CCIT Report #463 January 2004

The Power of Tuning: a Novel Approach for the Efficient Design of Survivable Networks

Ron Banner and Ariel Orda
Department of Electrical Engineering
Technion – Israel Institute of Technology
Haifa 32000, Israel
{banner@tx, ariel@ee}.technion.ac.il

Abstract— **Current survivability schemes typically offer two degrees of protection, namely *full protection* (from a single failure) or *no protection* at all. Full protection translates into rigid design constraints, i.e. the employment of disjoint paths. We introduce the concept of *tunable survivability* that bridges the gap between full and no protection. First, we establish several fundamental properties of connections with tunable survivability. With that at hand, we devise efficient connection establishment schemes for both 1:1 and 1+1 protection architectures and formally establish their optimality. Then, we show that the concept of tunable survivability gives rise to a novel *hybrid protection* architecture, which offers improved performance over the standard 1:1 and 1+1 architectures. Next, we investigate some related QoS extensions. Finally, by way of simulations, we demonstrate the advantage of tunable survivability over full survivability. In particular, we show that, by just slightly alleviating the requirement of full survivability, we obtain major improvements in terms of the "feasibility" as well as the "quality" of the solution.**

***Keywords*— Survivable Connections, Path Protection, Routing, Theory of Algorithms, Combinatorial Optimization.**

I. INTRODUCTION

In recent years, transmission capabilities have increased to rates of 10 Gbit/s and beyond [9]. With this increase, any failure may lead to a vast amount of data loss. Consequently, several survivability strategies have been proposed and investigated. These strategies are based on securing an independent resource for each potentially faulty network element [6]. This requirement usually translates into the establishment of pairs of disjoint paths. Two major survivability architectures that employ the use of (link) disjoint paths are the 1+1 and 1:1 protection architectures. In the 1+1 protection architecture, the data is concurrently sent on a pair of disjoint paths. The receiver picks the better path and discards data from the other path. In the 1:1 protection architecture, data is sent only on one (active) path, while the other (backup) path is activated by signaling only upon a failure on the active path.

Under the common single link failure model, the employment of disjoint paths provides full (100%) protection against network failures. However, in practice, this requirement is often too restrictive. Indeed, in many cases this requirement is infeasible (when pairs of disjoint paths do not exist [12]) and in other cases it is very limiting and results in the selection of inefficient routing paths [9]. Therefore, it has been noted that a milder and more flexible survivability concept is called for, which would relax the rigid requirement of disjoint paths[9]. However, to the best of our knowledge, no previous work has systematically addressed this problem.

In this study, we introduce the concept of *tunable survivability*, which provides a *quantitative measure* to specify the desired level of survivability. This concept allows any degree of survivability in the range 0% to 100% and, in contrast to the rigid requirement of disjoint paths, it offers flexibility in the choice of the routing paths; consequently, it enables to consider valuable *tradeoffs* for designing survivable networks, such as survivability vs. feasibility, survivability vs. available bandwidth, survivability vs. delay performance, etc.

We adopt the widely used single link failure model, which has been the focus of most studies on survivability e.g., [5],[7],[10],[14]. Tunable survivability enables the establishment of connections that can survive a single failure with any desired probability p. Such connections are termed *p-survivable*. More specifically, a p-survivable connection is a set of paths between some source and destination nodes such that, *upon* a single network failure, the probability to have at least one operational path is at least p. The following example illustrates the power of p-survivable connections with respect to the traditional scheme of disjoint paths.

Example 1: Consider the network described in Fig.1. Assume that each link $e_i \in E$ is associated with a bandwidth value b_i. Let the failure probabilities (*given* the event of a network failure) be 0.05 for all links but e_6, and 0 for e_6.

As no pair of disjoint paths from S to T exists in the network, the traditional survivability requirement is infeasible. On the other hand,