Introducing a Relative Priority for the Shared-Protection Schemes

Nizar Bouabdallah and Bruno Sericola

Abstract—One of the major challenges of optical network operators is ensuring the stringent levels of availability required by their highest class clients. To achieve this, we introduce relative priorities among the different primary connections contending for access to the shared-protection paths. In this paper, we provide an analytical model for the proposed priority-enabled scheme. As a key distinguishing feature from existing literature, we derive explicit analytic expressions for the average availability and service disruption rate for the different priority classes.

Index Terms—Optical networks, protection, performance analysis.

1 Introduction

Over the last decade, networks have been witnessing a perpetual growth in data traffic. Due to the new incumbent challenges, the operators are progressively migrating toward optical core networks taking advantage of the tremendous transmission capacity offered by the optical technology, thanks to the revolutionary Wavelength-Division Multiplexing (WDM) technology. Meanwhile, the relentless need for more capacity may have been satisfied. However, in such an environment, the cut of a fiber link can lead to a tremendous traffic loss.

In this regard, network survivability becomes a critical concern for operators. To alleviate this, backup resources are used to restore failed connections. These resources are usually shared among several primary connections to improve the network utilization. Generally, the primary connections are considered as equally important when contending for the use of the backup resources [1], [2], [3]. However, this solution is unsuitable from the service perspective. Indeed, the quality of service (QoS) required by different clients can be very different because of their diverse services' characteristics. For instance, banking services will require stringent reliability, whereas Internet Protocol (IP) best effort packet delivery services may be satisfied without a special constraint on reliability. One possible solution to provide different levels of reliability is to use a priority mechanism for the shared-protection scheme. Recently, some service differentiation schemes have been proposed in literature [4], [5]. The impact of these schemes on the system reliability is conducted based on simulations. Note that service reliability can be measured by means of two basic parameters: service availability and service disruption rate.

- *The authors are with the INRIA-IRISA, Campus Universitaire de Beaulieu F-35042, Rennes CEDEX, France.*
  E-mail: {nizar.bouabdallah, bruno.sericola}@inria.fr.

Manuscript received 23 June 2006; revised 23 Jan. 2007; accepted 3 May 2007; published online 29 May 2007.
For information on obtaining reprints of this article, please send e-mail to: tdsc@computer.org, and reference IEEECS Log Number TDSC-0086-0606.
Digital Object Identifier no. 10.1109/TDSC.2007.1008.

These parameters were assessed analytically in [6] for the single backup path (BP) shared-protection scheme (that is, the $1:N$ case). To achieve this, the authors adopted simplifying assumptions such as all the paths (primary and backup) are considered equally available (that is, having the same failure and repair rates). This assumption is not realistic. Specifically, the BP is usually chosen to be the longest link-disjoint path among all the precomputed ones. In so doing, the BP has the lowest availability compared to its associated primary paths. Likewise, the primary path availabilities of different connections are not the same and depend on their classes of service. Moreover, the authors in [6] limited their study to only two classes of service, whereas the majority of the standards deal with multiple (that is, more than two) classes of service. Generally, four classes of service can be identified (premium, gold, silver, and bronze), and according to the customer profiles' and the provided service, a finer classification can be achieved [7]. As our first main contribution, we therefore derive the reliability parameters associated to the $1:N$ protection scheme for multiple classes of service with different path availabilities.

Furthermore, [6] only handled the $1:N$ scheme. In view of this, as our second main contribution, in this paper, we give explicit formula of reliability parameters related to the multiple BPs protection scheme (that is, the $M:N$ scheme, with $M \geq 1$). In this case, analytical results are derived only for two classes of service.

Finally, previous works [4], [5], [6] suggested the use of strict preemptive priority to differentiate among different classes of connections. This kind of policy may not be desired, since it is extremely penalizing for low-priority classes. To cope with such limitation and as our third main contribution, we suggest a new policy that allows relative priority among different classes of service. We develop a new analytical model to analyze our proposal. Accordingly, all underlying policies are compared. We demonstrate that our new scheme provides operators much more flexibility to manage their networks in order to satisfy various client requirements. Specifically, it enables cost savings compared to existing solutions.

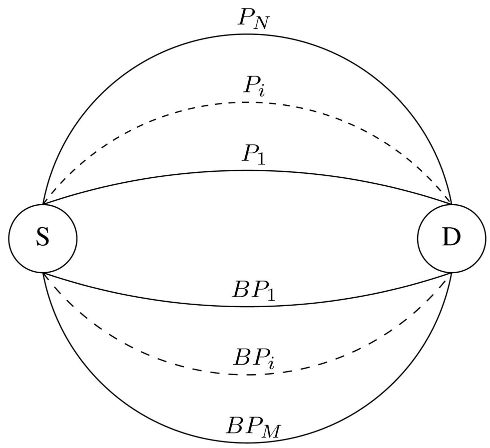

Fig. 1. N primary paths sharing M BPs (the $M:N$ system).

The rest of this paper is organized as follows: In Section 2, we propose and describe the priority-enabled shared-protection scheme. Section 3 discusses prior research related to this work. In Section 4, we introduce a mathematical model to evaluate the impact of the priority mechanism on the classical single BP shared-protection scheme (that is, the $1:N$ case). Throughout this section, we present numerical results to evaluate the benefits of the service differentiation feature introduced by the priority mechanism. In Section 5, the study is extended to multiple BPs shared-protection scheme (that is, the $M:N$ case). We introduce the concept of relative priority for the access to the BPs. Then, we derive the resulting reliability parameters considering two priority levels. Finally, conclusions are drawn in Section 6.

2 Relative Priority for Shared-Protection Scheme

Nowadays, widely deployed optical networks are circuit-switched oriented networks (that is, wavelength routed networks). In this context, shared protection has been established as a mechanism for providing considerable savings in terms of the number of provisioned wavelengths. Indeed, the shared protection is more efficient in terms of capacity allocation when compared with dedicated protection. However, what still lacks in the literature when dealing with shared protection is the capability to provide multiple grades of service for different clients with various requirements. This study therefore deals with QoS-aware shared-protection mechanisms in wavelength-routed networks (that is, WDM networks), where lightpaths with different QoS requirements can share wavelengths on their protection paths. Such a protection strategy needs to be implemented at the Generalized Multiprotocol Label Switching (GMPLS) control plane that calculates the primary and the BPs for each new arriving connection subject to its QoS requirements and the current network availability.

Assume N primary paths $P_1, \ldots, P_N$ sharing M BPs $BP_1, \ldots, BP_M$ (that is, the $M:N$ shared-protection system), as depicted in Fig. 1. In the classical shared-protection scheme (that is, without priority consideration), the first failed connection is recovered by the BPs until its primary path reparation, regardless of the QoS requirements of the subsequent failed connections.

On the other hand, considering the priority-enabled scheme, these connections are classified into K reliability classes $C_1, \ldots, C_K$, with N_i connections belonging to class C_i for $i = 1, \ldots, K$ and $N_1 + \cdots + N_K = N$. Moreover, class C_i has a higher priority than class C_j, as long as $i < j$.

The key idea behind our proposal is to provide each class C_i with a preemption quota M_i. Accordingly, a connection, say, t, belonging to class C_i loses its preemptive priority with respect to lower classes once the number of C_i-connections already restored by the BPs reaches the predefined quota M_i.

Doing so, once the primary path of the C_i-connection t breaks down, one of the free BPs, if any, is assigned to protect connection t, and restoration is ensured by switching t to the BP. Meanwhile, repair actions are performed on the primary path. Once repaired, the restored connection is switched back to its primary path.

On the contrary, if at the moment, the C_i-connection t fails, then all the operating BPs are already occupied, protecting other connections with higher priorities than t, and connection t becomes unavailable. However, if a part of the BPs is protecting connections belonging to classes with lower priorities than t (that is, classes with priority comprised between $i+1$ and K), then the one having the lowest priority is preempted immediately by connection t only if the current number of class C_i connections using the BPs is less than or equal to M_i. In this case, the preempted connection becomes unavailable, waiting for a BP to be released or for its primary path to be repaired.

The preemption is therefore allowed according to the current backup resources utilization. In view of this, this shared-priority scheme is called relative as opposed to the classical strict preemptive priority, where a connection t preempts instantaneously any other connection belonging to a lower priority class in order to access the BPs.

To illustrate the relative preemptive priority, we consider the simple example of two classes of service. Let us consider a connection t belonging to the high-priority class (that is, class C_1). As stated before, when t fails, and all the BPs are already occupied, it may be recovered by preempting a lower priority connection that already uses a BP. The preemption decision depends on the current backup resources utilization. Specifically, if the number of the first-class connections going through the BPs already exceeds a predefined quota M_1, $0 \leq M_1 \leq M$, then connection t is not allowed to preempt low-priority connections. Accordingly, connection t becomes unavailable. Otherwise, connection t preempts one of the possible low-priority connections using the BPs. If such connections do not exist, then connection t becomes unavailable. Note that when the threshold M_1 is set equal to M, we obtain the classical strict preemptive priority. Reducing the M_1 value penalizes the low-priority connections with respect to the high-priority ones less. When $M_1 = 0$, we get the classical shared-protection scheme without any service differentiation.

As stated before, our relative priority policy is more shifted toward shared protection in current deployed circuit-switched optical networks. In such networks, the bandwidth requirement of each lightpath can be a multiple of the wavelength capacity. Specifically, a C_i-connection requiring w wavelengths can be simply seen as w connections of class C_i, where each connection needs one wavelength.

This QoS-aware protection mechanism can also be extended to handle subwavelength granularity with minor modifications. Recall that nowadays, widely deployed all-optical wavelength routed networks are no more consistent with the packet switching philosophy of the Internet.

Specifically, in next-generation networks, packet-based data traffic of bursty nature will become prevalent. In this regard, many researches are now focusing on bringing the subwavelength switching concept into the optical domain. In such environment, multiple connections with various QoS levels could be groomed on the same lightpath. Once a lightpath fails, all the connections contained therein have to be rerouted through the backup resources based on their classes of service. According to whether traffic bifurcation among multiple BPs is allowed or not, each connection looks for available bandwidth on the backup resources and may be allowed to preempt lower class connections in order to be recovered.

3 Related Works

One commonly used control scheme to provide multiple grades of service in the context of shared resources is the partial sharing policy [8], [9]. This policy was originally proposed to protect each customer from the overloads caused by other customers sharing the same resource. To achieve this, each customer (class of service) is characterized by two parameters: the Upper Limit (UL) and Guaranteed-Minimum (GM) bounds on the number of requests that can be in the system simultaneously at any time. Accordingly, a new customer request is admitted to use the shared resources only if both conditions are satisfied:

- There is enough free capacity on the shared resource.
- After the admission of the new request, the number of requests from that customer is less than or equal to its specified UL bound.

Hence, using the partial sharing policy, the utilization of the shared resources (the BPs in our case of study) is completely controlled at the access level. Once a customer is admitted to use the shared resources, it occupies the allocated resource until the end of its service. In other words, the connection is never preempted by other connections.

However, using our relative priority policy, this connection could be preempted by a new failed connection with a higher class of service. This is the main key feature that differentiates our policy from the partial sharing one. In doing so, our policy does not limit the customer utilization of the shared resources to a hard bound, regardless of their current availabilities. Our policy allows rather customers to benefit from the total available shared capacity. Then, once the total shared capacity is exhausted, the lowest class customer could be preempted because of the arrival of a higher class client. The preemption feature therefore adds more flexibility to our sharing method. Accordingly, the availability of low-class connections is improved, compared to partial sharing policy, without affecting the high-class availabilities. Indeed, our method utilizes more efficiently the shared resources.

To illustrate this, we will present the simple example of $N = N_1 + N_2$ connections belonging to two classes of service ($N_1 \geq 2$ of class C_1 and $N_2 \geq 2$ of class C_2) sharing $M = 2$ BPs. First, the problem will be treated in the light of our proposed policy, considering that the relative priority is adopted within the network. Afterwards, it will be considered in the context of the partial sharing policy. In the former case, we consider the following configuration $M_1 = 1$ and $M_2 = 0$ to provide different grades of service between classes C_1 and C_2. The corresponding configuration using the partial sharing policy is $(UL_1 = 2, MG_1 = 1)$ and $(UL_2 = 1, MG_2 = 0)$.

As explained before, the latter partitioning policy tends to be relatively inefficient because the benefits of the complete sharing scheme are partially lost. In fact, this policy may block failed connections to access free backup resources, whereas this never happens with our proposed policy. Specifically, assume the case where two primary C2 connections are under failure simultaneously. If the partial sharing policy is taken into account, then only one connection would be recovered by the backup resources. The second connection will be rejected in spite of the fact that another BP still remains free. However, using our policy, both connections are recovered by the backup resources, thus increasing the availability of the lowest class connections.

4 Single Backup Path Shared-Protection Scheme

In this section, we present a mathematical model for the $1:N$ relative priority-enabled shared-protection scheme. In this case, the threshold M_i can take two values 0 or 1. When $M_i = 0$ for $i = 1, \ldots, K$, we get the classical $1:N$ shared-protection scheme without service differentiation, whereas $M_i = 1$ for $i = 1, \ldots, K$ leads to a strict preemptive priority among the different classes of service. Solving this latter model, we derive explicit expressions for the average availability and service disruption rate of each class of connections, resulting from deploying the aforementioned priority strategy. The availability of a connection is defined as the proportion of time that the connection is up, and the service disruption rate of a connection is defined as the rate at which an available connection becomes unavailable [6]. It is worth noting that in this particular case, where $M = 1$, we have to respect the following condition:

$$M_i \leq M_j, \quad \forall i > j$$

unless this can result in a better availability and disruption rate of class C_i compared to the higher class C_j, which is an absurdity. Hence, in its general form, the quotas M_i can be written as follows:

$$M_1 = 1, \ldots, M_j = 1, M_{j+1} = 0, \ldots, M_K = 0.$$

In this case, the connection belonging to classes $C_j, \ldots, C_K$ can be grouped in the same class j. Consequently, the system can be analyzed as the $1:N$ system with a strict preemptive policy among the j classes of service.

4.1 Basic Assumptions

We use the following classical assumptions in our study:

- A path (primary or backup) has only two states: it is either operating (up) or nonoperating (down).
- A primary path is said unavailable when it is down.

- A BP is said to be unavailable to restore connection t when it is down or already occupied with recovering a nonpreemptable connection t' (for example, connection t' with a higher priority than t).
- A connection t has only two states: it is either available or unavailable. t is unavailable only when its primary path and the BP are unavailable.
- For each path, the successive times to failure and repair times form an alternating renewal process.
- All the paths are supposed to be statistically independent from the failure and repair viewpoint. In other words, there are enough resources to repair simultaneously any number of failed connections. This is known in the literature as unlimited repair [10].

4.2 The Analytical Model

Assume that the N_i primary paths of each class C_i have identical failure and repair rates denoted, respectively, by λ_i and μ_i. The availability p_i of a primary path belonging to class C_i is thus

$$p_i = \frac{\mu_i}{\lambda_i + \mu_i}.$$

The unavailability q_i of a class C_i primary path is given by $q_i = 1 - p_i$.

In the same way, the BP is characterized by its own failure and repair rates λ_b and μ_b. The BP is generally chosen to be the longest link-disjoint path among the $N + 1$ precomputed ones. Its availability is

$$p_b = \frac{\mu_b}{\lambda_b + \mu_b}$$

and its unavailability is $q_b = 1 - p_b$.

In the following, we derive the analytic expressions for the availability and the service disruption rate for each connection according to its class of service. We begin by calculating the unavailability U_i of a connection t belonging to class C_i.

Theorem 1. *For every $i = 1, \ldots, K$, the unavailability U_i of a connection t belonging to class C_i is given by*

$$U_i = q_i - \Pr\{t \text{ is restored by the BP}\}.$$

Proof. For the sake of simplicity, we denote by $\{pp(t) = 0\}$ the event that the primary path of t is down, whose probability is equal to q_i if t belongs to class C_i. By definition of the unavailability of a connection t, we have $\{t \text{ is unavailable}\} \subseteq \{pp(t) = 0\}$, thus

$$\begin{aligned} U_i &= \Pr\{t \text{ is unavailable}\} \\ &= \Pr\{t \text{ is unavailable}, pp(t) = 0\} \\ &= \Pr\{pp(t) = 0\} - \Pr\{t \text{ is available}, pp(t) = 0\} \\ &= q_i - \Pr\{t \text{ is restored by the BP}\} \end{aligned}$$

which concludes the proof. □

Henceforth, we say that t is a C_i-connection if it belongs to class C_i. Since all the C_i-connections behave identically, U_i can also be written as

$$U_i = q_i - \frac{1}{N_i} \Pr\{\text{a } C_i\text{-connection is restored by the BP}\}.$$

Note that the BP restores a C_i-connection if and only if all the independent events A, B, and C occur, where

$$\begin{aligned} A &= \{\text{The BP is up}\}, \\ B &= \{\text{All the primary paths of } C_j\text{-connections} \\ &\qquad \text{with } j < i, \text{are up}\}, \\ C &= \{\text{At least one } C_i\text{-connection is down}\}. \end{aligned}$$

Hence, U_i can be written as

$$\begin{aligned} U_i &= q_i - \frac{1}{N_i} \Pr\{A, B, C\} \\ &= q_i - \frac{1}{N_i} \Pr\{A\} \Pr\{B\} \Pr\{C\} \\ &= q_i - \frac{1}{N_i} p_b \left(1 - p_i^{N_i}\right) \prod_{j=1}^{i-1} p_j^{N_j}. \end{aligned}$$

In the particular case where all the primary connections belong to the same class of service and, thus, have the same primary path availability $p_i = p$, the unavailability U of each connection becomes

$$U = q - \frac{p_b\left(1 - p^N\right)}{N}, \tag{1}$$

where $q = 1 - p$. The availability A of a connection is $A = 1 - U$.

According to (1), we can see that U is simply the unavailability of the primary path of t reduced by the value of the availability introduced by the shared BP. Moreover, it is easy to show that the availability A increases with p_b and decreases with the sharability index N. We easily get

$$\lim_{N \longrightarrow \infty} A = p.$$

Indeed, when N is large, the supplement of availability introduced by the BP becomes negligible. This behavior is illustrated in Fig. 2a, where we consider three variants of BPs with different values of cut rates. In this case, the primary paths are supposed to be identical and characterized by the same cut rate $\lambda = 1/250\ h^{-1}$. In addition, the mean time to repair $(MTTR = 1/\mu)$ of all the paths is considered equal to 12 hours (see [5]).

Let us now derive the service disruption rate S_i for each class C_i. Recall that the disruption rate of a C_i connection t is defined as the rate at which connection t becomes unavailable from an available state. Connection t is available in the following disjoint cases:

- The primary path of t is up.
- The primary path of t is down, and connection t is restored by the BP.

In the former case, connection t transitions to the unavailable state only if the primary path of t breaks down, and the BP is already unavailable to restore C_i-connections. In this case, t has no possibility of being restored by the BP, and it transitions to an unavailable state at a rate equal to the failure rate of its primary path λ_i. In the latter case, such transition happens when the BP fails at a rate equal to λ_b or when one of the higher priority primary connections breaks

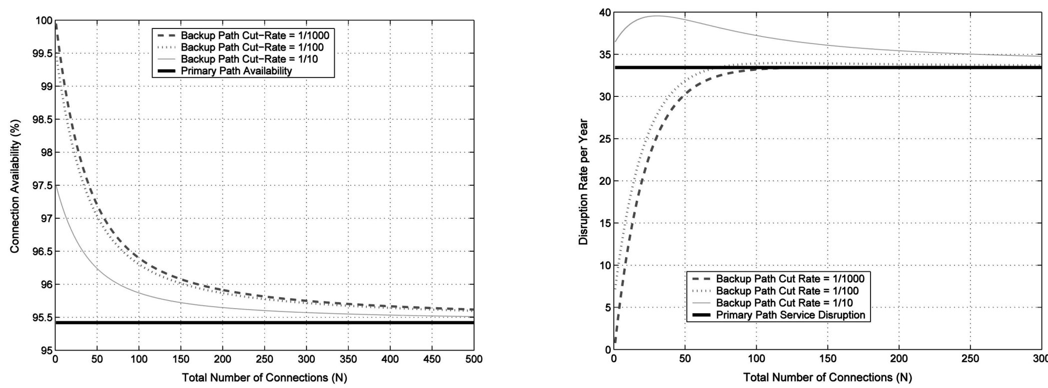

(a) (b)

Fig. 2. Classical (without priority) $1:N$ shared-protection scheme. (a) Availability. (b) Disruption rate (per year).

down and, thus, preempts the connection t at a rate equal to $N_1\lambda_1 + \cdots + N_{i-1}\lambda_{i-1}$. Hereafter, we write $\{pp(t) = 1\}$ to describe the event that the primary path of t is up. The expression of the disruption rate S_i is thus given by

$$S_i = \lambda_i \Pr\{pp(t) = 1, \text{the BP is unavailable}\} + \left(\lambda_b + \sum_{j=1}^{i-1} N_j\lambda_j\right) \Pr\{\text{t is restored by the BP}\}.$$

We already got the probability that connection t is restored by the BP, that is,

$$\Pr\{\text{t is restored by the BP}\} = \frac{1}{N_i} p_b\big(1 - p_i^{N_i}\big) \prod_{j=1}^{i-1} p_j^{N_j}.$$

Using the same argument, we obtain

$$\Pr\{pp(t) = 1, \text{the BP is unavailable}\} = p_i - p_b \prod_{j=1}^{i} p_j^{N_j}.$$

The disruption rate of a C_i-connection is then given by

$$S_i = \frac{1}{N_i} p_b \left(\lambda_b + \sum_{j=1}^{i-1} N_j\lambda_j\right) \big(1 - p_i^{N_i}\big) \prod_{j=1}^{i-1} p_j^{N_j} + \lambda_i \left(p_i - p_b \prod_{j=1}^{i} p_j^{N_j}\right).$$

In the particular case where all the connections belong to the same class of service, the disruption rate $S_i = S$ is given by

$$S = \frac{1}{N} \lambda_b p_b \big(1 - p^N\big) + \lambda p \big(1 - p_b p^{N-1}\big).$$

It is easy to show that S increases when p_b decreases and that we have

$$\lim_{N \longrightarrow \infty} S = \lambda p.$$

This behavior is illustrated in Fig. 2b. In fact, when N is large, the impact of the BP is negligible. One interesting finding is related to the impact of the BP availability p_b on the disruption rate. Specifically, the usage of a BP to protect a set of connections does not necessarily mean a better disruption rate performance. We can observe in Fig. 2b that unlike the availability, which is always improved, thanks to the backup protection (see Fig. 2a), the disruption rate may become worse due to the BP introduction.

For instance, when $\lambda_b = 1/10\ h^{-1}$, the service disruption rate experienced by a connection belonging to a 1:40 system is approximately equivalent to a mean of 37 service disruptions per year. Without protection, the same connection experiences fewer disruptions per year, with exactly a mean of 34 service disruptions. This happens because of the relatively high cut-rate value of the BP with respect to the primary path's cut rate $\lambda = 1/250\ h^{-1}$.

4.3 Numerical Results

In this section, we evaluate the benefits introduced by our priority scheme. To achieve this, we consider a scenario consisting of N primary connections sharing a common BP. We first consider the priority-enabled protection scheme (that is, $M_i = 1$ for $i = 1, \ldots, K$), with $K = 3$ classes of service. Each of the highest classes of service (that is, C_1 and C_2) contains only one connection. The remaining $N - 2$ connections, which are varied from 1 to 10, belong to the lowest class C_3. Then, for comparison purposes, the classical shared-protection scheme without service differentiation is applied to this scenario. It is important to note that the MTTR $(1/\mu)$ of all the paths is taken equal to 12 hours. We also consider a reference path cut rate $\lambda = \lambda_b = 0.0002\ h^{-1}$.

Fig. 3a shows, as expected, that increasing the number of the lowest priority primary connections does not affect the higher priority connections, which maintain the same availability levels. On the other hand, when using the classical scheme, all the connections are penalized, as they become less available. As such, the availability required by high-priority connections is not respected, despite the use of a BP. Specifically, according to [7], a gold client requests an availability of 99.999 percent. This availability is never achieved with the classical scheme. However, when the priority mechanism is enabled, this target availability can be obtained for the highest priority connections. This proves that by deploying the proposed scheme, the gold connections provisioning becomes possible in the network.

Likewise, Fig. 3b shows that when applying our priority scheme, the service disruption rate for the two highest priority connections remains unchanged when the number of sharing primary connections increases. Indeed, our

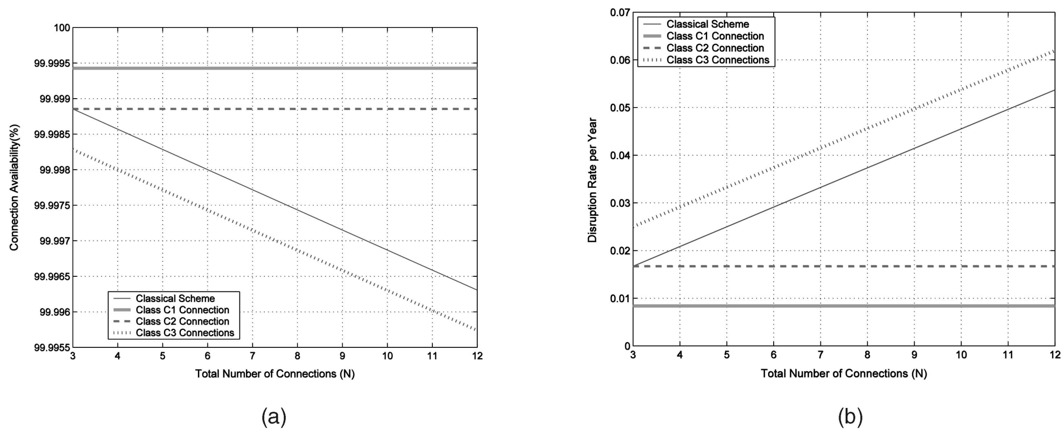

Fig. 3. Classical (without priority) and priority-enabled $1:N$ shared-protection schemes. (a) Availability. (b) Disruption rate (per year).

priority scheme is able to satisfy stringent high-priority connection requirements, whereas this objective cannot be achieved by the classical protection scheme.

5 Extension to Multiple Backup Paths Shared-Protection Scheme

This section extends the obtained results above to the general $M:N$ protection case. In other words, we derive the availability and service disruption rate of connections protected by $M \geq 1$ BPs. Such a scheme allows the system to recover against multiples failures, which become increasingly probable [11], [12] as the size and the complexity of optical mesh networks continue to grow. In what follows, the results are derived for the classical shared-protection schemes (without priority and with strict preemptive priority) and for our relative priority-enabled scheme.

5.1 The Analytical Model

Let us consider the $M:N$ priority-enabled shared-protection system proposed in Section 2, with N connections divided into two sets of reliability classes C_1 and C_2 with N_1 and N_2 connections, respectively. Connections of class C_1 have higher priority than connections of class C_2. We assume that all the primary paths behave identically, so they have the same availability $p = \mu/(\lambda + \mu)$, and we set $q = 1 - p$. These N primary connections are sharing M identical and independent BPs. We always use the assumptions described in Section 3.1. Let us consider the stochastic process $\{X(t), t \geq 0\}$ representing the up and down states of the paths. The states of that process are denoted by the triplet (n_1, n_2, m), where n_1 and n_2 are, respectively, the number of high- and low-priority failed primary paths, and m is the number of operating BPs. We denote by X the stationary state of the process $X(t)$ and by $p(n_1, n_2, m)$ the stationary distribution of states (n_1, n_2, m). Since all the paths are statistically independent, we have

$$p(n_1, n_2, m) = p_1(n_1) p_2(n_2) p(m),$$

where

$$p_1(n_1) = \binom{N_1}{n_1} q^{n_1} p^{N_1 - n_1}, p_2(n_2) = \binom{N_2}{n_2} q^{n_2} p^{N_2 - n_2},$$
$$p(m) = \binom{M}{m} p_b^m q_b^{M-m}.$$

In the following, we derive the analytic expressions for the availability and the service disruption rate for each connection according to its priority class. Again, we underline that the N_1 high-priority connections cannot preempt instantaneously all the other connections belonging to the lower priority classes in the utilization of the BPs. The preemption decision depends on the current system state. As explained in Section 2, a low-priority connection is preempted by a high-priority connection only if the number of high-priority connections currently restored by the BPs is less than or equal to a predefined quota M_1, which is such that $M_1 \leq N_1$.

5.1.1 Unavailability

We will begin by considering high-priority connections. Specifically, a connection t belonging to class C_1 is unavailable when both of the following conditions are verified:

- A: The primary path of t is down.
- B: t is not restored by the BPs.

The unavailability U_1 of a connection t of class C_1 is thus

$$\begin{aligned} U_1 &= \Pr\{A, B\} \\ &= \sum_{n_1=1}^{N_1} \sum_{n_2=0}^{N_2} \sum_{m=0}^{M} \Pr\{B | A, X = (n_1, n_2, m)\} \\ &\quad \times \Pr\{A | X = (n_1, n_2, m)\} p(n_1, n_2, m). \end{aligned}$$

Clearly, we have

$$\Pr\{A | X = (n_1, n_2, m)\} = \frac{n_1}{N_1}$$

and the cases where $\Pr\{B | A, X = (n_1, n_2, m)\}$ is nonnull are the following:

- $m \leq M_1$ and $m < n_1$. In this case, only m primary C_1 connections among the n_1 ones under failure can be restored. Thus, we have

$$\Pr\{B | A, X = (n_1, n_2, m)\} = 1 - \frac{m}{n_1}.$$

- $m > M_1$ and $n_1 \geq m$. In this case, the first set of M_1 BPs is assigned to the high-priority failed connections. The remaining $n_1 + n_2 - M_1$ failed connections belonging to both classes of service contend with equal priority to the remaining second set of $m - M$ BPs. As such, t is not restored by both sets of BPs with a probability

$$\Pr\{B|A, X = (n_1, n_2, m)\} = \frac{(n_1 - M_1)(n_1 + n_2 - m)}{n_1(n_1 + n_2 - M_1)}.$$

- $m > n_1 > M_1$, and $n_1 + n_2 > m$. The same arguments lead to the same result, that is

$$\Pr\{B|A, X = (n_1, n_2, m)\} = \frac{(n_1 - M_1)(n_1 + n_2 - m)}{n_1(n_1 + n_2 - M_1)}.$$

Putting together these results, we get

$$U_1 = \frac{1}{N_1}\left[\sum_{m=0}^{M_1}\sum_{n_1=m+1}^{N_1}(n_1 - m)p_1(n_1)p(m) + \sum_{m=M_1+1}^{M}\sum_{n_1=m}^{N_1}\sum_{n_2=0}^{N_2}(n_1 - M_1)\frac{n_1 + n_2 - m}{n_1 + n_2 - M_1}p(n_1, n_2, m) + \sum_{m=M_1+2}^{M}\sum_{n_1=M_1+1}^{m-1}\sum_{n_2=m-n_1}^{N_2}(n_1 - M_1)\frac{n_1 + n_2 - m}{n_1 + n_2 - M_1}p(n_1, n_2, m)\right].$$

Note that when $M_1 = 0$, we get the classical $M : N$ protection scheme without priority among the $N_1 + N_2 = N$ primary connections, which gives

$$U = \frac{1}{N}\sum_{m=0}^{M}\sum_{n=m+1}^{N}(n - m)p'(n)p(m), \tag{2}$$

where

$$p'(n) = \binom{n}{N}q^n p^{N-n}.$$

Following the same reasoning, we obtain the unavailability U_2 of a C_2-connection as

$$\begin{aligned} U_2 &= \Pr\{A, B\} \\ &= \sum_{n_1=0}^{N_1}\sum_{n_2=1}^{N_2}\sum_{m=0}^{M}\Pr\{B|A, X = (n_1, n_2, m)\} \\ &\quad \times \Pr\{A|X = (n_1, n_2, m)\}p(n_1, n_2, m). \end{aligned}$$

Clearly, we have

$$\Pr\{A|X = (n_1, n_2, m)\} = \frac{n_2}{N_2}$$

and the cases where $\Pr\{B|A, X = (n_1, n_2, m)\}$ is nonnull are the following:

- $m \leq M_1$ and $m < n_1$. In this case, all the BPs are used to restore C_1-connections; thus,

$$\Pr\{B|A, X = (n_1, n_2, m)\} = 1.$$

- $n_1 \leq M_1$, $n_1 \leq m$, and $n_1 + n_2 > m$. In this case, only $m - n_1$ low-priority primary connections among the n_2 ones under failure can be restored. Hence, we get

$$\Pr\{B|A, X = (n_1, n_2, m)\} = 1 - \frac{m - n_1}{n_2}.$$

- $M_1 < m \leq n_1$. In this case, M_1 BPs are affected to the high-priority connections. The remaining spare capacity (that is, $m - M_1$) is shared without any priority policy among the $n_1 + n_2 - M_1$ failed connections. Thus,

$$\Pr\{B|A, X = (n_1, n_2, m)\} = 1 - \frac{m - M_1}{n_1 + n_2 - m}.$$

- $M_1 < n_1 < m$ and $n_1 + n_2 > m$. The same arguments lead to the same result, that is,

$$\Pr\{B|A, X = (n_1, n_2, m)\} = 1 - \frac{m - M_1}{n_1 + n_2 - m}.$$

Putting together these results, we get

$$U_2 = \frac{1}{N_2}\left[\sum_{m=0}^{M_1}\sum_{n_1=m+1}^{N_1}\sum_{n_2=0}^{N_2}n_2 p(n_1, n_2, m) + \sum_{m=0}^{M}\sum_{n_1=0}^{M_1\wedge m}\sum_{n_2=m-n_1}^{N_2}(n_1 + n_2 - m)p(n_1, n_2, m) + \sum_{m=M_1+1}^{M}\sum_{n_1=m}^{N_1}\sum_{n_2=0}^{N_2}n_2\frac{n_1 + n_2 - m}{n_1 + n_2 - M_1}p(n_1, n_2, m) + \sum_{m=M_1+2}^{M}\sum_{n_1=M_1+1}^{m-1}\sum_{n_2=m-n_1}^{N_2}n_2\frac{n_1 + n_2 - m}{n_1 + n_2 - M_1}p(n_1, n_2, m)\right].$$

By setting $M_1 = 0$ and $N_1 + N_2 = N$, we get the unavailability of the classical $M : N$ protection scheme without priority, already given by relation (2).

5.1.2 Disruption Rate

We now consider the service disruption rate for each class of service. We first consider the high-priority connections. Thus, let t be a C_1-connection. Such connection is available in two cases:

- A: The primary path of t is up.
- $B = B_1 \cap B_2$, where B_1: The primary path of t is down, and $B_2 : t$ is restored by one of the BPs.

In case A, connection t becomes unavailable if its primary path fails (at rate λ), whereas all the BPs are unavailable for its restoration. In case B, connection t becomes unavailable if its restoring BP fails (at rate λ_b), and all the other BPs are unavailable for its restoration. We thus have the relation

$$S_1 = \lambda\Pr\{A, C\} + \lambda_b\Pr\{B_1, B_2, C\},$$

where the event C is defined by

- C: All the BPs are unavailable to restore t.

By conditioning on the stationary state $X = (n_1, n_2, m)$, we get

$$\Pr\{A, C\} = \sum_{n_1=0}^{N_1} \sum_{n_2=0}^{N_2} \sum_{m=0}^{M} \Pr\{C|A, X = (n_1, n_2, m)\} \times \Pr\{A|X = (n_1, n_2, m)\} p(n_1, n_2, m).$$

It is easy to check that

$$\Pr\{A|X = (n_1, n_2, m)\} = 1 - \frac{n_1}{N_1}$$

and

$$\Pr\{C|A, X = (n_1, n_2, m)\} = \begin{cases} 1 & \text{if } M_1 \le n_1 < m \text{ and } n_1 + n_2 \ge m \\ & \text{In this case, the high-priority class already uses its minimum quota } M_1, \text{ and all the BPs are occupied; hence, } t \text{ is not allowed to preempt protected low-priority connections subsequent to its primary path failure.} \\ 1 & \text{if } n_1 \ge m \\ 0 & \text{otherwise.} \end{cases}$$

In the same way, we have

$$\Pr\{B_1, B_2, C\} = \sum_{n_1=0}^{N_1} \sum_{n_2=0}^{N_2} \sum_{m=0}^{M} \Pr\{C|B_1, B_2, X = (n_1, n_2, m)\} \times \Pr\{B_2|B_1, X = (n_1, n_2, m)\} \Pr\{B_1|X = (n_1, n_2, m)\} \times p(n_1, n_2, m).$$

We get

$$\Pr\{B_1|X = (n_1, n_2, m)\} = \frac{n_1}{N_1}$$

and

$$\Pr\{C|B_1, B_2, X = (n_1, n_2, m)\} = \begin{cases} 1 & \text{if } M_1 < n_1 < m \text{ and } n_1 + n_2 \ge m \\ & \text{Note that unlike the probability } \Pr\{C|A, X = (n_1, n_2, m)\}, \text{ if } n_1 = M_1, \text{ then connection } t \text{ is allowed to preempt a protected low-priority connection following to its backup path failure since the quota } M_1 \text{ is not reached yet.} \\ 1 & \text{if } n_1 \ge m \\ 0 & \text{otherwise.} \end{cases}$$

In the same way, we obtain $\Pr\{B_2|B_1, X = (n_1, n_2, m)\}$

$$= \begin{cases} 1 & \text{if } n_1 \le M_1 \text{ and } n_1 \le m \\ 1 & \text{if } M_1 \le n_1 \le m \text{ and } n_1 + n_2 \le m \\ \frac{m}{n_1} & \text{if } m \le n_1 \text{ and } m \le M_1 \\ \frac{M_1}{n_1} + \left(1 - \frac{M_1}{n_1}\right)\left(\frac{m - M_1}{n_1 + n_2 - M_1}\right) & \text{if } (M_1 < m \le n_1) \text{ or } (M_1 < n_1 < m \text{ and } n_1 + n_2 > m) \\ 0 & \text{otherwise.} \end{cases}$$

This leads to the following expression of S_1:

$$\begin{aligned} S_1 = \frac{\lambda}{N_1} & \left[\sum_{m=0}^{M} \sum_{n_1=m}^{N_1} (N_1 - n_1) p_1(n_1) p(m) + \sum_{m=M_1+1}^{M} \sum_{n_1=M_1}^{m-1} \sum_{n_2=m-n_1}^{N_2} (N_1 - n_1) p(n_1, n_2, m) \right] \\ & + \frac{\lambda_b}{N_1} \left[\sum_{m=0}^{M_1} \sum_{n_1=m}^{N_1} m p_1(n_1) p(m) + \sum_{m=M_1+1}^{M} \sum_{n_1=m}^{N_1} \sum_{n_2=0}^{N_2} \left(M_1 + \frac{(n_1 - M_1)(m - M_1)}{n_1 + n_2 - M_1} \right) \times p(n_1, n_2, m) \right. \\ & \left. + \sum_{m=M_1+2}^{M} \sum_{n_1=M_1+1}^{m-1} \sum_{n_2=m-n_1}^{N_2} \left(M_1 + \frac{(n_1 - M_1)(m - M_1)}{n_1 + n_2 - M_1} \right) \times p(n_1, n_2, m) \right]. \end{aligned}$$

Based on this result, we can simply derive the expression of the disruption rate for the classical shared-protection scheme without priority by setting $M_1 = 0$ and $N_1 + N_2 = N$. In doing so, we get

$$S = \frac{1}{N} \sum_{m=0}^{M} \sum_{n=m}^{N} (\lambda(N - n) + \lambda_b m) p'(n) p(m).$$

Following the same reasoning, the disruption rate S_2 of a low-priority connection can be expressed as

$$S_2 = \lambda \Pr\{A, C\} + \lambda \sum_{n_1=0}^{M_1-1} (N_1 - n_1) \Pr\{B, C, X_1 = n_1\} + \lambda_b \sum_{n_1=0}^{M_1} n_1 \Pr\{B, C, X_1 = n_1\} + \lambda_b \Pr\{B, C\},$$

where X_1 is the first component of vector X, that is, the number of high-priority failed primary paths. Note that the events A, B, and C are those defined for the evaluation of S_1, but now, they concern a connection $t \in C_2$ and not $t \in C_1$. As done previously, we have

$$\Pr\{A, C\} = \sum_{n_1=0}^{N_1} \sum_{n_2=0}^{N_2-1} \sum_{m=0}^{M} \Pr\{C|A, X = (n_1, n_2, m)\} \times \Pr\{A|X = (n_1, n_2, m)\} p(n_1, n_2, m).$$

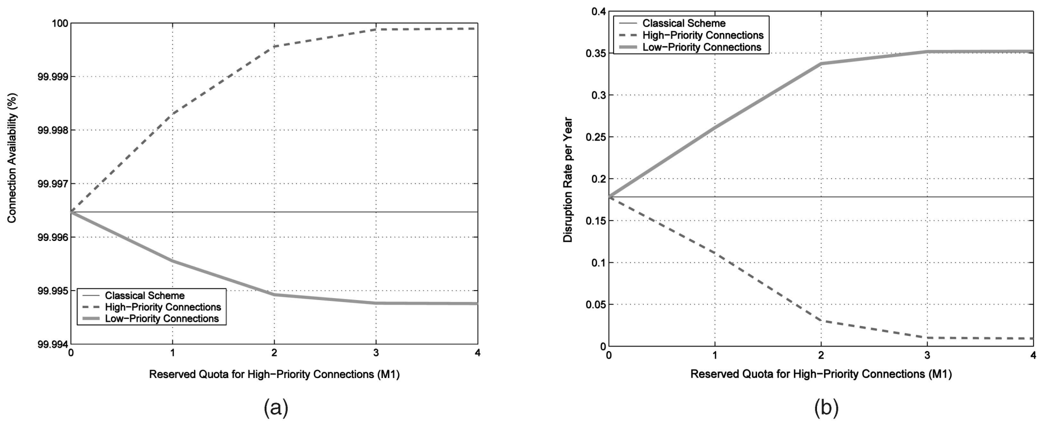

Fig. 4. Classical (without priority) and priority-enabled 4:12 shared-protection schemes, with four high-priority connections and eight low-priority ones. (a) Availability. (b) Disruption rate (per year).

It is easy to check that

$$\Pr\{A|X=(n_1,n_2,m)\}=1-\frac{n_2}{N_2}$$

and

$$\Pr\{C|A,X=(n_1,n_2,m)\}=\begin{cases}1 & \text{if } n_1+n_2\geq m\\ 0 & \text{otherwise.}\end{cases}$$

In the same way, using again the fact that $B=B_1\cap B_2$, we have

$$\begin{aligned}&\Pr\{B,C,X_1=n_1\}\\&=\sum_{n_2=1}^{N_2}\sum_{m=1}^{M}\Pr\{C|B,X=(n_1,n_2,m)\}\\&\times\Pr\{B_2|B_1,X=(n_1,n_2,m)\}\Pr\{B_1|X=(n_1,n_2,m)\}\\&\times p(n_1,n_2,m).\end{aligned}$$

We get

$$\Pr\{B_1|X=(n_1,n_2,m)\}=\frac{n_2}{N_2}$$

and $\Pr\{C|B,X=(n_1,n_2,m)\}$ is given by

$$\Pr\{C|B,X=(n_1,n_2,m)\}=\begin{cases}1 & \text{if } n_1+n_2\geq m\\ 0 & \text{otherwise.}\end{cases}$$

Finally, we get

$$\Pr\{B_2|B_1,X=(n_1,n_2,m)\}$$
$$=\begin{cases}1 \text{ if } n_1+n_2\leq m\\ \frac{m-n_1}{n_2} \text{ if } n_1\leq M_1 \text{ and } n_1<m \text{ and } n_1+n_2\geq m\\ \frac{m-M_1}{n_1+n_2-M_1} \text{ if } (M_1<m\leq n_1)\\ \qquad\text{or } (M_1<n_1<m \text{ and } n_1+n_2>m)\\ 0 \text{ otherwise}\end{cases}$$

and

$$\Pr\{B,C\}=\sum_{n_1=0}^{N_1}\Pr\{B,C,X_1=n_1\}.$$

The expression of the disruption rate S_2 is thus

$$\begin{aligned}S_2=\;&\frac{\lambda}{N_2}\left[\sum_{m=0}^{M}\sum_{n_1=0}^{N_1}\sum_{n_2=(m-n_1)^+}^{N_2}(N_2-n_2)p(n_1,n_2,m)\right.\\&\left.+\sum_{n_1=0}^{M_1-1}\sum_{m=n_1+1}^{M}\sum_{n_2=m-n_1}^{N_2}(m-n_1)(N_1-n_1)p(n_1,n_2,m)\right]\\&+\frac{\lambda_b}{N_2}\left[\sum_{n_1=0}^{M_1}\sum_{m=n_1+1}^{M}\sum_{n_2=m-n_1}^{N_2}(m-n_1)(n_1+1)p(n_1,n_2,m)\right.\\&+\sum_{n_1=M_1+1}^{N_1}\sum_{m=M_1+1}^{n_1\wedge M}\sum_{n_2=0}^{N_2}\frac{n_2(m-M_1)}{n_1+n_2-M_1}p(n_1,n_2,m)\\&\left.+\sum_{n_1=M_1+1}^{M-1}\sum_{m=n_1+1}^{M}\sum_{n_2=m-n_1}^{N_2}\frac{n_2(m-M_1)}{n_1+n_2-M_1}p(n_1,n_2,m)\right],\end{aligned}$$

where $(m-n_1)^+=\max(m-n_1,0)$.

5.2 Numerical Results

In this section, we evaluate the benefits introduced by our relative priority mechanism. Moreover, we study the impact of the quota M_1 on the availability and disruption rate of each class of service. To achieve this, we consider a scenario consisting of $N=12$ primary connections sharing $M=4$ BPs. We keep the same assumptions used in Section 3.3. That is, we consider a reference path cut rate $\lambda=\lambda_b=1/250\ h^{-1}$. In addition, we assume that the MTTR $(1/\mu)$ of all the paths is the same and equal to 12 hours. We first consider our priority-enabled scheme with $N_1=4$ high-priority and $N_2=8$ low-priority connections. The availability of each class is calculated for different values of M_1 based on the derived expressions of U_1 and U_2, and it is reported in Fig. 4a. Recall that when $M_1=0$, we deal with the classical shared-protection scheme without priority, and when $M_1=M$, we get the classical strict preemptive priority scheme.

Fig. 4a shows that the availability of high-priority connections increases with M_1, at the expense of a reduced

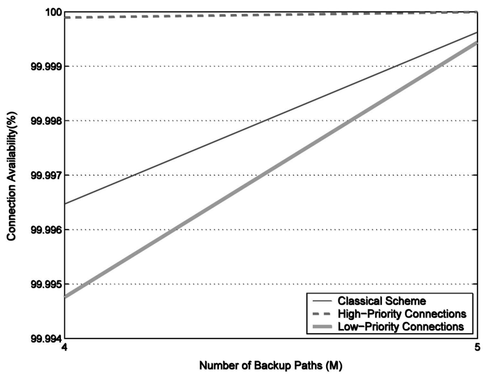

Fig. 5. Classical (without priority) and strict preemptive priority-enabled 4:12 shared-protection schemes, with four high-priority connections and eight low-priority ones.

availability of low-priority connections. Tuning this quota, we penalize more or less the low-priority connections. Indeed, the choice of M_1 depends mainly on the availability required by high-priority connections. For instance, according to [7], a gold client requests an availability of 99.999 percent. In this case, a value of $M_1 = 2$ is sufficient to obtain such availability and to respect at the same time the availability required by silver clients (that is, 99.99 percent according to [7]).

Fig. 4b depicts the service disruption rates for each class of service as a function of the threshold M_1. Unlike low-priority connections, the disruption rate of high-priority connections decreases with M_1. Again, this parameter can be fixed according to the requirements of each class of service. These results emphasize the flexibility and the gain introduced by our relative priority-enabled approach.

Indeed, our priority scheme is able to satisfy stringent high-priority connection requirements, whereas this objective cannot be achieved by the classical protection scheme without priority. In addition, thanks to the flexibility introduced by the parameter M_1, the satisfaction of strict high-priority connection requirements can be achieved while respecting the low-priority connection requirements. This feature is the key differentiator with respect to the classical strict preemptive strategy.

The observed availability results in Fig. 4a can be also interpreted from a practical-application perspective by using the following reasoning. Assume that the service provider offers two grades of service: gold and silver. According to the negotiated contracts, a gold client needs an availability of 99.999 percent, whereas a silver client requires an availability of 99.995 percent.

Fig. 4a shows that the classical protection scheme cannot satisfy the gold client requirements when only $M = 4$ BPs are provisioned. In this case, all the clients have the same availability, which is around 99.996 percent. This level of availability satisfies only the silver clients. To comply with the requirements of the gold clients, we need to provision $M = 5$ BPs instead of 4, as shown in Fig. 5. In doing so, both gold and silver clients are satisfied, thus benefiting from the same availability, which is slightly above 99.999 percent. However, this level is too high compared to the silver client requirements. As a direct result, more capacity may have been provided than needed, which is a cost factor for the service vendor.

One way to alleviate the aforementioned resource wastage problem is to provide multiple grades of service. To achieve this, we can use either our relative priority policy or the strict preemptive priority policy. Fig. 4a shows that the latter policy (that is, the case where $M_1 = M = 4$) differentiates between the two types of clients, but again, it does not meet with all the client requirements. On the contrary, with the classical scheme without priority, which satisfies only the silver clients when $M = 4$, the strict preemptive policy satisfies only the gold clients. Indeed, this policy is extremely severe with the lowest priority clients, resulting in the availability below the negotiated threshold 99.995 percent. To comply with both gold and silver clients under the strict preemptive priority, we need to provision $M = 5$ BPs (see Fig. 5), whereas this same target is achieved with only $M = 4$ BPs when using our relative priority policy, as shown in Fig. 4a. In this regard, our policy enables us to save backup resources with respect to both the classical protection scheme without priority and the strict preemptive policy.

In what follows, we investigate a second scenario, where the number of BPs is again set to $M = 4$. Half of this capacity is reserved for the $N_1 = 4$ high-priority connections (that is, $M_1 = 2$). Then, we vary the number of low-priority connections N_2 from 1 to 8. In other words, the total number of sharing connections N varies from 5 to 12.

Fig. 6a shows that increasing the number of low-priority connections does not affect the high-priority connections,

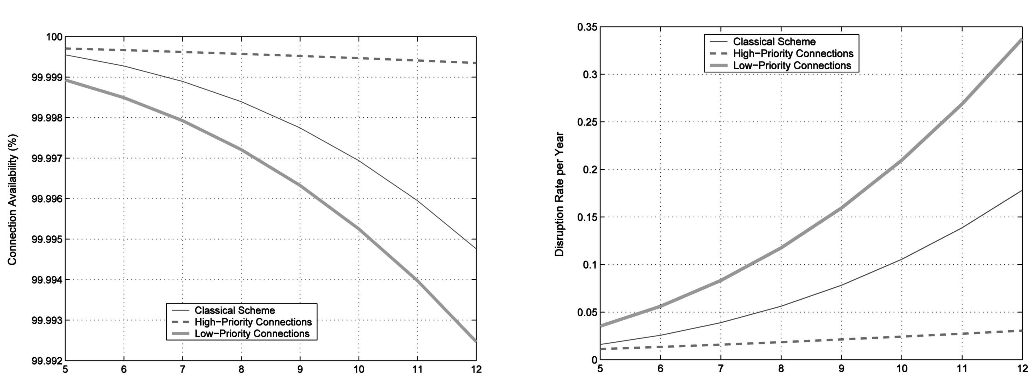

(a) (b)

Fig. 6. Classical (without priority) and priority-enabled $4 : N$ shared-protection schemes, with four high-priority connections, $N - 4$ Low-Priority ones, and $M_1 = 2$. (a) Availability. (b) Disruption rate (per year).

which maintain almost the same availability level. On the other hand, when priority is disabled, all the connections are penalized, as they become less available. As such, the availability required by high-priority connections (that is, 99.999 percent) is not respected once the total number of sharing connections exceeds 6, despite the use of multiple BPs. By using our priority mechanism, this issue is relieved. Indeed, high-priority connections always achieve the strict availability requirements of the gold clients, regardless of the increase of the number of low-priority connections. Likewise, Fig. 6b shows that by applying our priority scheme, the service disruption rate for the high-priority connections remains practically unchanged when the number of sharing primary connections increases.

6 Conclusion

In this paper, we presented a detailed mathematical model for the relative priority-enabled shared protection. We derived explicit analytic expressions for the average availability and service disruption rate for each class of service. Finally, we motivated the use of the proposed scheme, since it allows achieving stringent reliability requirements of gold clients without really penalizing the low-priority ones, whereas this target is hardly accomplished with the classical schemes.

References

[1] J. Zhang, K. Zhu, B. Mukherjee, and H. Zan, "A New Provisioning Framework to Provide Availability-Guaranteed Service in WDM Mesh Networks," *Proc. IEEE Int'l Conf. Comm. (ICC '03),* May 2003.

[2] A. Fumagalli, M. Tacca, F. Unghvary, and A. Farago, "Shared Path Protection with Differentiated Reliability," *Proc. IEEE Int'l Conf. Comm. (ICC '02),* Apr. 2002.

[3] M. Tacca, M. Tacca, A. Fumagalli, A. Paradisi, F. Unghvary, K. Gadhiraju, S. Lakshmanan, S.M. Rossi, A. de Campos Sachs, and D.S. Shah, "Differentiated Reliability in Optical Networks: Theoretical and Practical Results," *OSA/IEEE J. Lightwave Technology,* special issue on optical networks, vol. 21, no. 11, pp. 2576-2586, Nov. 2003.

[4] C.V. Saradhi, M. Gurusamy, and L. Zhou, "Differentiated QoS for Survivable WDM Optical Networks," *IEEE Optical Comm.,* vol. 42, no. 5, pp. S8-14, May 2004.

[5] S. Ramamurthy and B. Mukherjee, "A Review of Fault Management in WDM Mesh Networks: Basic Concepts and Research Challenges," *IEEE Networks,* vol. 18, no. 2, pp. 41-48, Mar. 2004.

[6] W. Fawaz, F. Martignon, K. Chen, and G. Pujolle, "A Novel Protection Scheme for Quality of Service Aware WDM Networks," *Proc. IEEE Int'l Conf. Comm. (ICC '05),* pp. 122-127, May 2005.

[7] W. Fawaz, B. Daheb, O. Audouin, B. Berde, M. Vigoureux, M. Dupond, and G. Pujolle, "Service Level Agreement and Provisioning in Optical Networks," *IEEE Comm. Magazine,* vol. 42, no. 1, pp. 36-43, Jan. 2004.

[8] F. Kamoun and L. Kleinrock, "Analysis of Shared Finite Storage in a Computer Network Node Environment under General Traffic Conditions," *IEEE Trans. Comm.,* vol. 28, no. 7, pp. 992-1003, July 1980.

[9] G.L. Choudhury, K.K. Leung, and W. Whitt, "Efficiently Providing Multiple Grades of Service with Protection against Overloads in Shared Resources," *AT&T Technical J.,* vol. 74, pp. 50-63, Aug. 1995.

[10] J.E. Angus, "On Computing MTBF for a K-Out-of-N:G Repairable System," *IEEE Trans. Reliability,* vol. 37, no. 3, pp. 312-313, Aug. 1988.

[11] D. Schupke and R. Prinz, "Performance of Path Protection and Rerouting for WDM Networks Subject to Dual Failures," *Proc. IEEE Optical Fiber Comm. Conf. (OFC '03),* Mar. 2003.

[12] S. Kim and S. Lumetta, "Evaluation of Protection Reconfiguration for Multiple Failures in WDM Mesh Networks," *Proc. IEEE Optical Fiber Comm. Conf. (OFC '03),* Mar. 2003.

Nizar Bouabdallah received the BS degree in telecommunications engineering from the Ecole Superieur des Communications, Tunis, Tunisia, in 2001, and the MS and PhD degrees in network and computer science from the University of Paris VI, Paris, in 2002 and 2004, respectively. He joined Alcatel Research Laboratories, Marcoussis, France, in 2002, while working on the PhD degree. In 2005, he was with North Carolina State University, Raleigh, as a postdoctoral fellow. He is currently a researcher at INRIA (Institut National de Recherche en Informatique et en Automatique). Since February 2007, he has been a visiting researcher in the School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada. His research interests include optical networking, wireless and sensor networks, performance evaluation, network planning and modeling, and control and management architectures.

Bruno Sericola received the PhD degree in computer science from the University of Rennes I in 1988. He has been with INRIA (Institut National de Recherche en Informatique et Automatique, a public research French laboratory) since 1989. His main research activity is in computer and communication systems performance evaluation, dependability and performability of fault-tolerant architectures, and applied stochastic processes.