and the DPP + BR + PR scheme, respectively. It was found that the time for solving the ILP_DPP_PR is relatively short (i.e., a few tens of milliseconds) and it is even lower than the value reported in [30]. This means that the proposed ILP model can compute optimal restoration paths within an acceptable time.

V. CONCLUSION

This work addresses the problem of guaranteeing high survivability levels in WDM transport networks in the presence of multiple failures. More specifically, the work focuses on double link failure scenarios and proposes two hybrid survivability schemes, namely, DPP + PR and DPP + BR + PR. They combine the backup reprovisioning concept with an end-to-end path restoration scheme with a focus on maximizing the connection availability without a significant impact on the blocking performance. ILP models formulated for the implementation of the survivability schemes are also presented.

The proposed schemes are evaluated against two benchmark solutions, namely, DPP (1∶2) and DPP + BR. Simulation results show that both proposed schemes achieve substantially better blocking probability performance than DPP (1∶2), while still maintaining acceptable connection availability levels. Furthermore, their performance in terms of connection availability is far better than DPP + BR. The average path computation time when a path restoration procedure is invoked with a set of input connection requests is less than 50 ms, meaning that the proposed ILP model for path restoration procedures ILP_DPP_PR is scalable even under high load (at least in small networks). Finally, the DPP + BR + PR scheme results in low connection downtime values and drops only half as many connections as DPP + PR under high load, which is an important performance parameter for network service providers.

ACKNOWLEDGMENTS

The research leading to these results was financed by the FP7/2007-2013 program under grant agreement no. 318137 (ICT-DISCUS). A preliminary version of this work was presented at the IEEE/OSA Optical Fiber Communication Conference and the National Fiber Optic

Engineers Conference (OFC/NFOEC), Anaheim, CA, USA, March 2013.

References

[1] B. Mukherjee, *Optical WDM Networks*, New York: Springer, Feb. 2006.

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

[3] J. Zhang and B. Mukheriee, "A review of fault management in WDM mesh networks: Basic concepts and research challenges," *IEEE Network*, vol. 18, no. 2, pp. 41–48, Mar.–Apr. 2004.

[4] S. Ramamurthy and B. Mukherjee, "Survivable WDM mesh networks part 1—Protection," in *Proc. IEEE INFOCOM*, New York, Mar. 1999, vol. 2, pp. 744–751.

[5] A. Giorgetti, F. Cugini, F. Paolucci, and P. Castoldi, "OpenFlow and PCE architectures in wavelength switched optical networks," in *Proc. Optical Network Design and Modeling (ONDM)*, Apr. 2012.

[6] S. Ramamurthy and B. Mukherjee, "Survivable WDM mesh networks, part II—Restoration," in *Proc. Int. Conf. on Communications (ICC)*, Vancouver, Canada, 1999, pp. 2023–2030.

[7] H. Choi, S. Subramaniam, and H.-A. Choi, "Loopback recovery from double-link failures in optical mesh networks," *IEEE/ACM Trans. Netw.*, vol. 12, no. 6, pp. 1119–1130, Dec. 2004.

[8] G. Mohan, C. S. R. Murthy, and A. K. Somani, "Efficient algorithms for routing dependable connections in WDM optical networks," *IEEE/ACM Trans. Netw.*, vol. 9, no. 5, pp. 553–566, Oct. 2001.

[9] M. Clouqueur and W. Grover, "Availability analysis of span-restorable mesh networks," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 4, pp. 810–821, May 2002.

[10] S. S. Lumetta and M. Medard, "Classification of two-link failures in all-optical networks," in *Proc. Optical Fiber Communication Conf.*, Vancouver, Canada, 2001.

[11] J. Zhang, K. Zhu, and B. Mukherjee, "Backup reprovisioning to remedy the effect of multiple link failures in WDM mesh networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 57–67, Aug. 2006.

[12] L. Song, J. Zhang, and B. Mukherjee, "A comprehensive study on backup-bandwidth reprovisioning after network-state updates in survivable telecom mesh networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 6, pp. 1366–1377, Dec. 2008.

[13] R. Ramamurthy, A. Akyamac, J.-F. Labourdette, and S. Chaudhuri, "Pre-emptive reprovisioning in mesh optical networks," in *Proc. OFC*, Atlanta, GA, Mar. 2003, vol. 2, pp. 785–787.

[14] S. Kim and S. Lumetta, "Evaluation of protection reconfiguration for multiple failures in WDM mesh networks," in *Proc. OFC*, Atlanta, GA, Mar. 2003, vol. 1, pp. 210–211.

[15] D. Schupke and R. Prinz, "Performance of path protection and rerouting for WDM networks subject to dual failures," in *Proc. OFC*, Atlanta, GA, Mar. 2003, pp. 209–210.

[16] E. Bouillet, J. F. Labourdette, R. Ramamurthy, and S. Chaudhuri, "Lightpath re-optimization in mesh optical networks," *IEEE/ACM Trans. Netw.*, vol. 13, no. 2, pp. 437–447, Apr. 2005.

[17] E. Bouillet, J. Labourdette, and G. Ellinas, "Local optimization of shared backup channels in optical mesh networks," in *Proc. OFC*, Atlanta, GA, Mar. 2003, pp. 774–775.

[18] C. Assi, W. Huo, A. Shami, and N. Ghani, "On the benefits of light-path reprovisioning in optical mesh networks," in *Proc. ICC*, Seoul, South Korea, May 2005.

[19] D. Leung, S. Arakawa, M. Murata, and W. D. Grover, "Re-optimization strategies to maximize traffic-carrying readiness in WDM survivable mesh networks," in *Proc. OFC*, Anaheim, CA, Mar. 2005.

[20] L. Ruan and T. Feng, "A hybrid protection/restoration scheme for two-link failure in WDM mesh networks," in *Proc. IEEE Global Telecommunications Conf.*, Dec. 6–10, 2010, pp. 1–5.

[21] J. Ahmed, C. Cavdar, P. Monti, and L. Wosinska, "A dynamic bulk provisioning framework for concurrent optimization in PCE-based WDM networks," *J. Lightwave Technol.*, vol. 30, no. 14, pp. 2229–2239, 2012.

[22] J. Ahmed and S. Nawaz, "POSE: A new discrete event optical simulator for the Java platform," in *Proc. NCET*, Karachi, Pakistan, Dec. 2004.

[23] M. C. Sinclair and M. J. O'Mahony, "COST 239: Initial network design and analysis," in *Proc. 36th RACE Concertation Meeting*, Brussels, Belgium, 1993, pp. 89–94.

[24] J. Ahmed, P. Monti, and L. Wosinska, "Benefits of connection request bundling in a PCE-based WDM network" (invited paper), in *European Conf. on Networks & Optical Communication (NOC2009)*, Spain, June 2009.

[25] Gurobi Optimizer [Online]. Available: http://www.gurobi.com/.

[26] J. Wang, L. Sahasrabuddhe, and B. Mukherjee, "Path vs. subpath vs. link restoration for fault management in IP-over-WDM networks: Performance comparisons using GMPLS control signaling," *IEEE Commun. Mag.*, vol. 40, no. 11, pp. 80–87, Nov. 2002.

[27] X. Shao, L. Zhou, W. Zheng, and Y. Wang, "Providing differentiated quality-of-protection for surviving double-link failures in WDM mesh networks," in *IEEE Int. Conf. on Communications (ICC)*, Glasgow, Scotland, June 2007, pp. 2180–2185.

[28] M. Ma, "Transport network evolution for advanced services (invited talk)," in *Proc. OFC/NFOEC*, Anaheim, CA, Mar. 2013.

[29] J. Ahmed, C. Cavdar, P. Monti, and L. Wosinska, "Survivability strategies for PCE-based WDM networks offering high reliability performance," in *Optical Fiber Communication Conf. and the Nat. Fiber Optic Engineers Conf. (OFC/NFOEC)*, Anaheim, CA, Mar. 2013.

[30] A. Giorgetti, L. Valcarenghi, F. Cugini, and P. Castoldi, "PCE-based dynamic restoration in wavelength switched optical networks," in *IEEE Int. Conf. on Communications (ICC)*, 2010.