

A Simple Priority Mechanism for 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 protection paths. In this letter, we propose 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.

I. INTRODUCTION

DUE to the perpetual growth in data traffic, the operators are progressively migrating toward optical core networks taking advantage of the tremendous transmission capacity offered by the optical technology. However, in such 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]. However, this solution is unsuitable from service perspective. One possible solution to provide different levels of reliability degrees is to use a priority mechanism for the shared-protection scheme. Recently, some service differentiation schemes have been proposed in literature [2]. 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: the service availability and service disruption rate.

These parameters were assessed analytically in [3] for the single backup path shared-protection scheme (i.e. 1 : N case). To achieve this, authors adopted simplistic assumptions, such as, the primary and backup paths are considered equally available (i.e. having the same failure and repair rates). This assumption is not realistic. Specifically, the backup path is usually chosen to be the longest link-disjoint path among all the precomputed ones. Doing so, the backup path 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, authors in [3] limited their study to only two classes of service, while the majority of the standards deal with up to 8 classes. As a main contribution, we therefore

Manuscript received September 20, 2006. The associate editor coordinating the review of this letter and approving it for publication was Prof. Maode Ma.

The authors are with INRIA (Institut National de Recherche en Informatique et Automatique), Rennes, France (email: {nizar.bouabdallah, bruno.sericola}@inria.fr).

Digital Object Identifier 10.1109/LCOMM.2007.061525.




Fig. 1. N primary paths sharing a backup path (1 : N system)

derive the reliability parameters associated to the 1 : N protection scheme for multiple classes of service with different path availabilities. The advantage of our model is to provide simple analytic expressions.

II. PRIORITY FOR SHARED-PROTECTION SCHEME

Assume N primary paths ($P_i, i = 1, \dots, N$) sharing the same backup path (BP) (i.e. 1 : N shared-protection system), as depicted in Fig. 1. In the classical shared-protection scheme, the first failed connection is recovered by the backup path until its primary path repair, 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, \dots, C_K , with N_i connections belonging to class C_i for $i = 1$ to K , and $N_1 + \dots + N_K = N$. Moreover, class C_i has a higher priority than class C_j as long as $i < j$.

Doing so, once the primary path of a connection t belonging to class C_i breaks down, the backup path, if it is available, is assigned to protect connection t and restoration is ensured by switching t to the backup path. 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 t fails the backup path is unavailable (i.e. already protecting another connection with the same or a higher priority than t ; or under failure), connection t becomes unavailable. However, if the backup path is protecting a connection belonging to a class of lower priority than t , it is immediately preempted by connection t . The preempted connection thus becomes unavailable, waiting for the backup path to be released or for its primary path to be repaired.

III. PERFORMANCE EVALUATION

In this Section, we present a mathematical model for the 1 : N priority-enabled shared-protection scheme. We derive explicit expressions for the average availability and service disruption rate of each class of connections. The availability of a connection is defined as the proportion of time the connection is up during its entire service; and the disruption rate of a connection is defined as the rate at which an available connection becomes unavailable [3].

A. 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 non operating (down).
- The backup path is said unavailable to restore connection t when it is down or already occupied recovering another connection with the same or 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 backup path are unavailable.
- For each path, the successive times to failure and repair times form an alternating renewal process.
- There are enough resources to repair simultaneously any number of failed connections. This is known in the literature as unlimited repair [4].

B. 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 $p_i = \mu_i/(\lambda_i + \mu_i)$. The unavailability q_i of a class C_i primary path is given by $q_i = 1 - p_i$.

In the same way, the backup path is characterized by its own failure and repair rates λ_b and μ_b . Its availability is $p_b = \mu_b/(\lambda_b + \mu_b)$ 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 . To simplify the writing, we shall denote the backup path by BP.

Theorem 1: For every $i = 1, \dots, K$, 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 sake of simplicity, we denote by $\{pp(t) = 0\}$ the event $\{t \text{ is down}\}$ 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. \blacksquare

Since all the C_i -connections behave identically, we get:

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

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

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

Hence, U_i can be written as:

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

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

$$U = q - \frac{p_b (1 - p^N)}{N} \quad (2)$$

where $q = 1 - p$.

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 backup path.

In the first case, connection t undergoes unavailable state only if the primary path of t breaks down and the backup path is already unavailable to restore C_i -connections. In this case, t transits to an unavailable state at a rate equal to the failure rate of its primary path, λ_i . In the second case, such transition happens when the backup path fails, at a rate equal to λ_b , or when one of the higher-priority primary connections breaks down and thus preempts the connection t , at a rate equal to $N_1 \lambda_1 + \dots + N_{i-1} \lambda_{i-1}$. Hereafter, we write $\{pp(t) = 1\}$ to describe the event $\{t \text{ is up}\}$. The expression of the disruption rate S_i is thus given by

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

We already got the probability that connection t is restored by the backup path, i.e.,

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

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

$$\begin{aligned} S_i &= \frac{1}{N_i} p_b \left(\lambda_b + \sum_{j=1}^{i-1} N_j \lambda_j\right) \left(1 - p_i^{N_i}\right) \prod_{j=1}^{i-1} p_j^{N_j} \\ &\quad + \lambda_i \left(p_i - p_b \prod_{j=1}^i p_j^{N_j}\right) \end{aligned} \quad (3)$$

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

$$S = \frac{1}{N} \lambda_b p_b (1 - p^N) + \lambda \cdot p (1 - p_b p^{N-1}) \quad (4)$$

It is easy to show that S increases when p_b decreases and that we have $\lim_{N \rightarrow \infty} S = \lambda \cdot p$. This behavior is illustrated in Fig. 2. In fact, when N is large, the impact of the backup path




Fig. 2. Disruption rate per year for the classical 1 : N shared-protection scheme

is negligible. One interesting finding, is that the impact of the backup path availability p_b on the disruption rate. Specifically, the usage of a backup path to protect a set of connections does not necessarily mean a better disruption rate performance. We can observe through Fig. 2 that, unlike the availability, which is always improved thanks to the backup protection (refer to (2)), the disruption rate may become worse due to the backup path introduction (See Fig. 2 case $\lambda_b = 1/10 h^{-1}$).

This example exhibits the critical choice of the backup path. It shows the importance of using the disruption rate as a QoS metric. In fact, two connections may have the same availability during their entire service periods; however, one of them may experience fewer network failures with longer service downtime for each failure, while the other may experience more network failures with shorter service downtime. Although the two connections have the same availability, they experience different service disruption rates, which may lead to different customer-perceived service qualities.

C. 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 backup path. We first consider the priority-enabled protection scheme, with three classes of service. Each of the two highest classes of service (i.e., C_1 and C_2) encloses 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 is applied to this scenario. It is important to note that the MTTR ($1/\mu$) of all the paths is considered equal to 12 hours [5]. We also consider a reference path cut-rate $\lambda = \lambda_b = 2 \cdot 10^{-4} h^{-1}$.

Figure 3 shows 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, availability required by high-priority connections is not respected although the use of a backup path. Specifically, according to [5], a gold client requests an availability of 99.999%. 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.




Fig. 3. Availability for the classical (without priority) and priority-enabled 1 : N shared-protection schemes




Fig. 4. Disruption rate (per year) for the classical (without priority) and priority-enabled 1 : N shared-protection schemes

Likewise, Fig. 4 shows that, applying our priority scheme, the service disruption rate for the two highest-priority connections remains unchanged when the number of remaining lower-priority connections increases. Indeed, our priority scheme enables to satisfy stringent high-priority connection requirements while this objective can not be achieved by the classical protection scheme.

IV. CONCLUSION

In this letter, we presented a detailed mathematical model for the 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 while this target is hardly accomplished with the classical scheme.

REFERENCES

- [1] M. Tacca *et al.*, "Differentiated reliability in optical networks: theoretical and practical results," *OSA/IEEE J. LightWave Technol.*, special issue on optical networks, vol. 21, no. 11, pp. 2576–2586, Nov. 2003.
- [2] C. V. Saradhi, M. Gurusamy, and L. Zhou, "Differentiated QoS for survivable WDM optical networks," *IEEE Optical Commun.*, May 2004.
- [3] K. C. W. Fawaz, F. Martignon, and G. Pujolle, "Analysis of protection techniques for quality of service aware WDM networks," in *Proc. IEEE International Conference on Communications (ICC'05)*.
- [4] 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, Mar. 1998.
- [5] W. Fawaz, B. Daheb, O. Audouin, B. Berde, M. Vigoureux, M. Dupond, and G. Pujolle, "Service level agreement and provisioning in optical networks," *IEEE Commun. Mag.*, vol. 42, no. 1, pp. 36–43, Jan. 2004.