Exploiting Excess Capacity for Survivable Traffic Grooming in Optical Backbone Networks

Ferhat Dikbiyik, Massimo Tornatore, and Biswanath Mukherjee

***Abstract*—Backbone networks usually have some excess capacity to accommodate traffic fluctuations and to avoid early capacity exhaustion. Network operators can exploit excess capacity (EC) in optical WDM backbone networks to support survivable traffic grooming, where connection requests are of sub-wavelength granularity and each provisioned request has to be protected from single-link failures. We investigate novel EC management techniques which can improve network performance, in terms of Service-Level Agreement (SLA) violations and bandwidth blockings, with no requirement of deploying additional capacity. We investigate exploiting and managing EC by the following techniques. *i*) *Pre-provisioning*: When traffic is light, network resources are reserved by a pre-provisioning scheme, i.e., a connection can be provisioned on reserved protected links to increase availability. We show that pre-provisioning also decreases connection setup time, an important metric for delay-sensitive services. *ii*) *Backup reprovisioning*: Since high-availability protection schemes usually consume more resources, connections in our solution can be switched to a protection scheme that provides lower availability (but higher resource efficiency) by reprovisioning backup resources when traffic increases. *iii*) *Hold-lightpath*: We propose a new "hold-lightpath" scheme to exploit EC which prevents the termination of pre-established (but unused) resources to increase availability and decrease connection setup time. We compare our techniques with traditional protection schemes for typical daily fluctuating traffic on typical backbone network topologies, and find that significant improvements can be achieved in terms of decreasing SLA violations, bandwidth blocking, and connection setup time.**

***Index Terms*—Optical network, survivability, traffic grooming, availability, excess capacity, protection, reprovisioning.**

I. Introduction

ANY operational network usually has some excess capacity (EC), viz., some unused capacity which has been deployed to avoid early exhaustion of resources (e.g., bandwidth and grooming ports). EC can be exploited to improve network performance such as connection availability (ratio between time for which a connection is available, i.e., the connection is carrying traffic, and the connection holding time).

In [1], we investigated the problem of exploiting EC where connections are not distinguished in terms of their protection needs and each connection occupies an entire wavelength channel. Ref. [2] is an enhancement of [1] by considering the changes in the states of the connections' downtimes. Note that exploiting EC is a largely unexplored problem, and we provide a brief survey on EC management studies in [1]. Here, in this work, we investigate the EC management problem where connections have heterogeneous bandwidth requirements, i.e., of sub-wavelength granularity, and the network operator has to efficiently aggregate low-speed connections into high-capacity wavelength channels, a process called traffic grooming [3]. Exploiting EC to improve robustness of networks with grooming capabilities is a new problem; and novel and efficient EC management techniques should be designed.

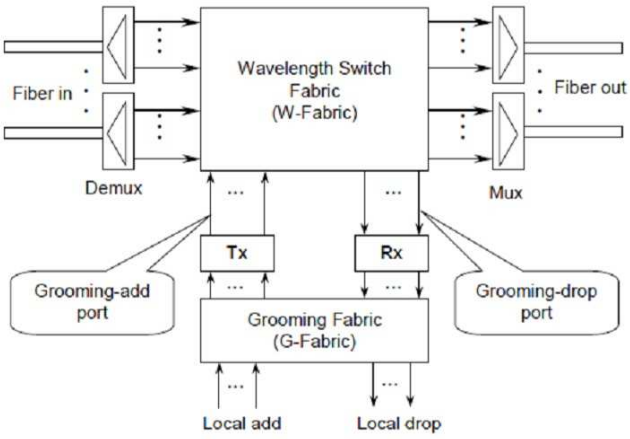

Fig. 1. Simplified grooming node architecture [4].

Figure 1 shows a simplified grooming-node architecture [4] which can switch traffic at wavelength granularity in a W-Fabric and finer granularity in a G-fabric. This node multiplexes, demultiplexes, and switches low-speed connections, each of which is provisioned and groomed on one lightpath (i.e., a full-capacity wavelength channel) or a sequence of lightpaths. Networks with grooming nodes have two types of constraints to provision a connection: both wavelengths and grooming ports are limited in number [4]. For networks with grooming capabilities, EC can be represented by both excess bandwidth and excess grooming ports.

Network robustness is typically provided by protection, i.e., by allocating some spare capacity to protect working traffic by either dedicated protection (dedicating backup resources to a connection or lightpath) or shared protection (backup

This work has been supported in part by the Defense Threat Reduction Agency (DTRA) under Grant No. HDTRA1-08-10-BRCWMD.

F. Dikbiyik is with Department of Computer Engineering, Sakarya University, Sakarya 54187 Turkey (e-mail: fdikbiyik@sakarya.edu.tr).

M. Tornatore and B. Mukherjee are with the University of California, Davis, CA 95616 USA (e-mail: bmukherjee@ucdavis.edu). F. Dikbiyik was also with the U.C. Davis when most of this work was performed.

M. Tornatore is also with Department of Electronics and Informatics, Politecnico di Milano, Milano 20133 Italy (e-mail: tornator@elet.polimi.it).

A short version of this paper was presented at the IEEE Globecom'11 conference in December 2011.

resources are shared among connections or lightpaths). A protection scheme which provides high availability to connections can decrease downtimes experienced by connections and therefore can reduce (or even eliminate) Service-Level Agreement (SLA) violations. The SLA, stipulated between a network operator and its customer, determines the allowed downtime for the connection, and the network operator has to pay a penalty for the downtime that exceeds the allowed downtime. Even though high-availability protection schemes can decrease connections' downtimes, they usually require high capacity (both in terms of bandwidth and grooming ports), so they may increase the rejection of connections. Refs. [4] and [5] showed availability and bandwidth blocking ratio (BBR)[1] performance of two fundamental protection paradigms for survivable traffic grooming: protection-at-lightpath (PAL) and protection-at-connection (PAC) levels, for shared and dedicated protection. i) PAL provides end-to-end protection by provisioning a backup path for each individual lightpath separately (note that connections may traverse a sequence of protected lightpaths, called p-lightpaths); ii) PAC provides end-to-end protection by providing a backup path at connection level. Both dedicated PAC (D-PAC) and dedicated PAL (D-PAL) provide higher availability (so they decrease SLA violations) but consume more resources (bandwidth and grooming ports), causing higher BBR and lower grooming efficiency than their shared counterparts (S-PAC and S-PAL). Thus, for backbone networks with grooming capabilities, there is a trade-off between SLA violations and BBR. This trade-off is explained in details in Section II. Note that these protection schemes differ from protection schemes mentioned in previous works [1], [2] in the sense that these are designed for survivable grooming by considering a series of lightpaths to support connections.

In a dynamic scenario, where connections arrive, hold for a while, and terminate, EC can be exploited to address this trade-off and to reduce the network operator's cost, both in terms of SLA violation penalties and loss due to bandwidth rejections. When traffic load is low, a large amount of EC may exist and is unutilized. In this case, dedicated protection can be suitable to benefit from its high availability. When EC in the network decreases (due to traffic increase), reprovisioning backup resources using shared protection frees some bandwidth and grooming ports, and helps to decrease BBR and increase grooming efficiency. After the reprovisioning, new connections can be provisioned by dedicated protection. Note that, while reprovisioning backup resources, the current state of the connections' accumulated downtimes should be considered. The downtime of some connections, if they switch to shared protection, might be close to exceeding their allowed downtime in their remaining holding time. These connections should be reprovisioned by dedicated protection, while other connections can be protected by shared protection.

EC in the network can be managed by this admitting-by-dedicated-and-reprovisioning-to-shared scheme. However, EC can be further exploited to improve network robustness by also using a link-based protection scheme (protecting each link by a dedicated backup lightpath). Link protection has been widely discussed in several works [6]–[8] without considering grooming, and these works showed that dedicated link protection provides high availability but consumes too many resources. In this study, we explore the opportunity to use link protection for survivable grooming, a method that is largely unexplored. When there is sufficient EC, we consider reserving resources before connections arrive such that each link has some protected lightpaths so that, when connections arrive, they can be provisioned on these protected links and there is no need to setup backup paths on the fly. We call this scheme *pre-provisioning* (also discussed in [1], but enhanced here for survivable grooming), and each lightpath on a protected link with a backup lightpath is referred as *1-link-p-lightpath*. Using pre-provisioning, when a new connection arrives, it can be easily provisioned and groomed on these 1-link-p-lightpaths, which *(i)* increases availability because link protection provides higher availability than path protection, and *(ii)* provides faster connection setup time, an important metric for delay-intolerant services, as there is no need for configuring optical crossconnects (OXCs).

As a broader goal, providing high-availability protection in advance may also help to better prepare the network against large-scale failures due to disasters, e.g., weapons of mass destruction (WMD) attacks, earthquakes, etc. To avoid early exhaustion of bandwidth and grooming ports due to excessive pre-provisioning, we study a statistical pre-provisioning approach, where we consider the traffic intensity between any node pairs.

Therefore, in this study, we investigate how to exploit the EC in a dynamic scenario by *i*) *pre-provisioning* resources by link protection when EC in the network is large and *ii*) *reprovisioning* backup resources of some connections by dedicated protection and others by shared protection based on the downtimes experienced by connections, when EC in the network reduces (due to more traffic). We show that pre-provisioning and reprovisioning allow to significantly decrease BBR and increase grooming efficiency by improving availability. We also propose a complementary EC management scheme, called *hold-lightpath*, which prevents the termination of pre-established (but unused) lightpaths, and therefore increases the availability by creating a segmented protection [9]. We also show that both pre-provisioning and hold-lightpath decrease connection setup time. To the best of our knowledge, this is the first study to consider all these parameters (availability, grooming efficiency, BBR, and connection setup time) to exploit EC for survivable traffic grooming with pre-provisioning, reprovisioning, and hold-lightpath.

This study is organized as follows. We compare survivable grooming techniques in Section II to show our motivation. We provide the problem statement in Section III. Section IV explains our EC management scheme and its key steps (pre-provisioning, provisioning, and backup reprovisioning). We propose hold-lightpath scheme in Section V and show illustrative numerical examples in Section VI. Section VII concludes the study.

[1]BBR is the ratio of total amount of rejected bandwidth to requested bandwidth. It is an useful metric when studying connections of different bandwidth needs.

II. Comparison of Survivable Traffic-Grooming Techniques

References [4] and [5] compare PAC and PAL schemes through numerical examples in various parameters. However, they do not explicitly make a comparison between dedicated and shared protection schemes. Here, we provide an analysis of availability provided by the survivable traffic grooming techniques and show the amount of downtime they cause through an example to explain the SLA violations. We also show the amount of resources they require to understand their BBR potentials.

A. Availability Analysis

1) PAL: As mentioned above, PAL is forming connections over protected lightpaths (each called a p-lightpath) and if a link on a lightpath fails, the connection is switched to that lightpath's backup lightpath. Thus, availability of a connection (A_t^c) protected by PAL can be specified as $A_t^c = \prod_{\eta \in H_t} A_\eta^{pl}$, where H_t is the set of p-lightpaths to form the connection t and A_η^{pl} is the availability of the p-lightpath η. For D-PAL, A_η^{pl} is similar to the availability of a connection protected by Dedicated Path Protection [10], so it can be expressed as:

$$A_\eta^{pl} = A_\eta^{w_{pl}} + (1 - A_\eta^{w_{pl}}) A_\eta^{b_{pl}} \tag{1}$$

where $A_\eta^{w_{pl}}$ and $A_\eta^{b_{pl}}$ are primary and backup paths of η, respectively. Note that the availability of a path is the product of the availabilities of the links on that path[2]. For S-PAL, availability of a p-lightpath can be calculated through its unavailability U_η^{pl} similar to availability calculation of a connection protected by Shared Path Protection [11]. So, we can express U_η^{pl} by:

$$U_\eta^{pl} = U_\eta^{w_{pl}} (U_\eta^{b_{pl}} + 0.5(\sum_{\nu \in V_\eta} U_\nu^{w_{pl}})) \tag{2}$$

where V_η is the set of p-lightpaths that share the backup resources with p-lightpath η. So, we can formulate the availability of a connection protected by PAL as:

$$A_t^{D-PAL} = \prod_{\eta \in H_t} A_\eta^{w_{pl}} + (1 - A_\eta^{w_{pl}}) A_\eta^{b_{pl}} \tag{3a}$$

$$A_t^{S-PAL} = 1 - \sum_{\eta \in H_t} U_\eta^{pl} \tag{3b}$$

2) PAC: PAC forms connections over a set of unprotected lightpaths to establish its primary path and forms its backup path on a different set of lightpaths. Let L_t^p and L_t^b be the sets of lightpaths on primary and backup paths, respectively. Then, we can similarly define the availabilities as follows:

$$A_t^{D-PAC} = \prod_{l \in L_t^p} A_l + (1 - \prod_{l \in L_t^p} A_l) \prod_{l \in L_t^b} A_l \tag{4a}$$

$$A_t^{S-PAC} = 1 - \sum_{l \in L_t^p} U_l (\sum_{l \in L_t^b} U_l + 0.5 \sum_{l \in L_t^s} U_l) \tag{4b}$$

[2]Here, we mean the statistical availability of the links for convenience, which is the long-term availability. However, a link's actual availability is usually less than its statistical availability during a typical holding time of a connection.

where L_t^s is the set of lightpaths that are on the primary paths of the connections that share backup resources with connection t, while A_l and U_l are the availability and unavailability of lightpath l, respectivley.

Both Eqs. (3b) and (4b) become equivalent to their dedicated counterparts when there is no sharing of backup resources (i.e., $|V_\eta| = 0,\ \forall \eta \in H_t$ in Eq. (3b) and $|L_t^s| = 0$ in Eq. (4b)). When the sharing increases in time (naturally encouraged by the shared protection schemes to avoid capacity exhaustion), the availability of shared protection schemes decreases and downtime experienced by a connection increases (and so does the risk of SLA violation).

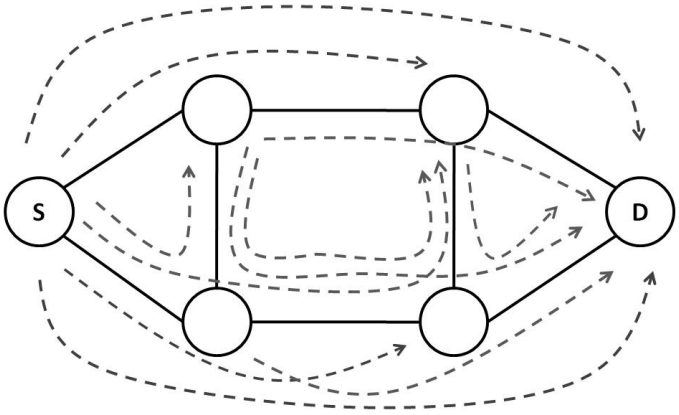

Fig. 2. A 6-node partial network with candidate lightpaths to form a connection from node S to D.

Figure 2 shows a 6-node partial network (the links of the full network are not shown), where there is a connection request from node S to D. Besides the lightpaths that can be formed on each link, i.e., between the end nodes of the links (not shown in the figure), there are some candidate lightpaths (shown with dashed lines) to form the connection protected by either PAC or PAL. For each scheme, there might be alternate paths based on the free resources (bandwidth and grooming ports). Figure 3 shows downtime of the connection protected by PAL or PAC with alternate choice of lightpaths to form the connection and backup resources. These values are calculated by using Eqs. (3) and (4) with the following assumptions: (*i*) each link has availability of 0.99 and (*ii*) the primary path of the connections that share backup resources with the requested connection has 2 hops on average. The horizontal dashed lines show some examples of allowed down times (ADT) normalized by the connection holding time w.r.t. specific target availabilities. The downtimes above these lines indicate SLA violations for which network operator has to pay some penalty. Figs. 3(a) and 3(b) show that dedicated-protection schemes have less violations compared to their shared counterparts and violations increase as the sharing increases.

B. BBR Analysis

Both D-PAC and S-PAC approaches require a grooming-add port and a grooming-drop port for each lightpath established, regardless of whether it is on the primary or backup path. However, the bandwidth reserved for backup on a lightpath

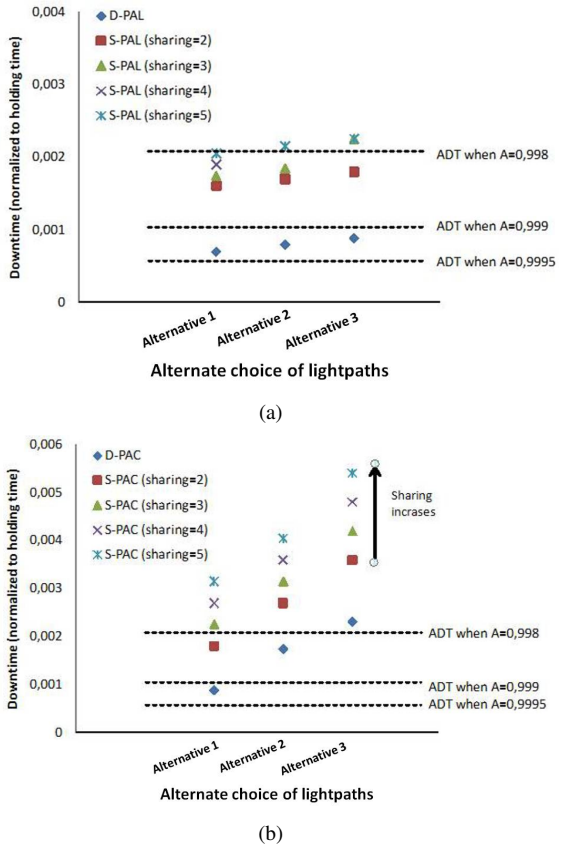

Fig. 3. Downtime (normalized by the connection holding time) of the connection protected by (a) PAL and (b) PAC with alternate paths.

can be shared among multiple connections for S-PAC. For instance, a lightpath with free bandwidth of $5B$ can be used as backup capacity by MAS$\times 5 \times B$ number of connections each of which requires bandwidth B, where MAS is maximum allowed sharing [12] (e.g., 5), as long as the other sharing requirements are met (e.g., link-disjointness of the primary paths). However, for D-PAC, the same lightpath can be used as a backup for 5 connections, each with required bandwidth of B. To support more bandwidth, it has to find another lightpath or establish a new one by consuming new grooming ports.

For S-PAL, backup lightpaths do not even require additional grooming ports, because, when the primary path of a p-lightpath fails, it can switch to backup lightpath with all the connections carried on the primary path and can use the grooming ports of the primary path. Similar to S-PAC, bandwidth on backup lightpaths can be shared among multiple connections. On the other hand, D-PAL requires two grooming-add and -drop ports for a p-lightpath and backup resources cannot be shared.

Despite their advantage of reducing SLA violations, dedicated schemes' excessive bandwidth and grooming-port consumptions show a significant disadvantage in terms of high BBR compared to shared schemes.

III. Problem Statement

As showed above, shared and dedicated survivable grooming techniques have conflicting advantages, i.e., the shared schemes can provide lower BBR with the risk of high SLA violation and the dedicated schemes can significantly reduce or even eliminate SLA violations, but can cause high bandwidth blocking. To decrease the SLA violations, one can use availability-aware approaches (e.g., [13]), which provision lightpath on most reliable paths, for shared schemes. For instance, if the primary resources of a connection are on very reliable links, then the connection can be protected by a shared scheme, while if the primary resources are not on very reliable links, then the connection can be protected by a dedicated scheme. The availability-aware approaches use the statistical availability of links, which shows the links' uptime with respect to a long duration of time compared to typical holding time of a connection. Thus, the decision on whether a connection should be protected by a dedicated or shared scheme must be handled by a dynamic method.

For instance, if a connection protected by dedicated protection has not experienced any downtime for a long time, then it can be switched to shared protection to lower the BBR; and if a connection protected by shared scheme has experienced some downtime, then it can be switched to dedicated protection to lower the risk of SLA violation. Here, we consider to exploit EC in the network for this adaptive protection scheme by admitting connections with dedicated protection and switch their protection schemes based on the downtimes they experienced. Even though, in previous works [1], [2], we propose EC management approaches, where each connection occupies an entire wavelength channel, applying them to grooming-capable networks is not straightforward because survivable grooming techniques are quite different than the protection techniques considered in [1], [2]. Also, EC, in this work, is defined by considering both excess bandwidth and grooming ports. Here, we also propose a novel pre-provisioning approach to exploit EC with link protection (which is designed for grooming-capable networks for the first time) and a new complementary scheme, hold-lightpath, to decrease SLA violations and to have shorter setup time.

Note that the optimal solution of the problem is a complex scheduling problem that requires knowledge of the connection arrival/departure times and link failure/repair times, which are not available in practice. Surely, the optimal solution for the single pieces of EC management scheme, namely pre-provisioning and reprovisioning, can be formulated. However, in our previous works [1], [2], we show that such optimal solutions for these features, even for the connections that require full wavelength channel, are not scalable. Since EC management may require several reprovisionings in a dynamic scenario. and the finer granularity of bandwidth (i.e., traffic grooming) is considered in this work, optimal solution would be even more intractable for large network instances. Therefore, we provide heuristics to reach our goal, namely lower SLA violations and BBR.

The general problem we investigate here is to improve network robustness, in terms of decreasing the overall penalties that may be paid by network operator due to SLA violations and loss due to bandwidth blocking, by exploiting EC. In the next section, we explain how to manage EC to solve this problem.

IV. Excess Capacity Management

The flow chart in Fig. 4 shows the proposed EC management scheme for grooming (ECM). After pre-provisioning of resources, ECM waits for a connection arrival. A connection is admitted based on its availability requirements and the amount of reserved resources. During provisioning (shown by dashed lines), if reserved or free resources do not suffice to provision the connection, backup reprovisioning and pre-provisioning may be triggered to rearrenge and free network resources. The key steps of ECM, namely pre-provisioning, provisioning, and backup reprovisioning, and the key parameters, namely A_t, A_{th}, and ξ_{sd}, are explained below.

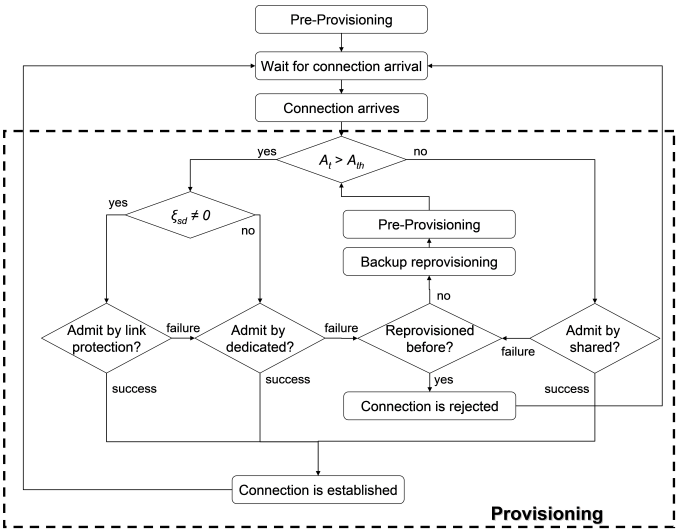

Fig. 4. Excess capacity management (ECM).

A. Statistical Pre-Provisioning

Figure 5 compares traditional protection schemes, PAC and PAL (Fig. 5(a)), with the link protection (Fig. 5(b)) used here for pre-provisioning on a 7-node network, where the labels on the nodes show the numbers of add/drop grooming ports used. The solid and dashed lines show links with link availabilites of 0.99 and 0.98, respectively. Initially, there are no connections in the network, and then a connection request (t_1) from node C to node F arrives. In Fig. 5(a), a working lightpath on the path C-B-D-F (w_1) and a backup lightpath on the path C-E-G-F (b_1) are lit, and the connection is provisioned on these lightpaths. Note that, for both PAC and PAL, b_1 is used as backup. Then, another connection request (t_2) from node A to node F arrives. A working lightpath on the path A-B-D-F is lit. Assuming b_1 has enough spare capacity, PAC can provision a new backup lightpath on path A-C ($b_{2(PAC)}$) and groom the connection on b_1 and $b_{2(PAC)}$ for backup. If PAL is used for protection, it requires a separate lightpath for each working lightpath, so it provisions a lightpath on path A-C-E-G-F ($b_{2(PAL)}$) for backup. Figure 5(b) shows an example of link protection for survivable grooming of the same connection requests. When t_1 arrives, three one-link lightpaths ($w_1, w_2,$ and w_3) are lit up on links C-B, B-D, and D-F, respectively, and each lightpath is protected by a dedicated backup path ($b_1, b_2,$ and b_3). Then, t_1 is provisioned on these lightpaths. When t_2 arrives, it can be groomed on a new 1-link-p-lightpath on link A-B (w_4 protected by b_4) and pre-established lightpaths w_2 and w_3. Intuitively, link protection requires more resources, both in terms of capacity and grooming ports (which is not an issue if there is enough EC), but it provides higher availability. For instance, based on the availability calculation provided in various studies (e.g., [1], [10]) for path and link protections, link protection in this example increases connection availability of t_1 from 0.9983 to 0.9984 and connection availability of t_2 from 0.9977 to 0.9986.

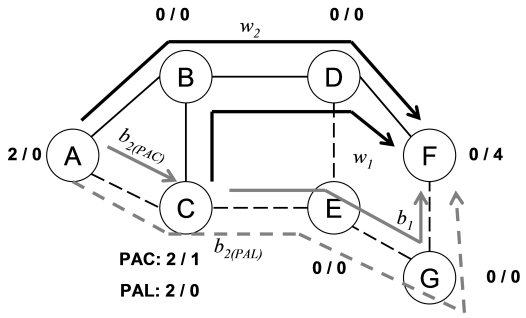

(a)

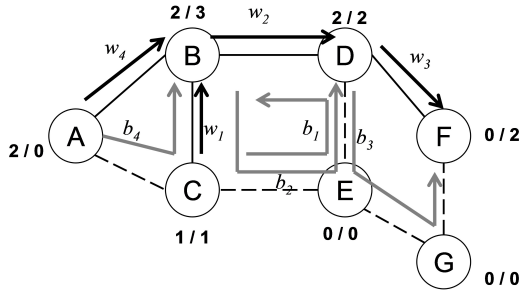

(b)

Fig. 5. Protection with (a) PAC or PAL, and (b) link protection.

Link protection can be provided before a connection arrives by pre-provisioning, i.e., reserving 1-link-p-lightpaths. Besides increasing availability, pre-deployment of lightpaths can also decrease resource usage and setup time of connections [1], [14]. However, pre-provisioning of 1-link-p-lightpaths for each link might not be possible. For instance, in Fig. 5(b), if node D has only two grooming add (or drop) ports, then 1-link-p-lightpath on link D-E can not be provisioned. Thus, we consider statistical pre-provisioning which provisions 1-link-p-lightpaths on the links of shortest paths between source-destination pairs with high traffic intensities. We develop an algorithm for pre-provisioning shown in Alg. 1 for a given network topology graph $G(V, E)$, where V is the set of nodes and E is the set of links, and given the set of source-destination pairs denoted by $\Psi = \{\psi =< s, d, \phi_{sd} >\}$ where *s*, *d*, and ϕ_{sd} are source, destination, and traffic intensity between s-d pair, respectively. Our approach starts from the s-d pair with highest traffic intensity and finds k-shortest paths from source to destination based on the available capacity and selects the path which has the largest bottleneck (the link which has least

free capacity on a path is the bottleneck). Then, it computes the maximum number of lightpaths (ξ_{sd}) that can be pre-provisioned on this path for the s-d pair in Step 6, where $\bar{\phi}_{sd}$ is ϕ_{sd} normalized to $\max_{sd\in\Psi}\phi_{sd}$ so that there will be more 1-link-p-lightpaths on the paths of s-d pairs with high traffic intensities. The pre-provisioning algorithm then takes a free wavelength on each link on the path and tries to find a backup path for it. In our ECM scheme, pre-provisioning is performed when some resources are freed by termination of a wavelength or reprovisioning of backup resources.

Algorithm 1 Pre-Provisioning

1: Sort node pairs in Ψ w.r.t. their traffic intensities (ϕ_{sd}).
2: Update link costs as follow:

$$C(e) = \begin{cases} \infty & F(e) = 0 \\ 1 + \varepsilon \times (W(e) - F(e)) & (o.w.) \end{cases} \quad (5)$$

where $W(e)$ and $F(e)$ are total and free wavelengths on link *e*, respectively; ε is a small number (e.g., 10^{-5}).
3: **for all** $\psi \in \Psi$ **do**
4: Find *k*-shortest paths for ψ. Let K_{sd} be the set of these *k* number of shortest paths.
5: Update cost of each path $p \in K_{sd}$ with $C(p) = (\min_{e\in p} F(e))^{-1}$ and take minimum-cost path (p^{sd}_{min}).
6: Compute possible number of lightpaths, $\xi_{sd} = \lfloor \bar{\phi}_{sd}/C(p^{sd}_{min}) \rfloor$ where $\bar{\phi}_{sd}$ is ϕ_{sd} normalized to $\max_{sd\in\Psi}\phi_{sd}$.
7: **if** $\xi_{sd} \neq 0$ **then**
8: **for all** $e \in p^{sd}_{min}$ **do**
9: Take ξ_{sd} number of free wavelengths on link *e* (let Ω_e be set of these wavelengths on link *e*).
10: **for all** $w \in \Omega_e$ **do**
11: Find a backup path for wavelength w (b_w).
12: **if** b_w exists (i.e., a free wavelength exists on each link of b_w) and there are more than one free grooming-drop port at the ingress node and more than one free grooming-add port at the egress node of link *e*, **then**
13: Provision a lightpath l^w_e on wavelength w on link *e*, and ligthpath $l^{b_w}_e$ on backup path b_w.
14: Assign $l^{b_w}_e$ to l^w_e (i.e., everything is sent through w will also be sent through b_w).
15: Increment $R(e)$ (number of reserved wavelengths on link *e*)
16: **end if**
17: **end for**
18: **end for**
19: **end if**
20: **end for**

B. Provisioning

We assume that connections have specific availability requirements stated in their SLA. The connections that have low availability requirements can be provisioned by shared protection. An availability threshold (A_{th}) can be used to determine if a connection requires high or low availability. However, if connection request $t(s_t, d_t, B_t, A_t)$ arrives, where s_t, d_t, B_t, and A_t are the source, destination, bandwidth, and availability requirement of connection t, and if $A_t \geq A_{th}$ and there are enough reserved 1-link lightpaths on s-d pair's shortest path ($\xi_{sd} \neq 0$), this connection can be protected by link protection. If there are no resources reserved for this s-d pair due to low traffic intensity or reserved resources are not sufficient to provide link protection, then this connection can be provisioned by dedicated protection. If network resources do not suffice to provide dedicated protection (for $A_t \geq A_{th}$) or shared protection (for $A_t \leq A_{th}$), then ECM triggers backup reprovisioning followed by a pre-provisioning. If after the reprovisioning, the connection request cannot be provisioned because of lack of resources, then it is rejected. A flowchart of the proposed provisioning scheme is shown in Fig. 4 inside the dashed rectangle.

C. Backup Reprovisioning

Due to high resource consumption of link- and dedicated-protection schemes used in our pre-provisioning and provisioning phases, network resources might get exhausted as more connections arrive. Hence, reprovisioning of backup resources would be required to free some network resources. Reprovisioning in WDM networks has been studied [12], [15], [16] to reconfigure lightpaths or connections requiring full wavelength channel. Ref. [17] proposes a method for backup reprovisioning for survivable grooming where backup resources are reprovisioned after a failure for vulnerable connections (PAC) or lightpaths (PAL). Backup reprovisioning proposed here can be applied for both PAL and PAC, but we focus on PAC[3] while preliminary works for PAL can be found in [18].

Here, we consider a global backup reprovisioning scheme where some connections' backup resources are reprovisioned by dedicated protection, while others' are reprovisioned by shared protection, depending on the connections' downtime tolerances for their remaining holding time. Ref. [19] introduced the *urgency level* (UL) concept which captures these parameters. In [19], UL is defined as a function of the allowed number of failures (ANF), remaining holding time (RHT), downtime exceeding allowed donwtime (SDT), and SLA violation penalty. ANF of connection t at time of reprovisioning describes the risk of SLA violation and is defined as:

$$ANF_t = \left\lfloor \frac{(1 - A_t) \times h_t - DT_t}{MTTR} \right\rfloor \quad (6)$$

where A_t is the target availability specified in SLA (e.g., 0.999, 0.9999, etc.), h_t is holding time of connection t, DT_t is down time of connection t at the time of the reprovisioning, and $MTTR$ is mean time to repair of a failure. When $ANF = 0$, a connection cannot afford any more failures. In this case, a connection needs more protection (i.e., UL is high). We modify UL definition to capture the downtime tolerances of

[3]Ref. [4] proposes two PAC schemes: mixed-PAC (MPAC) and separated-PAC (SPAC) where the first allows a lightpath to carry primary and backup traffic and the latter dictates that a lightpath can only carry primary or backup traffic. Here, we consider MPAC, but extending it on SPAC is intuitive.

connections without incurring penalty parameters introduced in [19]. The UL of connection t is given by:

$$UL_t = \begin{cases} RHT_t/ANF_t, & \text{if } ANF_t > 0 \quad (7a) \\ RHT_t, & \text{if } ANF_t = 0 \quad (7b) \\ SDT_t \times RHT_t & \text{if otherwise} \quad (7c) \end{cases}$$

where (7a) means that connection t can afford failure(s) and its UL is proportional to its RHT and inversely proportional to how many more failures it can afford, (7b) shows that connection d cannot afford any failures and its UL depends on its RHT, and (7c) shows that connection t's SLA is already violated and urgency depends on its SDT and RHT.

UL suggests which type of protection scheme (dedicated or shared) a connection needs. Thus, ECM reprovisions connections' backup resources by dedicated protection if their UL is above a certain threshold (UL_{th}). ECM first frees backup resources (removing backup routes of connections and then terminating lightpaths which do not carry any traffic), then sorts the existing connections with respect to their ULs in descending order. Starting from the most urgent connection, ECM reprovisions them by dedicated protection if their ULs are above UL_{th} and by shared protection if their ULs are below UL_{th}[4]. Algorithm 2 shows the heuristic for reprovisioning. Providing a backup path (grooming on a sequence of new and/or existing lightpaths) based on available resources, where a connection's primary path (a sequence of lightpaths) is given, is well studied in the literature for both dedicated and shared protections [4], [5], so details are omitted.

Algorithm 2 Backup Reprovisioning

1: Create set of connections to be reprovisioned (T_{rp}) by adding all existing connections ($T_{rp} = T$)
2: Free all backup resources.
3: Take first connection t from T_{rp}.
4: **if** $UL_t > UL_{th}$, **then**
5: Find backup resources to provide dedicated PAC [5].
6: **if** enough resources exist to provide dedicated protection for connection t, **then** provision dedicated backup resources, **else** go to Step 8.
7: **else**
8: Find backup resources by shared PAC [4].
9: **if** enough resources exist to provide shared protection for connection t, **then** provision backup resources, **else** terminate algorithm and return failure.
10: **end if**
11: Remove connection t from T_{rp}, **if** $|T_{rp}| = 0$, **then** terminate, **else** go to Step 3.

V. Hold-Lightpath Scheme

In a dynamic traffic scenario, when all connections traversing the same lightpath terminate, this lightpath is typically terminated also. In this study, we can hold this empty lightpath, since this kind of exploitation of excess grooming ports helps to increase the availability by creating a segmented protection [9], [20] (explained below) and to decrease setup time of future connections. Thus, we introduce the hold-lightpath scheme where the lightpaths are held, even though there is no connection supported, until a backup reprovisioning event occurs.

Figure 6 illustrates the hold-lightpath scheme. A connection request t_1 (1, 5, STS-48c, A_{t_1}) is provisioned on a primary (w_1) and a backup lightpath (b_1) (Fig. 6(a)), and, before the new connection request t_2 arrives, t_1 is terminated. Fig. 6(b) shows the provisioning of new connection request t_2 (1, 10, STS-12c, A_{t_2}) without hold-lightpath scheme. New lightpaths (w_2 and b_2) are lit up from node 1 to node 10 to provision t_2. Figure 6(c) shows the provisioning of t_2 with hold-lightpath scheme. After termination of t_1, we hold the lightpaths w_1 and b_1 shown in Fig. 6(a) and light up new lightpaths (w_3 and b_3) from node 5 to node 10 (one for primary and one for backup). The primary and backup routes of t_2 are groomed on (w_1, w_3) and (b_1, b_3) sequences. This scheme creates segmented protection that increases connection availability [20]. For instance, in Fig. 6(b) (without hold-lightpath scheme), connection availability of t_2 can be estimated by $a_{w_2} + (1 - a_{w_2})a_{b_2}$, where a is availability of a path, while, in Fig. 6(c), connection availability of t_2 is $(a_{w_1} + (1 - a_{w_1}a_{b_1}))(a_{w_3} + (1 - a_{w_3})a_{b_3})$. If each link has 0.99 availability, then hold-lightpath increases t_2's connection availability from 0.99855 to 0.99999.

Assuming a connection is set up after primary and backup paths are established, setup time is equal to provisioning (and/or grooming) the connection on backup paths, as follows:

$$\sum_{b \in B} (1 + m_b) \times M + X_b \left[(1 + m_b) \times C + D_b\right] \quad (8)$$

where B is set of backup lightpaths; m_b and D_b are the number of links and propagation delay on path b, respectively; X_b is a binary number, which is equal to 1 if b is a new lightpath and 0 if b already exists; M is message processing delay at each node; and C is configuration delay at each optical crossconnect (OXC). For instance, the setup time for t_2 in Fig. 6(b) is $6C+6M+D_{b_2}$. With hold-lightpath scheme (Fig. 6(c)), where $D_{b_1} + D_{b_3} = D_{b_2}$, setup time is $6M+4C+D_{b_3}$, which is $2C + D_{b_1}$ shorter than without hold-p-lightpath scheme.

Surely, hold-lightpath scheme may increase grooming ports usage, but in the context of excess capacity management, this problem can be solved by backup reprovisioning where, before the reprovisioning, all the empty lightpaths held can be terminated.

VI. Illustrative Numerical Examples

We conduct numerical studies on a US-wide network (Fig. 7(a)), with 32 wavelengths/link in each direction and wavelength conversion (i.e., each node has an OEO switch). Capacity of each wavelength is STS-192 (10 Gpbs). The number of grooming ports at a node is set to the number of wavelengths times its nodal degree. Connection-arrival rate fluctuates during the day with relative generated traffic loads based on the population served by source and destination nodes, and different time zones are taken into consideration. The traffic intensity between an s-d pair is also calculated

[4]We consider UL_{th} to be the average of ULs of the existing connections.

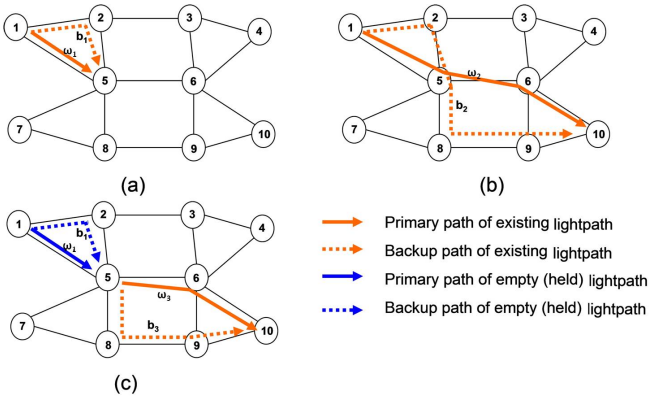

Fig. 6. (a) A connection request t_1 (1, 5, STS-48c, A_{t_1}) is provisioned. (b) t_1 is terminated and a new connection request t_2 (1, 10, STS-12c, A_{t_2}) is provisioned. (c) t_1 is terminated and lightpaths are held and t_2 is provisioned on these lightpath and new ones from nodes 5 to 10.

proportionally to the ratio of population served by the source and destination nodes over total population. Figure 7(b) gives an example of typical daily fluctuation of traffic with exponentially-distributed holding time (mean = one hour), generated by mimicking real traffic fluctuations (e.g., [21]. The number of connection requests for different bandwidth requirements follow the distribution STS-192c: STS-96c: STS-48c: STS-21c: STS-12c: STS-3c = 1: 2: 4: 10: 10: 20 [22]. The load of the network is equal to average arrival time $\times$ mean holding time $\times$ average bandwidth normalized to STS-192c. The SLA availability targets of connections are distributed as follows: 0.9999: 0.9995: 0.999: 0.99: 0.95 = 1: 5: 15: 30: 50. Links experience independent failures, and each link is in one of two categories with equal probability: failure-prone link (MTTR and mean time between failures (MTBF) are uniformly distributed over [1,4] and [100,400] hours, respectively), and rare-failure link (MTTR and MTBF are uniformly distributed over [4, 8] and [1000, 4000] hours, respectively) [23].

We consider the following penalty model which captures both penalty due to SLA violations and loss due to bandwidth blocking:

$$\text{Penalty} = \sum_{t \in T_b} (\delta_1 \times P_t \times B_t) + \sum_{t \in T_a} (\delta_2 \times B_t \times P_t \times SDT_t) \quad (9)$$

where δ_1 and δ_2 are monetary loss due to rejection of a unit bandwidth and SLA violation penalty per bandwidth per unit time, respectively (typically $\delta_1 \geq \delta_2$); T_a and T_b are the connections admitted and blocked, respectively; and P_t is penalty coefficient of connection t shown in Table I (SLA violations or bandwidth blocking of the connections that require higher availability cause higher penalty [19]).

TABLE I
PENALTY COEFFICIENT (P_t)

Availability Target	0.9999	0.9995	0.999	0.99	0.95
Penalty coefficient (P_t)	$3.0x$	$2.5x$	$2.0x$	$1.5x$	$1.0x$

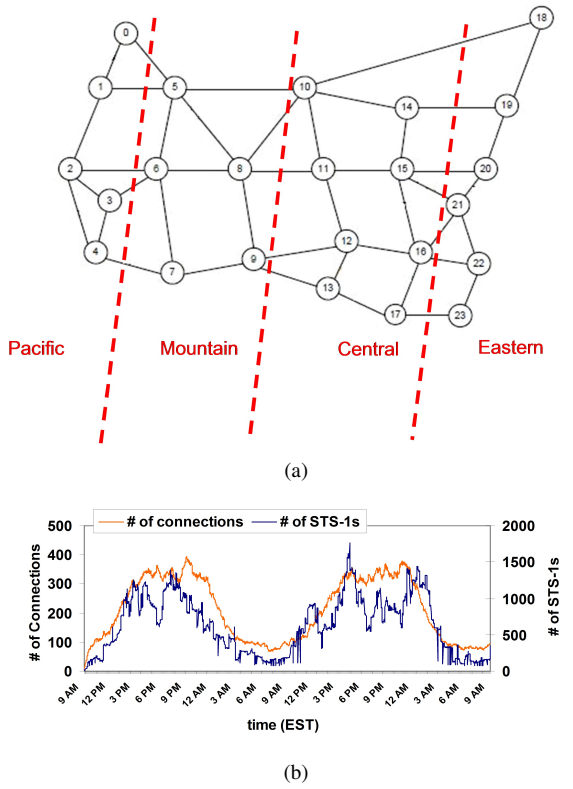

Fig. 7. (a) Representative US-wide network and (b) sample traffic variations over a two-day period.

A. Comparison

We compare PAC combined with our ECM approach (EC-PAC) with traditional approaches: dedicated PAC (D-PAC) and shared PAC (S-PAC), under the assumption that they also reprovision backup resources when they face resource exhaustion. Figure 8 shows cumulative downtime exceeding allowed downtimes (stated in SLAs) of connections and BBR. D-PAC shows very low SLA violations with rapidly increasing BBR with increasing network load. Since sharing decreases capacity consumption, S-PAC shows lower BBR compared to D-PAC, but causes higher SLA violations. EC-PAC shows low SLA violations for low loads. For high loads (>100 Erlangs), EC-PAC introduces more sharing, which causes increase in SLA violations, but it helps to decrease bandwidth blocking (as low as BBR for S-PAC). Therefore, EC-PAC exploits low SLA-violation performance of D-PAC for low loads, and low BBR advantage of S-PAC for high loads.

Figure 9 shows penalty results for different ratios of penalty parameters δ_1 and δ_2 (for δ_1/δ_2 =1 and 100 in Figs. 9(a) and (b), respectively). For lower loads, D-PAC shows lower penalty than S-PAC, and higher penalty for higher loads. Our approach EC-PAC exploits advantages of both schemes (low availability of D-PAC and low BBR of S-PAC) and shows low penalty close to D-PAC for low loads and to S-PAC for high loads. In fact, because of the hold-lightpath scheme and link protection provided by pre-provisioning, it shows much lower penalty than both D-PAC and S-PAC approaches. When δ_1/δ_2 increases, the penalty increases for both D-PAC and EC-PAC, but EC-PAC still adapts the state of protection by

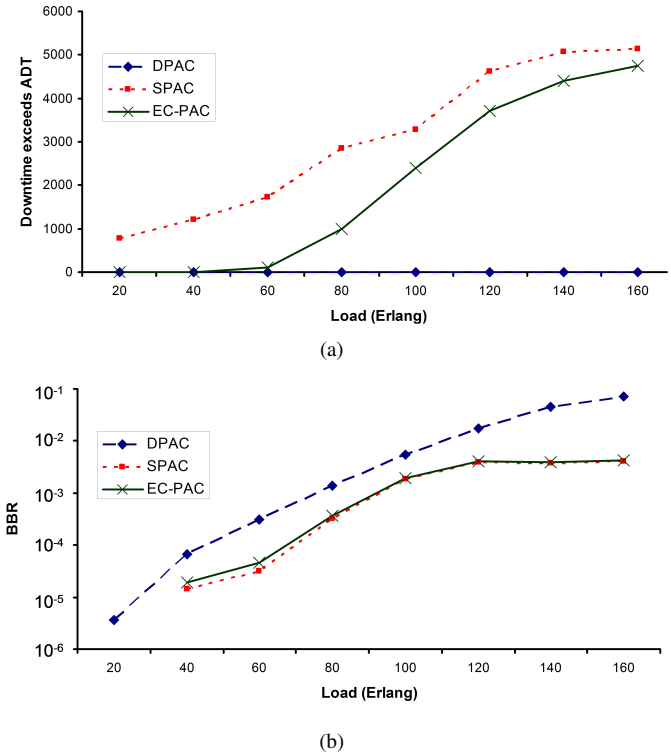

Fig. 8. Comparison of (a) SLA violation and (b) BBR.

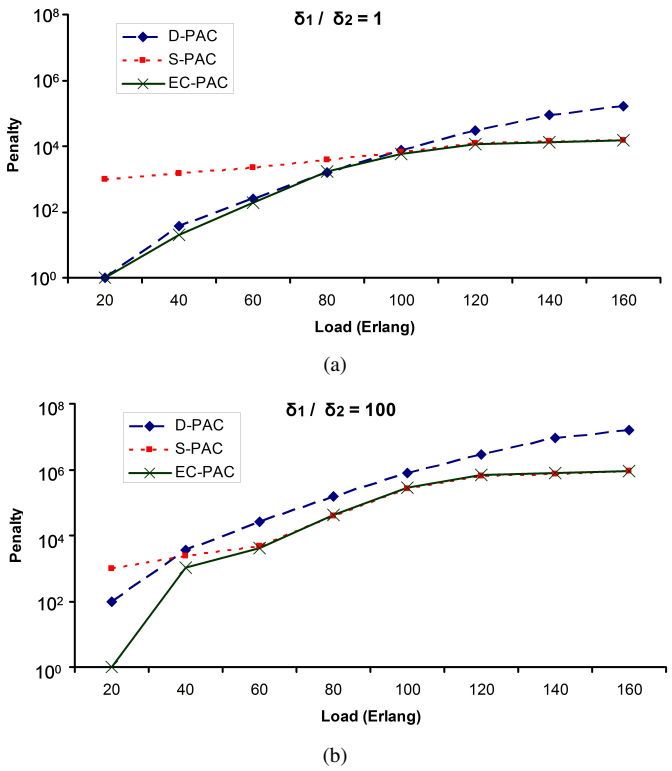

Fig. 9. Penalty results with cost ratio of (a) 1 and (b) 100.

using different protection schemes (D-PAC, S-PAC, and link protection) and provides the lowest penalty.

Figure 10 shows grooming-port utilization on network nodes over a two-day period when network load is 100 Erlang. Since EC-PAC exploits excess grooming ports by pre-provisioning and hold-lightpath schemes, its grooming port utilization is higher than D-PAC and S-PAC. However, the link utilization (ratio between bandwidth utilized for connections on a link over total bandwidth) of EC-PAC is adaptive to excess bandwidth. Figure 11 shows link utilization over the two-day period. Link utilizations for both D-PAC and S-PAC fluctuate with the traffic profile. When new traffic arrives (e.g., between 9 AM and 3 PM), EC-PAC shows high link utilization because it provides link protection on pre-provisioned network resources, which are reserved when traffic load is low (e.g., between 12 AM and 6 PM). When more traffic arrives and excess resources decrease (e.g., between 3 PM and 12 AM), EC-PAC introduces sharing with reprovisioning by considering the connections' ULs to avoid bandwidth blocking. Thus, for peak hours, link utilization decreases with EC-PAC. This adaptive feature of our ECM scheme provides less penalty for any network load.

The average connection setup times (calculated by Eq. (8)) for EC-PAC were found to be very small (around 1 ms) compared to traditional approaches (D-PAC around 12 ms, and S-PAC around 10 ms), because pre-provisioning and hold-lightpath schemes singificantly reduce setup time.

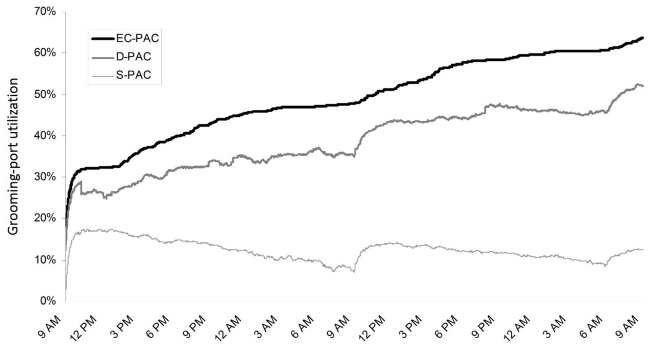

Fig. 10. Grooming-port utilization over a two-day period (network load = 100 Erlang).

B. *Effects of Pre-provisioning and Hold-Lightpath*

Figure 12 shows downtime reduction compared to S-PAC obtained by ECM with hold-lightpath and pre-provisioning (ECM-H-PP), ECM with hold-lightpath (ECM-H), ECM with pre-provisioning (ECM-PP), and ECM without either of hold-lightpath or pre-provisioning (ECM). Both pre-provisioning and hold-lightpath schemes help to decrease SLA violations, and ECM with both schemes significantly decreases SLA violations. The BBR results are omitted, because BBR of ECM approach slightly increases by pre-provisioning and hold-lightpath scheme (around 1.2%).

VII. Conclusion

In this study, we investigated the problem of exploiting excess capacity (EC), in terms of both bandwidth and groom-

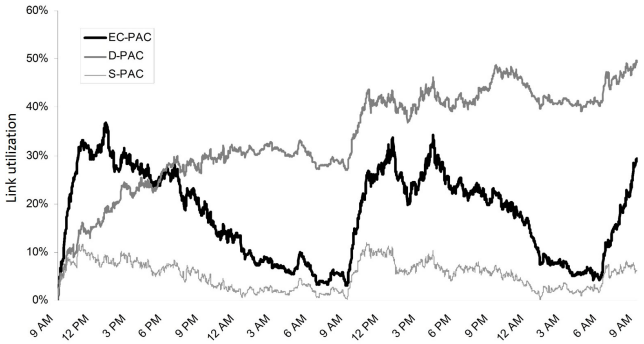

Fig. 11. Link utilization over a two-day period (network load = 100 Erlang).

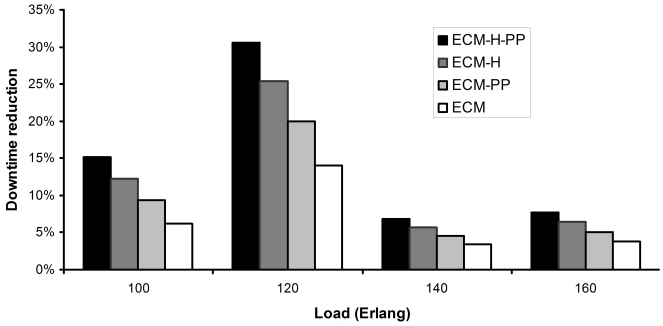

Fig. 12. Downtime reduction gained by pre-provisioning and hold-lightpath.

ing ports, to reduce penalties which may incur because of bandwidth blocking and SLA violations for survivable traffic grooming in optical WDM backbone networks, where connection requests are of sub-wavelength granularity and each provisioned request has to be protected from single-link failures. We developed an EC management (ECM) scheme which exploits EC by providing link protection in advance (pre-provisioning) statistically, admitting connections with appropriate protection scheme for sub-wavelength traffic (link protection, dedicated protection, or shared protection), and reprovisioning backup resources to avoid capacity exhaustion when traffic load increases. We also proposed a complementary method, namely hold-lighpath scheme, to increase connection availability and decrease connection setup time. The illustrative numerical examples showed that the ECM approach, by adapting the protection schemes depending on the EC in the network, reduces penalty paid by network operator for blocking and SLA violations, and decreases the connection setup time significantly.

REFERENCES

[1] F. Dikbiyik, L. Sahasrabuddhe, M. Tornatore, and B. Mukherjee, "Exploiting excess capacity to improve robustness of WDM mesh networks," *IEEE/ACM Trans. Netw.*, vol. 20, no. 1, pp. 114–124, Feb. 2012.

[2] F. Dikbiyik, M. Tornatore, L. Sahasrabuddhe, and B. Mukherjee, "Exploiting excess capacity, part II: Differentiated services under traffic growth," *submitted to IEEE/ACM Trans. Netw.* (under review). [Preliminary versions of this work were presented at the PS'10 and ANTS'10 conferences.].

[3] K. Zhu and B. Mukherjee, "Traffic grooming in an optical WDM mesh network," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 122–133, Jan. 2002.

[4] C. Ou, K. Zhu, H. Zang, L. Sahasrabuddhe, and B. Mukherjee, "Traffic grooming for survivable WDM Networks-shared protection," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 9, pp. 1367–1383, Nov. 2003.

[5] C. Ou, K. Zhu, J. Zhang, H. Zhu, B. Mukherjee, H. Zang, and L. Sahasrabuddhe, "Traffic grooming for survivable WDM Networks-dedicated protection," *OSA J. Opt. Netw.*, vol. 3, no. 1, pp. 50–74, Jan. 2004.

[6] B. Mukherjee, *Optical WDM Networks*. New York, NY: Springer, 2006.

[7] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE Network*, vol. 14, no. 6, pp. 16–23, Nov./Dec. 2000.

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

[9] D. Xu, Y. Xiong, and C. Qiao, "Novel algorithms for shared segment protection," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 8, pp. 1320–1331, Oct. 2003.

[10] D. Arci, G. Maier, A. Pattavina, D. Petecchi, and M. Tornatore, "Availability models for protection techniques in WDM networks," in *Proc. DRCN'03*, Oct. 2003, pp. 158–166.

[11] L. Zhou, M. Held, and U. Sennhauser, "Connection availability analysis of shared backup path-protected mesh networks," *J. Lightw. Technol.*, vol. 25, no. 5, pp. 1111–1119, May 2007.

[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] J. Zhang, K. Zhu, H. Zhang, N. Matloff, and B. Mukherjee, "Availability-aware provisioning strategies for differentiated protection services in wavelength-convertible WDM mesh networks," *IEEE/ACM Trans. Netw.*, vol. 15, no. 5, pp. 1111–1119, May 2007.

[14] J. M. Simmons, "Cost vs. capacity tradeoff with shared mesh protection in optical-bypass-enabled backbone networks," in *Proc. OSA NFOEC*, Anaheim, CA, Mar. 2007.

[15] E. Bouillet, J. 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.

[16] C. Assi, W. Huo, A. Shami, and N. Ghani, "On the benefits of lightpath reprovisioning in optical mesh networks," in *Proc. IEEE ICC*, Seoul, Korea, May 2005.

[17] C. Assi, W. Huo, and A. Shami, "Multiple link failures survivability of optical networks with traffic grooming capability," *Computer Comm.*, vol. 29, no. 18, pp. 3900–3912, 2006.

[18] F. Dikbiyik, M. Tornatore, and B. Mukherjee, "Exploiting excess capacity for survivable traffic grooming in optical WDM backbone networks," in *Proc. IEEE Globecom'11*, Houston, TX, Dec. 2011.

[19] M. Xia, M. Tornatore, C. U. Martel, and B. Mukherjee, "Service-centric provisioning in WDM backbone networks for the future internet," *J. Lightw. Technol.*, vol. 27, no. 12, pp. 1856–1865, June 2009.

[20] A. Todimala and B. Ramamurthy, "A dynamic partitioning sub-path protection routing technique in WDM mesh networks," in *Proc. ICCC*, Washington, DC, Aug. 2002.

[21] C. Labovitz. (2009) Arbor networks. [Online]. Available: http://ddos.arbornetworks.com/2009/08/what-europeans-do-at-night/

[22] S. Huang, M. Xia, C. Martel, and B. Mukherjee, "Survivable multipath traffic grooming in telecom mesh networks with inverse multiplexing," *IEEE/OSA J. Opt. Comm. Netw.*, vol. 2, no. 8, pp. 545–557, Aug. 2010.

[23] M. Xia, M. Tornatore, C. U. Martel, and B. Mukherjee, "Risk-aware provisioning for optical WDM mesh networks," *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 921–931, June 2011.