Computer Networks 82 (2015) 4–19

Contents lists available at ScienceDirect

Computer Networks

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

The *Spine* concept for improving network availability

Abdulaziz Alashaikh [a], Teresa Gomes [b], David Tipper [a,*]

[a] Graduate Telecommunications and Networking Program, University of Pittsburgh, Pittsburgh, PA 15260, United States
[b] Department of Electrical and Computer Engineering, University of Coimbra/INESC Coimbra, 3030-290 Coimbra, Portugal

ARTICLE INFO

Article history:
Received 16 June 2014
Received in revised form 4 December 2014
Accepted 4 February 2015
Available online 5 March 2015

Keywords:
Availability
Quality of resilience classes
Network design

ABSTRACT

Telecommunications networks need to guarantee that all node pairs involved in critical service communications are highly available. Here we adopt a novel approach to the problem of how to provide high levels of availability in an efficient manner. The basic idea is to embed at the physical layer a high availability set of links and nodes (termed the *spine*) in the network topology to support protection and routing in providing end-to-end availability. We first explore the spine concept through simple topologies illustrating the potential benefits of the approach in improving the overall network availability and the capability to support quality of resilience classes. Then, we study how the structural properties of a network topology can be used to determine heuristics to select a suitable spine and compare this with the case where all network components have the same availability. This is followed by a numerical based study comparing the heuristics with all possible spanning tree based spines for sample topologies. Our results demonstrate how to best design a physical network to support protection methods in achieving high levels of availability efficiently.

1. Introduction

Communication networks are one of the critical national infrastructures upon which society depends [1], thus it is imperative that they are highly available and resilient to failure. In particular networks need to guarantee that all node pairs involved in critical service communications (e.g., financial transactions, emergency calls, smart grid communications, etc.) have a high end-to-end availability. The traditional approach to improving availability in systems is to add parallel redundancy [2], which in the context of typical optical backbone networks would imply adding additional links and possibly nodes to the network topology to support additional parallel routes. However, from a service provider's perspective communication networking is increasingly becoming a commodity type of business with severe cost constraints limiting improvements to network availability. Thus adding links to nationwide or continent wide backbone networks simply to improve availability is difficult to economically justify. Furthermore, only a small number of users and services need very high levels of availability and these users/services produce only a small fraction of the total network traffic. Unfortunately, the small amount of high availability traffic drives the network design giving rise to a free rider scenario where the majority of customers get a higher level of availability than they need or are willing to pay for. Hence, from a service provider's perspective, there is a need to support classes of quality of resilience (QoR) in a fashion similar to quality of service classes. The basic concept is to categorize traffic into classes and provide different levels of availability and fault protection for each class. The goal of providing QoR classes is to just meet availability requirements without over-engineering the network for the lowest classes of traffic. Providing quality of resilience classes has been mentioned in the current literature in

* Corresponding author.
E-mail addresses: azizoozi@gmail.com (A. Alashaikh), teresa@deec.uc.pt (T. Gomes), tipper@tele.pitt.edu (D. Tipper).

http://dx.doi.org/10.1016/j.comnet.2015.02.020
1389-1286/

both a qualitative and quantitative fashion (see the survey paper [3]). Most of the current work focuses on providing QoR classes by using different restoration mechanisms per traffic type in a particular network layer (e.g., WDM, MPLS, etc.). For example, providing gold, silver and bronze service classes by giving the gold traffic dedicated reserved bandwidth backup paths, silver class shared backup path restoration and the bronze class no protection relying on rerouting after failure. Other approaches in the literature address the problem from upper layers by having an overlay network maximize path diversity and dual homing or preconfiguring logical rings or protection cycles (p-cycles) between nodes for gold class traffic. These approaches suffer from the crosslayer mapping issues discussed in the literature as without full knowledge of the physical layer and the mappings between layers no hard guarantees on availability can be provided (i.e., due to fault propagation) [4]. Essentially, the current approach is to take the physical network availability as a given and deploy redundancy and restoration techniques at various layers to provide QoR classes with different fault recovery capabilities and availabilities.

We believe that high availability must begin at the physical layer and work it's way up the various layers. Note that typical service provider optical backbone networks are at least two-way connected supporting some number of disjoint/partially disjoint paths between node pairs at the physical layer. Here we assume the network topology is fixed and the cost of adding links/nodes is prohibitive. We propose an innovative technique of embedding a higher availability sub-structure into the network at the physical layer to improve the overall network availability without substantial modifications to the topology. We term the high availability sub-structure portion of the network the *spine* [5]. The spine would connect those nodes with traffic needing a high level of availability and provide a basis for differentiated classes of resilience. For example, the highest quality of resilience class traffic would be routed on the spine or use the spine as a backup path. The nodes, link interfaces and links on the network spine would have higher availability than the equipment that is not part of the spine. This provides *levels of availability differentiation at the physical level* which can be leveraged with restoration techniques, logical virtual network topology routing, cross layer mapping and other methods to further differentiate resilience classes and provide an extended range of availability guarantees. One can think of the spine approach as assuming a restoration method (path restoration) or set of restoration methods (i.e., no protection, shared backup path, dedicated backup path, etc.) is to be used, then determining how should availabilities be assigned to the physical network components to best support the availability requirements.

The higher availability of the spine, in comparison to the non-spine part of the network, can be accomplished using a variety of techniques. For example, on the spine more expensive equipment can be utilized that is arranged and configured to provide high availability (e.g. hot standby line card, redundant fans, etc.) with redundant equipment deployed locally in parallel as needed (e.g., hot standby fiber in physically diverse duct, etc.). Also, the equipment along the spine can be situated to increase the mean time to failure (MTTF) using a number of techniques such as longer back up power supplies, better heating/cooling, stronger outside cabinets, underground cabling instead of above ground, etc. In a similar vein, methods can be employed to reduce the mean time to repair (MTTR) along the spine. For instance, one can follow best practices and training procedures as determined by several government and trade organizations (e.g., NRIC, FCC, TIA) and standards bodies (e.g., ITU) [6]. The operator can pre-position spare parts, equipment, software and test equipment along the spine. Similarly, the network operations center (NOC) can more closely monitor the spine portion of the network. Additionally, the operator can assign the most experienced staff to the operations, administration and management (OAM) of the spine portion of the network. Many of the methods above are employed in other critical infrastructures (e.g., the power grid) and industries and studies show that the average MTTR can be reduced by 5–25% resulting in a significant improvement in the availability. Of course exactly which combination of techniques (hardware, equipment siting, workforce training, etc.) is adopted to improve the reliability of the spine will depend on the cost versus benefit structure of the network owner. Even using techniques to improve the MTTF and MTTR of links and nodes that comprise the spine, we assume additional protection, either end-to-end, segment or local [7] is needed to achieve the desired level of end-end availability for the most stringent QoR class.

In this paper, we explore the spine concept and its potential advantages. We will assume that high availability communication services between all node pairs in the network is required and thus the spine will be a spanning tree. We show that if intelligently deployed, the spine approach can be leveraged to provide higher overall average end-to-end availability or lower downtime per year efficiently. The remainder of the paper is organized as follows. Section 2 presents the spine concept through analysis of a simple network, which shows that the spine can be advantageous from the point of view of the average end-to-end availability and support a wider spread in the availabilities. In Section 3 we study how the structural properties of the network topology can be used to determine heuristics to select a suitable spine and compare this with the case where all network components have the same availability. This is followed by a numerical based study comparing the heuristics with all possible spanning tree based spines in Section 4. Section 5 studies the sensitivity of the heuristic spine selection methods, the effects of heterogenous availabilities and other practical issues including cost. Our conclusions and future work are given in Section 6.

2. The *Spine* concept

The spine concept is to take a physical network topology graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ which consists of a set of $\mathcal{N}$ nodes and a set $\mathcal{L}$ of links (undirected edges), then embed a sub-structure $\mathcal{G}_s = (\mathcal{N}_s, \mathcal{L}_s)$ with higher availability in a fashion