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 polynomial (optimal) connection establishment schemes for both 1:1 and 1+1 protection architectures. 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, 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.

1. 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) 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., [4],[5],[7],[10],[11],[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[1]. 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. Let the failure probabilities be 0.001 for all links. The failure probability of each link upon an event of a failure is

---

[1] The probability is defined to be under the condition of a failure since survivability is the capability of the network to maintain service *upon* an event of a failure [8].

$p_i = \dfrac{0.001}{1 - (1 - 0.001)^{10}} = 0.1$. *As no pair of disjoint paths from S to*

T exists in the network, the traditional survivability requirement is infeasible. Suppose now that we are satisfied with connections that upon a failure remain operational with a probability of at least 0.9. In that case, it is easy to see that a connection that consists of the paths $\pi_1 = (S,a,b,f,T)$ *and* $\pi_2 = (S,e,d,f,T)$ *fits since the only (single) failure that can damage both paths is a failure in* $e_{10} = (f,T)$*; therefore, since the link* e_{10} *fails with a probability of 0.1 (upon a failure), the connection* (π_1, π_2) *is 0.9-survivable. Now suppose that we are satisfied with* $(0.9)^2$ *-survivable connections. In that case it is easy to see that for* $\pi_3 = (S,a,b,d,f,T)$ *and* $\pi_4 = (S,c,d,f,T)$ *the connections* (π_3, π_4)*,* (π_2, π_3) *and* (π_2, π_4) *can also be used; thus, substantially increasing the space of feasible solutions. Finally, assume that we are satisfied with* $(0.9)^4$ *-survivable connections. In that case it is easy to see that single paths like* π_1 *or* π_2 *turn also to be feasible solutions.*
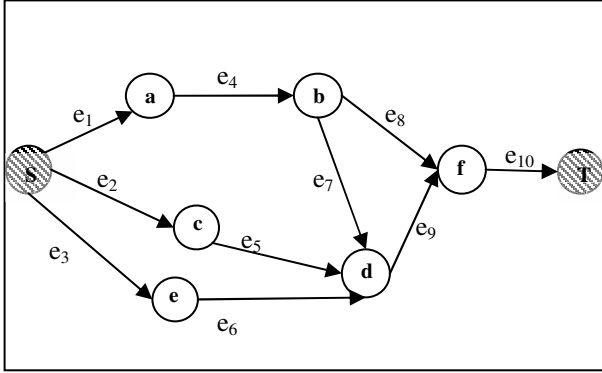

Figure 1. A reference network for the discussion of *p*- survivable connections

Through comprehensive simulations on random internet networks we demonstrate the major power of the tunable survivability concept. In essence, we show that, at the price of a *negligible* reduction in the level of survivability, we obtain a *major* increase in the bandwidth as well as the feasibility of the solutions.

Motivated by the above results, we investigate the tunable survivability concept from several different aspects and for different protection architectures. To that end, we first establish several fundamental properties of *p*-survivable connections. In particular, we prove that, if it is possible to establish a *p*-survivable connection with some supported bandwidth *B* through *more than two paths*, then it is also possible to establish such a connection (i.e., with the same probability *p* and bandwidth *B*) through *exactly two paths*.[1] Hence, in this study, we

---

[1] While this is a trivial property for *disjoint paths* under the single link failure model, it is far from trivial, and actually quite surprising, for paths that may be non-disjoint.

focus on survivable connections that consist of exactly two paths. Next, for both the 1+1 and the 1:1 protection architectures, we design efficient schemes for the establishment of *p*-survivable connections. Basically, for each protection architecture, we propose two types of survivability schemes: schemes that aim at *widest p-survivable connections* (i.e., *p*-survivable connections with maximum bandwidth) and schemes that aim at maximum survivability (i.e., connections with the maximum probability to survive single failures). We also show that each of the proposed schemes can be enhanced in order to consider QoS requirements. Finally, we show that all schemes achieve the *optimal* solution and are *computationally efficient*.

Next, we turn to show that the concept of tunable survivability gives rise to a third protection architecture, which is an hybrid between 1:1 protection and 1+1 protection. This new architecture is shown to have several important advantages over both the 1:1 and the 1+1 protection architectures; moreover, we show that the schemes that we have established for achieving either widest or most survivable connections in the case of 1:1 protection achieve the same goals in the case of hybrid protection.

The rest of this paper is organized as follows. In Section 2, we introduce some terminology and formally define the concept of tunable survivability. In Section 3, we investigate several properties of connections with tunable survivability. In Section 4, we design efficient schemes that establish *most survivable* and *widest p-survivable* connections for the 1:1 and 1+1 protection architectures. In section 5, we introduce the Hybrid Protection architecture, demonstrate its advantages and establish corresponding algorithmic schemes. In Section 6, we show how our schemes can be enhanced in order to consider QoS requirements. Section 7 presents simulation results that demonstrate the advantages of tunable survivability. Finally, Section 8 summarizes our results and discusses directions for future research.

2. Model and problem formulation

A *network* is represented by a directed graph $G(V,E)$, where *V* is the set of nodes and *E* is the set of links. Let $N = |V|$ and $M = |E|$. A *path* is a finite sequence of nodes $\pi = (v_0, v_1, \cdots v_h)$, such that, for $0 \le n \le h-1$, $(v_n, v_{n+1}) \in E$. A path is *simple* if all its nodes are distinct.

Given a source node $s \in V$ and a target (destination) node $t \in V$, the set $P^{(s,t)}$ is the collection of all directed paths from the source *s* to the target *t*.

Each link $e \in E$ is assigned a *weight* $w_e \in \mathbb{Z}^+$, a *bandwidth* $b_e \in \mathbb{Z}^+$ and an independent *failure probability* $p_e \in [0,1]$. We note that, since survivability schemes

consider recovery *upon* the event of a failure in the network [4], p_e is the probability that, *given* a (single) failure event in the network, the link e is the failed component. Under the single line failure model, it is straightforward to obtain the probabilities $\{p_e\}$ out of *a priori* link failure probabilities. The latter are often estimated out of available failure statistics of each network component [4].

We consider a *link state* routing environment, where each source node has a (precise) image of the entire network.

Definition 1: Given a (non-empty) path π, its *bandwidth* $B(\pi)$ is defined as the bandwidth of its bottleneck link, namely, $B(\pi) \triangleq \underset{e \in \pi}{Min}\{b_e\}$.

A link is classified as *faulty* upon its failure; it remains faulty until it is *repaired*. We say that a link $e \in E$ is *operational* if it is not faulty. Likewise, we say that a path π is *operational* if it has no faulty link i.e., for each $e \in \pi$, link e is operational.

Definition 2: Given a network $G(V,E)$ and a pair of source and destination nodes *s* and *t*, a *survivable connection* is a pair of paths $(\pi_1, \pi_2) \in P^{(s,t)} \times P^{(s,t)}$.[1]

We say that a connection (π_1, π_2) is *operational* if either π_1 or π_2 are operational. Moreover, as survivability is defined to be the capability of the network to maintain service continuity in the presence of failures [8], we quantify the quality of (tunable) survivable connections as their probability to remain operational in the presence of failures. This is formalized as follows.

Definition 3: Given are a network $G(V,E)$, a failure probability $p_e \geq 0$ for each link $e \in E$, and a survivable connection (π_1, π_2). We say that (π_1, π_2) is a *p-survivable connection* if, *upon a link failure*, it remains operational with a probability of at least *p*. The value of *p* is then termed as the *survivability level* of the connection.[2]

Definition 3 formalizes the notion of tunable survivability. Note that, under the single link failure model, any pair of disjoint paths is a 1-survivable connection.

[1] As was already mentioned, we will show that there is no advantage in the employment of more than two paths; hence, the definition focuses on two paths.

[2] Note that the *a-priory* probabilities that a *p*-survivable connection is operational is (considerably) larger than *p*. Specifically, it is equal to $p \cdot \left[1 - \prod_{e \in E}(1 - \widetilde{p_e})\right]$, where $\widetilde{p_e}$ is the *a-priory* probability that a link e fails.

We now quantify the bandwidth of a survivable connection. We consider first a connection (π_1, π_2) under the standard (full) survivability requirement. This means that π_1 and π_2 are disjoint, namely $\pi_1 \cap \pi_2 = \phi$. Obviously, for 1+1 protection, the maximum protected traffic rate that can be transferred via (π_1, π_2) is the minimum available bandwidth on any of the two paths. That is, the bandwidth of the connection (π_1, π_2) is $\min\{B(\pi_1), B(\pi_2)\} = \underset{e \in \pi_1 \cup \pi_2}{\min}\{b_e\}$. However, for connections with tunable survivability, paths are not necessarily disjoint. Therefore, for the 1+1 protection architecture, the total traffic rate that traverses links that belong to both π_1 and π_2 is twice the rate that traverses links that belong to only one out of the two paths. Accordingly, the available bandwidth of a survivable connection with respect to 1+1 protection is defined as follows.

Definition 4: Given a survivable connection (π_1, π_2), its *bandwidth with respect to the 1+1 protection architecture* is the maximum $B \geq 0$ such that $2 \cdot B \leq b_e$ for each $e \in \pi_1 \cap \pi_2$ and $B \leq b_e$ for each $e \in (\pi_1 \cup \pi_2) \setminus (\pi_1 \cap \pi_2)$.

In contrast to 1+1 protection, in 1:1 protection only one duplicate of the original traffic is carried at any given time. Therefore, the only restriction here is that traffic rate should not exceed the bandwidth of any of the links in $\pi_1 \cup \pi_2$. Accordingly, we formulate the bandwidth of a survivable connection with respect to the 1:1 protection architecture as follows.

Definition 5: Given a survivable connection (π_1, π_2), its *bandwidth with respect to the 1:1 protection architecture* is the maximum $B \geq 0$ such that $B \leq b_e$ for each $e \in \pi_1 \cup \pi_2$.

For a source-destination pair, there might be several *p*-survivable connections. Among them, we may be interested in those that have the best "quality". The following definitions correspond to two important quality criteria namely, maximum survivability and maximum bandwidth.

Given a network $G(V,E)$ and a pair of nodes *s* and *t*, we say that a *p*-survivable connection $(\pi_1, \pi_2) \in P^{(s,t)} \times P^{(s,t)}$ is a *most survivable connection* if there is no $\widehat{p}$-survivable connection $(\widehat{\pi_1}, \widehat{\pi_2}) \in P^{(s,t)} \times P^{(s,t)}$ such that $\widehat{p} > p$; *p* is then termed the *maximum level of survivability*. Next, we say that a *p*-survivable connection (π_1, π_2) is the *widest p-*

survivable connection for the 1+1 (alternatively 1:1) protection architecture if it is a p-survivable connection that has the largest bandwidth with respect to the 1+1 (correspondingly, 1:1) protection architecture. In section 6 we shall define additional quality criteria.

Finally, note that, whereas the widest p-survivable connection depends on the considered protection architecture, a most survivable connection for one architecture is also such for the other architecture.

3. Properties of Survivable Connections

In this section we establish several fundamental properties of survivable connections. We begin with a rather straightforward quantification of the probability of a survivable connection to remain operational upon a failure.

We are given a network $G(V,E)$ and a survivable connection $(\pi_1, \pi_2) \in P^{(s,t)} \times P^{(s,t)}$. Under the single link failure model, a link that is not common to both paths can never cause a survivable connection to fail. Similarly, a failure in a common link, causes a failure of the entire connection. Hence, the survivable connection (π_1, π_2) is operational *iff* for each $e \in \pi_1 \cap \pi_2$ it holds that e is operational, i.e., all the links that are common to both paths are operational. Therefore, the probability that a survivable connection remains operational upon a link failure is equal to the probability that all its common links are operational upon that failure. Thus, since link failure probabilities are independent, it holds that the probability that all common links are operational under the condition of a failure is equal to the product of their success probability under the condition of a failure. This is summarized as follows.

Property 1 Given are a survivable connection (π_1, π_2), and for each $e \in E$, a failure probability p_e. The probability that (π_1, π_2) is operational upon a failure event is equal to $\prod_{e \in \pi_1 \cap \pi_2} (1 - p_e)$.

We now turn to present a rather surprising property that shows that the employment of more than two paths is worthless. Consider a more general protection framework that admits *any* (≥ 2) number of paths. Basically, we show that, in any network and for each survivability constraint $0 \leq p \leq 1$, if there exists a p-survivable connection that admits *more than two paths*, then there exists a p-survivable connection that admits *exactly two paths*. Moreover, we show that the bandwidth of the widest p-survivable connection in protection frameworks where connections are allowed to employ any number of paths *is not larger* than the band-width of the widest p-survivable connection that is limited to at most two paths.

Remark 1: For completeness, we note that a p-survivable connection in protection frameworks that admit more than two paths is a collection of paths $(\pi_1, \pi_2, \cdots, \pi_k) \in P^{(s,t)} \times P^{(s,t)} \times \cdots \times P^{(s,t)}$ that has a probability of at least p to have at least one operational path after a failure. The bandwidth of such a connection with respect to the 1:1 protection architecture (i.e., in the case where the traffic is sent only over a single path) is the maximum $B \geq 0$ such that $B \leq b_e$ for each $e \in \bigcup_{i=1}^{k} \pi_i$. Similarly, the bandwidth of such a connection with respect to the 1+1 protection architecture (i.e., in the case where the traffic is carried independently over each path) is the maximum $B \geq 0$ such that $n \cdot B \leq b_e$ for each link $e \in E$ that is common to some n paths out of $(\pi_1, \pi_2, \cdots, \pi_k)$.

We are now ready to formulate two fundamental properties of survivable connections; the first corresponds to widest p-survivable connections and the second to most survivable connections. Due to space limits the proof of both properties is omitted and can be found in [2].

Property 2: Let $(\pi_1, \pi_2, \cdots, \pi_k) \in P^{(s,t)} \times P^{(s,t)} \times \cdots \times P^{(s,t)}$ be the most survivable connection in $G(V,E)$ and let $(\overline{\pi_1}, \overline{\pi_2}) \in P^{(s,t)} \times P^{(s,t)}$ be the most survivable connection in $G(V,E)$ that consists of at most two paths. The survivability level of $(\overline{\pi_1}, \overline{\pi_2})$ is not smaller than that of $(\pi_1, \pi_2, \cdots, \pi_k)$.

Property 3: Let $(\pi_1, \pi_2, \cdots, \pi_k) \in P^{(s,t)} \times P^{(s,t)} \times \cdots \times P^{(s,t)}$ be the widest p-survivable connection in $G(V,E)$ with respect to the 1:1 (alternatively, 1+1) protection architecture. There exists a p-survivable connection $(\overline{\pi_1}, \overline{\pi_2}) \in P^{(s,t)} \times P^{(s,t)}$ that has at least the bandwidth of $(\pi_1, \pi_2, \cdots, \pi_k)$ with respect to the 1:1 (correspondingly, 1+1) protection architecture.

The above key observations show that there is no incentive to define survivable connections that consist of more than two paths. Therefore, under the standard single link failure model, this finding indicates an important network design rule in terms of survivability.

4. Establishing *p*-survivable connections

In this section we show how to construct *p*-survivable connections for the 1+1 and 1:1 protection architectures. In view of the findings of the previous section, we focus on survivable connections that consist of at most two paths. We begin with the establishment of widest *p*-survivable connections and most survivable connections for the 1+1 protection architecture.

4.1. Establishing survivable connections for the 1+1 protection architecture

The first step towards the establishment of either widest *p*-survivable or most survivable connections is the development of an efficient algorithm that, for any $B \geq 0$, establishes a survivable connection with a bandwidth of at least *B* that has the maximum probability to remain operational upon a link failure. We term such a connection as the *most survivable connection with a bandwidth of at least B*.

Remark 2 Finding the most survivable connection with a bandwidth of at least *B* is beneficial per se. For example, in cases where the traffic demand γ is known in advance, it may be desired to establish a connection with a bandwidth of at least γ that has the maximum probability to remain operational upon a failure.

4.1.1. Establishing most survivable connections with a bandwidth of at least B.
We now establish an efficient algorithm that, for any $B \geq 0$, outputs the most survivable connection that has a bandwidth of at least *B*. Given a network $G(V, E)$, a pair of nodes *s* and *t*, a bandwidth constraint $B \geq 0$, and, for each link $e \in E$, a bandwidth $b_e \geq 0$ and a failure probability $p_e \geq 0$, the algorithm reduces the problem of finding the most survivable connection with a bandwidth of at least *B* into an instance of the Min Cost Flow problem [1]. In essence, the construction is based on a network transformation that considers three different cases, as illustrated in Figure 2. In the case of a link $e \in E$ with a bandwidth $b_e < B$, it follows by definition (Def. 4) that link *e* cannot be used by any survivable connection that has a bandwidth of at least *B*. Therefore, this link can be discarded from the network without any influence on the optimal solution. On the other hand, each link $e \in E$ that satisfies $b_e \geq 2 \cdot B$ can concurrently be used by both of the connection's paths in order to establish a survivable connection with a bandwidth of at least *B*. In that case, the corresponding link is transformed into two
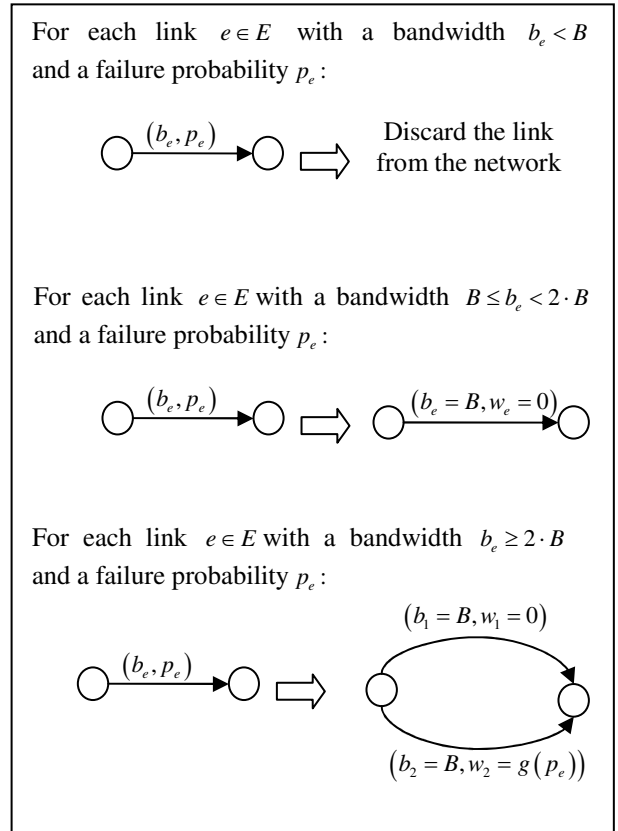

Figure 2. Finding the most survivable connection with a bandwidth of at least B (for the 1+1 protection architecture) by a reduction to the Min Cost Flow problem.

parallel links, each with a link bandwidth of *B*; however, whereas the first link is assigned with a zero weight, the other link is assigned with a weight that is a function ($g(p_e)$) of the link's failure probability (p_e). The reason for that stems from Property 1 (of the previous section) that shows that the degree of survivability of each connection is solely determined by its common links. More specifically, only when both of the connection's paths share the same link *e*, the link's failure probability p_e should be considered. Indeed, a Min Cost Flow (where "cost" is "weight") over the constructed network ensures that, when a single path traverses link *e*, the incurred cost is zero, whereas when both paths traverse through *e*, the cost $g(p_e)$ depends on the failure probability p_e ($g(p_e)$ shall be specified in the following). The third case corresponds to links that satisfy $B \leq b_e < 2 \cdot B$. In that case, at most one path with a bandwidth *B* can traverse through such a link without violating the link bandwidth b_e. Thus, these links can be transformed into links that have a bandwidth *B* without any effect on the optimal solution. In addition, since these links can be used by at most one path, their failure

probabilities should not be considered and therefore the transformed links are assigned zero weight.

Denote the transformed network as $\widetilde{G}\left(\widetilde{V},\widetilde{E}\right)$. The algorithm computes a min-cost flow $\{f_e\}$ with a flow demand of $2\cdot B$ units over the network $\widetilde{G}\left(\widetilde{V},\widetilde{E}\right)$ by employing any standard Min Cost Flow algorithm that returns an integral link flow when all link bandwidths $\{b_e\}$ are integral (see [1]). Since all link bandwidths in $\widetilde{G}\left(\widetilde{V},\widetilde{E}\right)$ are integral in B, the link flow $\{f_e\}$ is B-integral i.e., f_e is a multiple of B for each $e\in E$. Therefore, since the total traffic equals to $2\cdot B$ flow units, the *flow decomposition algorithm* [1] can be applied in order to decompose the link flow $\{f_e\}$ into a flow over two paths π_1,π_2 such that each carry B flow units from s to t. Moreover, since the flow has minimum cost, it follows that

$$\sum_{\widetilde{e}\in\widetilde{E}} f_{\widetilde{e}}\cdot w_{\widetilde{e}} = \sum_{e\in\pi_1\cap\pi_2} B\cdot g\left(p_e\right) = B\cdot\sum_{e\in\pi_1\cap\pi_2} g\left(p_e\right) \quad \text{has}$$

minimum value. Thus, $\sum_{e\in\pi_1\cap\pi_2} g\left(p_e\right)$ has minimum value. Finally, if we define $g\left(p_e\right)\triangleq -\ln\left(1-p_e\right)$ for each $e\in E$, the algorithm defines a pair of paths π_1,π_2 that minimizes $-\sum_{e\in\pi_1\cap\pi_2}\ln\left(1-p_e\right) = -\ln\prod_{e\in\pi_1\cap\pi_2}\left(1-p_e\right)$ and therefore maximizes $\ln\prod_{e\in\pi_1\cap\pi_2}\left(1-p_e\right)$. Thus, the connection $\left(\pi_1,\pi_2\right)$ maximizes $\prod_{e\in\pi_1\cap\pi_2}\left(1-p_e\right)$ which, according to Property 1, equals to the probability that the connection $\left(\pi_1,\pi_2\right)$ is operational upon a failure. The formal description of the algorithm, termed *Algorithm B-Width Most Survivable Connection*, appears in [2].

The following theorem shows that, for every $B\geq 0$, our algorithm establishes the most survivable connection with a bandwidth of at least B.

Theorem 1: Given are a network $G\left(V,E\right)$, a pair of nodes s and t, a bandwidth constraint $B\geq 0$, and, for each link $e\in E$, a bandwidth $b_e\geq 0$ and a failure probability $p_e\geq 0$. If there exists a survivable connection with a bandwidth of at least B, then Algorithm B-Width Most Survivable Connection returns the most survivable connection with a bandwidth of at least B; otherwise, the algorithm fails.

Due to space limits the proof is omitted. It is based on the ideas that were described above and can be found in [2].

4.1.2. Establishing most survivable and widest p-survivable connections. Finally, we are ready to construct most survivable connections and widest p-survivable connections for the 1+1 protection architecture. As is easy to see, the most survivable connection with a bandwidth of at least *B*=0 is in essence also a most survivable connection. As the corresponding problem is a special case of the problem that was addressed in the previous subsection, in this section we only focus on the establishment of widest *p*-survivable connections.

In order to establish the widest *p*-survivable connection, we employ Algorithm B-Width Most Survivable Connection. Specifically, given a network and a survivability constraint *p*, we search for the largest value of *B* such that the most survivable connection with a bandwidth of at least *B* is a *p*-survivable connection i.e., has a probability of at least *p* to remain operational upon a link failure. Obviously, this strategy is attractive only if we consider a small number of bandwidth constraints before we get to the bandwidth of the widest *p*-survivable connection. Fortunately, in the following we show that it is sufficient to consider $O\left(\log N\right)$ bandwidth constraints in order to find the bandwidth of the widest *p*-survivable connection.

First, we observe that, for every given network, the bandwidth of the widest *p*-survivable connection belongs to a set of at most $2\cdot M$ values. To see this, recall that the bandwidth of each survivable connection $\left(\pi_1,\pi_2\right)$ with respect to the 1+1 protection architecture, is defined as the maximum $B\geq 0$ such that $2\cdot B\leq b_e$ for each $e\in\pi_1\cap\pi_2$ and $B\leq b_e$ for each $e\in\left(\pi_1\cup\pi_2\right)\backslash\left(\pi_1\cap\pi_2\right)$. Hence, if the survivable connection $\left(\pi_1,\pi_2\right)$ admits a link $e\in E$, then by definition, its bandwidth with respect to the 1+1 protection, is not larger than either $\dfrac{b_e}{2}$ (for $e\in\pi_1\cap\pi_2$) or b_e (for $e\in\left(\pi_1\cup\pi_2\right)\backslash\left(\pi_1\cap\pi_2\right)$). Moreover, it follows by definition that there exists at least one link $e\in\pi_1\cup\pi_2$ such that the bandwidth of $\left(\pi_1,\pi_2\right)$ is either $\dfrac{b_e}{2}$ or b_e. Therefore, each survivable connection in $G\left(V,E\right)$ has a link $e\in E$ whose bandwidth is either $\dfrac{b_e}{2}$ or b_e. In particular, the bandwidth of the widest *p*-survivable connection in the network, denoted as B^*, must belong to the set $\mathbb{B}\triangleq\left\{\dfrac{b_e}{k}\big|e\in E,\ k=1,2\right\}$, which consists of at most $2\cdot M$ members.

Remark 4 Note that we can employ a *binary search* over the set $\mathbb{B}$ in order to find the value of B^*. Indeed, for each $B \in \mathbb{B}$, if the most survivable connection with a bandwidth of at least B is a p-survivable connection then so are all the other most survivable connections with bandwidths of at least $B', B' \leq B$; on the other hand, when the most survivable connection with a bandwidth of at least B is *not* a p-survivable connection, then none of the most survivable connections with bandwidth of at least $B'', B'' > B$, is a p-survivable connection.

The formal specification of the algorithm appears in [2].

Finally, we consider the complexity incurred by the establishment of most survivable connections and widest p-survivable connections. To that end, we denote by $T(N, M)$ the running time of any standard min-cost flow algorithm for an N-nodes M-links network. Since Algorithm B-Width Most Survivable Connection solves a single instance of the min-cost flow problem, the complexity of establishing most survivable connections and widest p-survivable connections is $O(T(N, M))$ and $O(T(N, M) \cdot \log N)$, respectively.

Remark 5 We note that it is possible to solve the min-cost flow problem in $O((M \cdot \log N) \cdot (M + N \cdot \log N))$ operations [1]; hence, we can establish widest p-survivable connections and most survivable connections within a total complexity of $O(M^2 \cdot \log^2 N + M \cdot N \cdot \log^3 N)$ and $O(M^2 \cdot \log N + M \cdot N \cdot \log^2 N)$, respectively.

4.2. Establishing survivable connections for the 1:1 protection architecture

We turn to establish survivable connections for the 1:1 protection architecture. Obviously, the most survivable connection in the 1+1 protection architecture is the same as that of the 1:1 protection architecture; therefore, we will only consider the establishment of widest p-survivable connections for the 1:1 protection architecture. Moreover, as the establishment of the widest p-survivable connection with respect to the 1:1 protection architecture is quite similar as for the 1+1 protection architecture, we only sketch the main ideas.

As before, we begin by finding a solution to the dual problem of establishing the most survivable connection with a bandwidth of at least B (however, this time the bandwidth is computed according to the 1:1 protection architecture). This is based on a reduction that is similar to the one presented in Fig 2. However, as the band-

width of any survivable connection (π_1, π_2) for the 1:1 protection architecture is defined as the largest $B \geq 0$ such that $B \leq b_e$ for each $e \in \pi_1 \cup \pi_2$, it follows that only two cases should be considered in the reduction, namely $b_e < B$ and $b_e \geq B$. More specifically, as before, all the links that satisfy $b_e < B$ should be discarded from the network since they cannot be used in order to construct a survivable connection with a bandwidth of at least B. However, in contrast to the solution of the 1+1 protection architecture, all other links can be concurrently employed by the pair of paths that constitute the survivable connection. More precisely, the only difference between the reduction that corresponds to the 1+1 protection architecture and the reduction that corresponds to the 1:1 protection architecture, is the type of links that can be used by both paths; namely, whereas in the 1+1 protection architecture, the most survivable connection with a bandwidth of at least B cannot employ a link $e \in E$ that satisfies $B \leq b_e < 2 \cdot B$ for both paths, in the 1:1 protection architecture such a link can be common to both paths. The reduction for the 1:1 protection architecture is illustrated in Fig. 3.

For each link $e \in E$ with a bandwidth $b_e < B$ and a failure probability p_e:

(b_e, p_e) → Discard the link from the network

For each link $e \in E$ with a bandwidth $b_e > B$ and a failure probability p_e:

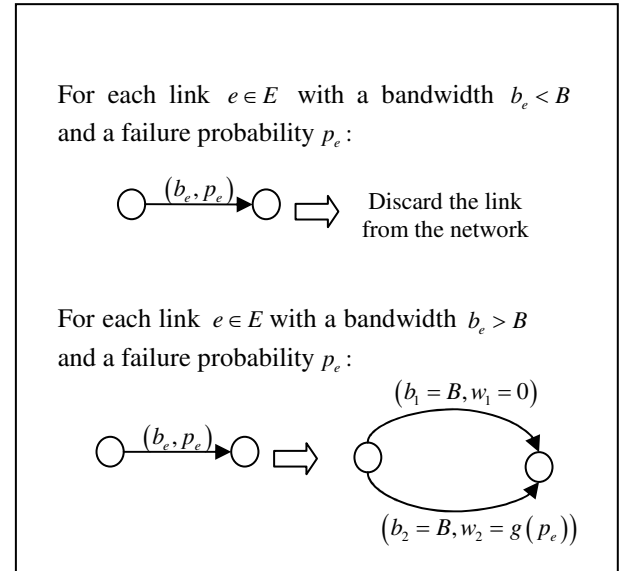(b_e, p_e) → $(b_1 = B, w_1 = 0)$ $(b_2 = B, w_2 = g(p_e))$

Figure 3. Finding the most survivable connection with a bandwidth of at least B for the 1:1 protection architecture by a reduction to the Min Cost Flow problem.

As before, given a scheme for constructing most survivable connections with a bandwidth of at least B, we employ a binary search in order to find the largest B such that the most survivable connection with a bandwidth of at least B is a p-survivable connection. However, this time the bandwidth of the widest p-survivable connection belongs to the set $\{b_e | e \in E\}$, which consists of at most M elements (as opposed to the previous case where it belongs to a set of at most $2 \cdot M$ elements). To see this, note that, by definition, the bandwidth of the

survivable connection (π_1, π_2) with respect to the 1:1 protection architecture is the bandwidth of its bottleneck link i.e., $\min_{e \in \pi_1 \cup \pi_2} \{b_e\}$. Therefore, the bandwidth of each survivable connection with respect to the 1:1 protection architecture is determined by some link in $e \in E$ i.e., it belongs to $\{b_e | e \in E\}$.

5. A Hybrid protection architecture

Thus far, we have focused on the 1+1 and 1:1 protection architectures. However, the tunable survivability concept gives rise to an efficient third protection architecture, which is a *hybrid* approach that combines the 1:1 and 1+1 protection architectures. More specifically, given a survivable connection (π_1, π_2) with a traffic demand γ, we present a new architecture that, for a connection (π_1, π_2), transfers γ flow units over the links in $\pi_1 \cap \pi_2$, as in 1:1 protection, while over the links in $(\pi_1 \cup \pi_2) \setminus (\pi_1 \cap \pi_2)$, it transfers γ flow units, as in 1+1 protection. This new architecture is illustrated through the following example.

Example 2: Consider the network depicted in Fig. 4. Suppose that we are given a survivable connection (π_1, π_2) such that $\pi_1 = (e_1, e_3, e_4)$ and $\pi_2 = (e_2, e_3, e_5)$. Hybrid Protection transfers one duplicate of the original traffic through link $e_1 \in \pi_1$ and another duplicate through link $e_2 \in \pi_2$. While both duplicates arrive to node u, only the <u>first</u> to arrive is assigned to link $u \to v$ and the other one is discarded. When the duplicate that was assigned to $u \to v$ arrives to v, Hybrid Protection transfers one duplicate through link $e_4 \in \pi_1$ and another through link $e_5 \in \pi_2$. Finally, as with 1+1 protection, node t considers only the duplicate that is the first to arrive. Note that such an assignment of traffic to links is <u>not</u> a flow.
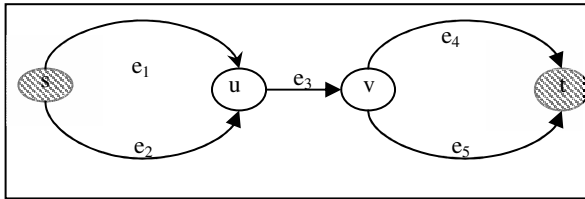

Fig. 4: The Hybrid Protection Architecture

Hybrid Protection has several important advantages. First, it reduces the congestion of all links that are shared by both paths with respect to 1+1 protection. At the same time, upon a link failure, it has a faster restoration time than 1:1 protection. Finally, it provides the *fastest* propagation of data with respect to the propaga-

tion time of *all* paths that can be constructed out of the links in $\pi_1 \cup \pi_2$. We demonstrate this property on the above example. Assume that the link propagation delays satisfy $d_{e_1} < d_{e_2}$ and $d_{e_5} < d_{e_4}$. Then, by construction, node u assigns the incoming flow of link e_1 over link e_3, and node t considers only the duplicate of link e_5. Thus, data is propagated through the path $\pi = (e_1, e_3, e_5)$, which has the *minimum propagation delay* among all the paths that can be constructed out of the links in $\pi_1 \cup \pi_2$; in particular, the delay of the path π is smaller than the delays of π_1 and π_2.

The above advantages notwithstanding, the implementation of the Hybrid Protection architecture requires additional nodal capabilities in comparison with the 1+1 and 1:1 architectures. To see this, note that node u in the example must be able to discard all the duplicates that it encounters for the second time i.e., the duplicates that contain data that was already sent to node v. This is in contrast to the 1+1 protection architecture, where such functionality is required only from the destination, and the 1:1 protection architecture, where this functionality is not required at all.

Finally, note that the Hybrid Protection architecture transfers through each link exactly one duplicate of the original traffic. Hence, the maximum traffic rate that can be transferred through a survivable connection (π_1, π_2) with respect to Hybrid Protection is bounded by $\min_{e \in \pi_1 \cup \pi_1} \{b_e\}$. In other words, the bandwidth of the survivable connection (π_1, π_2) with respect to Hybrid Protection is the maximum $B \geq 0$ such that $B \leq b_e$ for each $e \in \pi_1 \cup \pi_2$. Since this is precisely the definition of bandwidth with respect to 1:1 protection, the widest *p*-survivable connection with respect to Hybrid Protection is also the widest *p*-survivable connection with respect to 1:1 protection. Hence, we can employ the solution for 1:1 protection in order to establish widest *p*-survivable connections for Hybrid Protection. Nevertheless, it is important to note that, while 1:1 protection assigns traffic only to the links that belong to either π_1 or π_2, Hybrid Protection assigns traffic to *all* the links in $\pi_1 \cup \pi_2$.

6. Quality of Service Extensions

For any pair of nodes in a given network, there might be several widest *p*-survivable connections as well as several most survivable connections. Among them, we may be interested in those that optimize some QoS targets, such as end-to-end delay, jitter, cost, etc. Such (additive) metrics can be represented by *weights*

$\{w_e\}$. In [2] we investigate most survivable and widest p-survivable connections that have the minimum total weight. More precisely, given a network and a survivability constraint p, we show in [2] how to modify the schemes of Section 4 above, so as to establish widest p-survivable connections as well as most survivable connections that have a minimum total weight $\sum_{e \in \pi_1 \cup \pi_1} w_e$.

7. Simulation Results

The goal of this section is to demonstrate *how much* we gain by employing tunable survivability instead of traditional "full" survivability. To that end, we first compare between the maximum bandwidth of survivable connections that consist of a pair of disjoint paths (i.e.,1-survivable connections) and the maximum bandwidth of p-survivable connections, where $p \in [0,1)$. Then, we compare between the feasibility of both approaches i.e., the incidences where the establishment of pairs of disjoint paths is impossible and the incidences where the establishment of p-survivable connection is impossible. Through comprehensive simulations, we show that, at the price of a marginal reduction in the common requirement of 100% protection, a major increase in bandwidth as well as in feasibility is accomplished.

We generated two types of random networks: network topologies that follow the four power laws defined by [3] (henceforth: *power-law* topologies), and networks with a *flat topology* i.e., Waxman networks [13] (henceforth: *flat* topologies). Then, we constructed 10,000 random networks for each combination of the following three items: (a) the degree of survivability $p \in [0,1]$; (b) the type of protection architecture (i.e., either 1+1 or 1:1); and (c) the class of random networks (i.e., either power-law or flat). For each network, we identified a source-destination pair. We then conducted the following measurements: (1) We measured the number of networks $N(p)$ that admits a p-survivable connection among the 10,000 networks; we then derived the *feasibility ratio* $\rho_N(p) \triangleq \frac{N(p)}{N(1)}$; (2) for each of the $N(1)$ networks that admit 1-survivable connections, we measured the ratio $\frac{B(p)}{B(1)}$, where $B(p)$ denotes the bandwidth of the widest p-survivable connection, and derived the corresponding *bandwidth ratio* $\overline{\rho_B(p)}$,

which is the average value of $\frac{B(p)}{B(1)}$ over the corresponding $N(1)$ networks.

In all runs, we assumed that the link bandwidths are distributed uniformly in [5,150] MB/sec and the failure probability of each link is distributed normally with a mean of 1% and a standard deviation of 0.3%. Our construction of flat and power-law topologies followed the lines of [13] and [12] respectively. The precise way that we generated each type of random network is specified in [2].

We turn to present our results. First, we note that the value $N(1)$ i.e., number of networks that admitted 1-survivable connections, was in the range 4,000-7,000 (out of 10,000), hence the samples were always significant. In Fig. 5 we depict the bandwidth ratio $\overline{\rho_B(p)}$ versus the level of survivability $p \in [0.95,1]$ for 1:1 protection. In particular, we show that, with a reduction of 2% in the requirement of full survivability,[1] the bandwidth is increased by 51% for Waxman networks and 100% for power law networks. Due to space limits the results that corresponds to the 1+1 protection architecture are omitted from this version and can be found in [2].
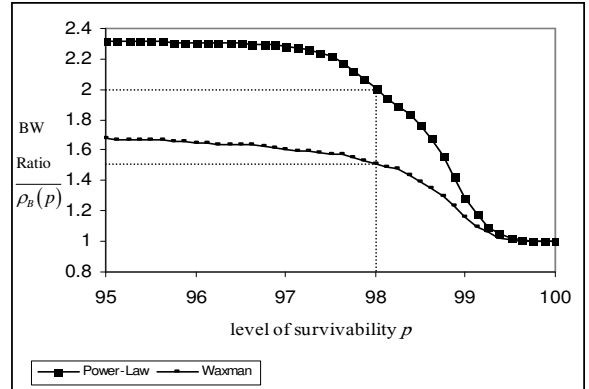

Figure 5. The average ratio between the bandwidths of widest p-survivable connections and widest 1-survivable connections in the 1:1 protection architecture.

In Fig. 6, we depict the ratio between the number of networks that have at least one feasible p-survivable connection and the number of networks that have at least one feasible 1-survivable connection; to that end, we present the feasibility ratio $\rho_N(p)$ versus the level of survivability $p \in [0.95,1]$. Note that the feasibility ratio is independent of the employed protection architecture; therefore, the corresponding results hold for

---

[1] We emphasize that these are 2% *given* the event of a network failure. Hence, the *a-priory* probability is much lower.

both protection architectures. Also, note that, with a reduction of 2% in the requirement of full survivability, the feasibility ratio is increased by 54% for Waxman networks and by 127% for power law networks.
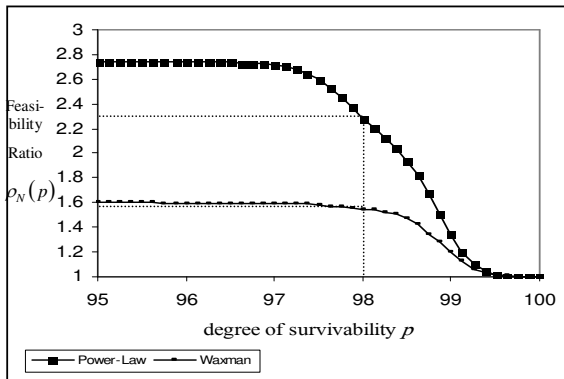

Figure 6. The ratio between the number of networks with at least one feasible p-survivable connection and the number of networks with at least one feasible 1-survivable connection.

8. Conclusions

Standard survivability schemes enhance the ability to recover from network failures by establishing pairs of disjoint paths. However, in practice, this approach is too restrictive and often leads to the selection of poor routing paths (if any). In this work, we have proposed a novel quantitative approach for network survivability. The new approach allows to alleviate the rigid path disjointedness requirement, which considers only full (100%) protection, into a weaker requirement, which can be tuned to accommodate any desired degree (0%-100%) of survivability. Just as in the standard approach, we have shown that the new approach can also be accommodated by efficient polynomial (optimal) schemes. However, as opposed to the original approach, the new approach allows a flexible choice of the desired degree of survivability, hence enabling to consider important tradeoffs. Moreover, since a 1-survivable connection is also p-survivable (for any value of p), our approach always offers a solution of at least (and usually a higher) quality than the traditional approach.

We have characterized several properties of the new approach. In particular, we established that, under the single link failure model, there is no benefit in establishing survivability schemes that employ more than two paths per connection. Since the single link failure assumption is practically valid in many cases of interest, this finding suggests an important network design rule in terms of survivability.

We evaluated the power of the new approach through comprehensive simulations. Our results clearly demonstrate the advantages of tunable survivability over full survivability. In particular, all measurements have shown that, by alleviating the traditional requirement of full survivabil-

ity by just 2%[1], we obtained major improvements in the quality of the solutions. Effectively, this indicates that (traditional) full protection levies an excessive price.

Finally, we have shown that the tunable survivability approach gives rise to a new protection architecture that poses several advantages over current architectures; moreover, the new architecture was shown to admit efficient optimal schemes.

The above notwithstanding, the practical deployment of the tunable survivability approach still posses several challenges. As mentioned, the hybrid protection architecture requires additional capabilities from transit nodes and the efficient implementation of these capabilities is an interesting issue for future work. At a more general level, the distributed implementation of all our algorithms as well as the development of simpler heuristic schemes are two major issues that have to be considered in practice.

In summary, while much is still to be done towards the actual deployment of the tunable survivability approach, this study provides ample and firm evidence of its major benefits and potential practical feasibility.

9. References

[1] R. K. Ahuja, T. L. Magnanti and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall, 1993.

[2] R. Banner and A. Orda, *The Power of Tuning: A Novel Approach for the Efficient Design of Survivable Networks*, CCIT Report No. 463, Electrical Engineering Dept., Technion, Haifa, Isreal, January 2004. Availble from: ftp://ftp.technion.ac.il/pub/supported/ee/Network/bo.pdf

[3] M. Faloutsos, P. Faloutsos, and C. Faloutsos, *On Power-law Relationships of the Internet Topology*, in Proceedings of ACM SIGCOMM, cambridge, MA, September 1999.

[4] A. Fumagalli and M. Tacca, *Optimal Design of Optical Ring Networks with Differentiated Reliability (DiR)*, in Proceedings of the International Workshop on QoS in multiservice IP networks, Rome, Italy, Janurary 2001.

[5] Michael T. Frederick and Arun K. Somani, *A Single-Fault Recovery Strategy for Optical Networks Using Subgraph Routing*, in Proceedings of the 7th IFIP Working Conference on Optical Network Design and Modeling, February 2003.

[6] P. Ho, J. Tapolcai, and H. T. Mouftah, *On Achieving Optimal Survivable Routing for Shared Protection in Survivable Next-Generation Internet*, IEEE Transactions on Reliability, March 2004.

[7] M. Kodialam and T. V. Lakshman, *Restorable Dynamic Quality of Service* Routing, IEEE Communication Magazine, vol. 40, no. 6, June 2002.

[8] W. Lai, Ed. and D. McDysan, Ed., *Network Hierarchy and Multilayer Survivability*, IETF RFC 3386, November 2002.

[9] G. Maier, A. Pattavina, S. De Patre, and M. Martinelli, *Optical Network Survivability: Protection Techniques in the WDM Layer*, Photonic Networks Communications, vol. 4, no. 3-4, July-Dec. 2002.

[10] G. Mohan and A. K. Somani, *Routing Dependable Connections with Specified Failure Restoration Guarantees in WDM Networks,* in Proceedings of IEEE Infocom, Tel Aviv, Israel, April 2000.

[11] G. Mohan and C.S.R. Murthy, *Lightpath Restoration in WDM Optical Networks*, IEEE Network , Volume: 14, Issue: 6 , Nov.-Dec. 2000.

[12] C. R. Palmer and J. G. Steffan, *Generating Network Topologies that Obey Power Laws*, in Proceedings of the Global Internet Symposium, (Globecom) 2000, San Fransisco, CA, November 2000.

[13] B. M. Waxman, *Routing of Multipoint Connections*, IEEE Journal on Selected Areas in Communications, 6:1617-1622,1988.

[14] Dahai Xu and Chunming Qiao, *Distributed Partial Information Management (DPIM) Schemes for Survivable Networks - Part 2*, in Proceedings of IEEE Infocom, NY, June 2002.

---

[1] and much less in terms of the *a-priory* probability.