Computer Networks 56 (2012) 3150–3162

Contents lists available at SciVerse ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Trading availability among shared-protected dynamic connections in WDM networks

Diego Lucerna [a], Massimo Tornatore [b,*], Biswanath Mukherjee [c], Achille Pattavina [b]

[a] Huawei Technologies Italia S.r.l., Via Lorenteggio 257, 20152 Milano (MI), Italy
[b] Department of Electronics and Information, Politecnico di Milano, Via Ponzio 34-35, 20121 Milan, Italy
[c] Department of Computer Science, University of California, Davis, CA 95616, USA

ARTICLE INFO

Article history:
Received 14 October 2011
Received in revised form 17 April 2012
Accepted 18 April 2012
Available online 3 May 2012

Keywords:
Optical network
WDM
Elastic availability
Holding time
Shared protection
SLA Violation Risk

ABSTRACT

Novel automatized management systems for optical WDM networks promise to allow customers asking for a connection (i.e., a bandwidth service) to specify on-demand the terms of the Service Level Agreement (SLA) to be guaranteed by the Network Operator (NO). In this work, we exploit the knowledge, among the other Service Level Specifications (SLS), of the *holding time* and of the *availability target* of the connections to operate shared-path protection in a more effective manner.

In the proposed approach, for each connection we monitor the actual downtime experienced by the connection, and, when the network state changes (typically, for a fault occurrence, or a connection departure or arrival), we estimate a new updated availability target for each connection based on our knowledge of all the predictable network-state changes, i.e., the future connection departures. Since some of the connections will be ahead of the stipulated availability target in their SLA (credit), while other connections will be behind their availability target (debit), we propose a mechanism that allows us to "trade" availability "credits" and "debits", by increasing or decreasing the shareability level of the backup capacity. Our approach permits to flexibly manage the availability provided to living connections during their holding times.

The quality of the provided service is evaluated in terms of availability as well as probability of violation of availability target stipulated in the SLA (also called *SLA Violation Risk*), a recently-proposed metric that has been demonstrated to guarantee higher customer satisfaction than the classical statistical availability. For a typical wavelength-convertible US nationwide network, our approach obtains significative savings on Blocking Probability (BP), while reducing the penalties due to SLA violations. We also analytically demonstrate the proposed scheme can be highly beneficial if the monitored metric is the SLA Violation Risk instead of the availability.

1. Introduction

In optical WDM networks survivability mechanisms are needed to avoid that a failure of a network element may cause significant losses of revenue for those customers that run their service over the bandwidth provided by means of the optical paths. These revenue losses are then reclaimed in the form of penalties to be paid by the Network Operator (NO). Different protection mechanisms to ensure survivability in WDM networks have been proposed [1]: among them, Shared-Path Protection (SPP) is one of the most-adopted options, because of its desirable resource efficiency [2].

Recently, many new applications are emerging with requirements of large bandwidth over relatively short

* Corresponding author.
E-mail addresses: diego.lucerna@gmail.com (D. Lucerna), tornator@elet.polimi.it (M. Tornatore), mukherje@cs.ucdavis.edu (B. Mukherjee), pattavina@elet.polimi.it (A. Pattavina).

1389-1286/$ - see front matter
http://dx.doi.org/10.1016/j.comnet.2012.04.021

and predictable periods of time: let us consider, e.g., video distribution of important sport or social events, or the massive data transfer for backup, storage or e-science purposes. Network technology and the bandwidth market are developing to provide the flexible platform the new applications are asking for. In particular, new architectures and routines for user-controlled on-demand optical circuit provisioning [3], typically based on automatic or web-based interfaces at the management plane (MP) [1], will enable the on-line specification of the Service Level Agreement (SLA) terms to be guaranteed (with different price range) by the Network Operator (NO). In other words, users may be able to specify the QoS terms [3] of their connection requests, e.g., the availability target (A_T) or the holding-time.

In particular, the NO should be able to guarantee with a high probability that the stipulated A_T is respected, in order to avoid penalties; at the same time, the NO aims at increasing its profit, i.e., it wants to maximize the number of connections (or bandwidth) provisioned. But, in case of SPP and dynamic traffic, the NO must carefully monitor the actual availability provided to the customer. In fact, whenever a new SPP connection is routed, the NO must not only verify if the A_T of the incoming connection is satisfied, but also it must check if the A_T's of the existing connections are still respected despite the increased sharing of backup resources. A typical solution to avoid penalties is to employ an *availability-guaranteed provisioning* approach [4,2]: in this case, the NO provisions a connection only if the network can provide a path with a long-term theoretical availability (that can be a priori evaluated) that is equal or larger than the A_T target.

In this work, we present a novel availability-guaranteed provisioning method that *dynamically manages* the availability provided to SPP-protected connections during their holding time. Our approach (i) leverages on the information about connection departures (future departures imply decreased sharing and in turn more availability) and (ii) is able to dynamically "trade" availability from connections ahead of their A_T to connections behind their A_T. Note that the proposed scheme *do not reprovision backup resources* (and, consequently, it avoids the additional control overhead required by backup reprovisioning), but it only operates on the sharing of the backup resources.

When a new connection has to be routed, the target availability of each living connection in the network is re-estimated by considering the experienced downtime and the remaining holding time: e.g., if a connection has not been affected by failures, the NO can be considered in "credit" of availability with respect to the customer and the customer's availability requirements can be opportunely decreased by the NO, as long as the original target A_T is still respected. On the other hand, if a connection has undergone an outage period and it is getting close to its maximal acceptable downtime, then it can be considered in "debit" of availability with respect to the NO, and its current availability requirement should be increased by the NO in the attempt to match the original stipulated A_T. In brief, the proposed approach dynamically "transfers" availability from connections that have an availability debit to connections in credit of availability by allowing more or less sharing of backup resources and it helps the NO to meet connections SLAs. Although the proposed methodology can be applied, to any SPP algorithm, in the following we show the effectiveness of our approach through simulative experiments using the Availability-Guaranteed Provisioning algorithm (AGP) presented in [2]. Since our proposed method requires the knowledge of the connection holding-time we will refer to it as the Holding-Time-Aware (HTA) method.

Furthermore, recent studies [5–7] have shown that using the theoretical long-term availability to evaluate the quality of a connection provided over a short period of time (e.g., a period comparable with the average failure interval) may not be enough to characterize the quality of provisioning in terms of SLA satisfaction (more details will be given in Section 2). So, in this paper, we also analytically discuss how to extend the proposed approach to the case where the monitored metric is the SLA Violation Risk (i.e., the probability that, given a certain availability target and certain statistical availability associated to the path, the offered connection does satisfy the availability target) when our trading-availability method is adopted. Some preliminary results are also provided.

The rest of the paper is organized as follow. Section 2 overviews some background work on availability-guaranteed SPP strategies. In Section 3 we formally state the availability-guaranteed SPP problem and we briefly describe an existing solution for availability-guaranteed SPP, called AGP. In Section 4 our new HTA method to dynamically trade availability among SPP connections is presented. Section 5 discusses how to extend the HTA approach to include the new SLA Violation Risk metric. In Section 6 we compare and evaluate by means of simulations our HTA methodology vs. the basic AGP approach. In Section 7, we draw the conclusion. In the Appendix A, a rigorous approach to evaluate the availability in a SPP network scenario is presented.

2. Prior work

This paper provides novel contributions on two complementary, bus distinct, lines of research in the field of shared-path protection: (1) how to route of availability-guaranteed shared-path-protected connections and (2) how to evaluate the SLA Violation Risk, or interval availability, for availability-guaranteed SPP routing.

2.1. Availability-guaranteed SPP

We start by considering the problem of dynamic routing with Shared-Path Protection (SPP). To enable dynamic provisioning of SPP connections, a network-control element (say, e.g., the Path Computational Element, PCE) needs to compute two link-disjoint paths, a dedicated working path and a shared backup path, for each incoming connection request, based on the current network state. SPP routing algorithms are usually based on two-step approaches (e.g., [8]), which compute separately the working and the backup path, using shortest-path or K-shortest-path algorithms [9] that minimize the total link costs. Generally, link costs