University of New Mexico
UNM Digital Repository

Electrical and Computer Engineering ETDs | Engineering ETDs

7-2-2012

Design and evaluation of network survivability schemes for correlated multi-failure scenarios

Oscar A. Díaz

Follow this and additional works at: https://digitalrepository.unm.edu/ece_etds

Recommended Citation
Díaz, Oscar A.. "Design and evaluation of network survivability schemes for correlated multi-failure scenarios." (2012).
https://digitalrepository.unm.edu/ece_etds/61

This Thesis is brought to you for free and open access by the Engineering ETDs at UNM Digital Repository. It has been accepted for inclusion in Electrical and Computer Engineering ETDs by an authorized administrator of UNM Digital Repository. For more information, please contact disc@unm.edu.

Oscar A. Diaz

Candidate

Electrical and Computer Engineering

Department

This thesis is approved, and it is acceptable in quality and form for publication:

Approved by the Thesis Committee:

Dr. Nasir Ghani , Chairperson

Dr. Majeed M. Hayat

Dr. Marios S. Pattichis

Design and Evaluation of Network Survivability Schemes for Correlated Multi-Failure Scenarios

by

Oscar A. Diaz

Bachelor of Telecommunication Engineering, 2010

THESIS

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Master of Science
Computer Engineering

The University of New Mexico

Albuquerque, New Mexico

May, 2012

©2012, Oscar A. Diaz

Dedication

To my father, mother, brother, sister, girlfriend and friends ...

Acknowledgments

I would like to acknowledge and extend my heartfelt gratitude to the following persons who have made the completion of this work possible:

My advisor, Professor Nasir Ghani, for his help, dedication, research advisement, and more importantly, for becoming a very good friend,

Professor Majeed Hayat, for all the help he gave me since the beginning of this process,

Mr. Feng Xu, who is an important part of this research and helped me when I needed, as well as Ms. Kaile Liang for help with the editing,

Zach Gould and Hunter Riley, who made this time in Albuquerque pass by like being back home in my lovely Concepcion,

All my friends, in particular my high school friends, who always put a smile on my face,

My parents and family, for their endless love and support,

And most especially, to my girlfriend and future copilot of my life, Nicolle Lathrop. I would not be here without her support and love.

Design and Evaluation of Network Survivability Schemes for Correlated Multi-Failure Scenarios

by

Oscar A. Diaz

Bachelor of Telecommunication Engineering, 2010

M.S., Computer Engineering, University of New Mexico, 2012

Abstract

Wireline high-speed networks have become a critical part of modern cyberinfrastructures and provide the base substrates to support a full range of higher-layer user services and applications. Indeed, a wide range of technologies have been deployed in these domains, ranging from ultra-fast *Internet Protocol* (IP) packet routing systems to multi-wavelength optical switching nodes. Today these setups provide immense levels of traffic scalability, reaching well into the 100s of gigabits/second and even terabits/second ranges. Owing to this growth, network survivability is now a central concern, as even a single link or node failure can cause widespread service disruption for thousands of users or more.

Now over the years, a full range of network survivability schemes have been developed for packet routing and optical switching networks. Indeed, the open research literature lists many types of solutions here, broadly classified as pre-fault protection and post-fault restoration strategies. The former schemes pro-actively setup backup (redundant) resource pools to overcome anticipated failure events. Meanwhile the

latter strategies are more reactive by design and attempt to re-establish connectivity after failures. By and large, the bulk of these solutions are only concerned with single failure recovery, i.e., either at the link or node level. In general, these are the most common types of faults events experienced in operational networks. However, recent developments and considerations are pushing the need for more capable schemes to recover from *multiple* failure events, i.e., as occurring during natural disasters, massive power outages, and *weapon of massive destruction* (WMD) type attacks. Indeed, these types of scenarios are much more challenging, as they induce large numbers of *correlated* failures which can quickly overwhelm most traditional single-failure recovery schemes.

Along these lines, some recent studies have looked at network recovery under massive correlated network failures. The key idea here is to introduce probabilistic risk information into the path provisioning (routing, protection) processes in order to minimize vulnerability to random failures. However, even though these schemes can reduce connections failure rates, they yield very high resource inefficiencies (usage consumption). In turn, these concerns will inhibit their adoption in most practical network settings, as operators have to balance the need for improved resiliency with revenue generation. To address this challenge, this thesis proposes a novel multi-failure survivability scheme that *jointly* incorporates both risk mitigation and *traffic engineering* (TE) efficiency objectives. In particular, the approach leverages multi-path routing strategies to first compute a selection of diverse working/backup path pairs and then uses ranking methods to select the most balanced combination. This framework applies graph-theoretic principles and hence can readily be integrated into real-world traffic provisioning systems.

The performance of the proposed solution is evaluated using discrete event simulation techniques for a variety of network topologies and compared against several existing schemes. Overall findings show that the scheme yields notably improved

survivability rates as compared to vanilla traffic engineering policies. At the same time, it also gives much better operational resource efficiencies versus existing probabilistic risk reduction routing strategies. Hence network carriers can fully leverage this new design to achieve much-improved reliability for critical data flows without sacrificing operational revenues.

Contents

List of Figures

Acronyms

BBR bandwidth blocking rate

DES discrete event simulation

DWDM dense wavelength division multiplexing

DiffServ differentiated service

GUI graphical user interface

HC hop count

HSLP hybrid selective link protection

HSSP hybrid selective segment protection

IAT inter arrival time

ILP integer linear program

INLP integer nonlinear program

IP Internet Protocol

JPP joint path pair

LAN local area network

LB load balancing

MF multi-failure

MPLS multiprotocol label switching

Acronyms

MSMF maximum survivability under multiple failures

MSPA maximum survivability protection algorithm

pdf probability density function

p-cycle protection cycle

p-SRLG probabilistic shared risk link group

QoS quality of service

RM risk minimization

SDP shortest disjoint path

SHR self-healing rings

SP shortest path

SRLG share risk link group

TE traffic engineering

TEPP traffic engineer path pair

WDM wavelength-division multiplexing

WMD weapon of massive destruction

Chapter 1

Introduction

High-bandwidth networks have become an indispensable part of modern society. These infrastructures provide the base communication substrates from which to implement a full range of end-user services. As a result many commercial, government, and defense organizations maintain and operate a full range of high-speed network infrastructures to meet the growing needs of their users. For example, at the lowest fiber level, optical *dense wavelength division multiplexing* (DWDM) backbones are being used to route gigabit-level lightpath connections across large regional and national distances. Meanwhile, overlying *Internet Protocol* (IP) and Ethernet technologies are also being deployed to provide finer granularity bandwidth control. Indeed, many of these setups support full *quality of service* (QoS) provisions via standardized frameworks such as IP *multiprotocol label switching* (MPLS) [1] and *differentiated service* (DiffServ) [2].

1.1 Background

As network infrastructures have grown and expanded, survivability concerns have also become more critical. Namely, given the increased scalability of many new technologies, even a single failure can now cause massive service disruptions for user services. For example, a DWDM link failure can easily disrupt hundreds of gigabits/sec of users flows.

In light of the above, a wide range of network failure recovery schemes have been developed [3]. Broadly, these strategies can be classified as *pre-provisioning* or *post-provisioning*, shown in Figure 1.1. Specifically, the former are usually termed as "protection" approaches and pre-compute backup routes for all (or part of) the working end-to-end routes. For example, point-to-point link protection schemes have been widely used in optical DWDM backbones to provide rapid recovery [4]. Additionally, many protection strategies have also been developed for end-to-end path recovery [5]. These schemes compute backup routes for working (primary) connections and perform appropriate switch-overs upon failure detection. Within these strategies, many researchers have also looked *shared* protection variants to further improve resource utilization on idle backup routes [3]. However, shared protection algorithms are strictly premised on single link failure assumptions, i.e., no more than one link will fail at any given time. Meanwhile, post-provisioning schemes are usually termed as "restoration" and do not implement pre-computation or pre-reservation of recovery resources. Instead, active post-fault computation and signaling is done to recover failed routes. These schemes, generally, cannot guarantee any level of fault recovery (even against single link failures) and provide slower recovery. However, restoration usually gives much better resource utilization.

Although there are many existing network recovery solutions, they are mostly geared towards single link/node failure recovery. As such, these approaches will

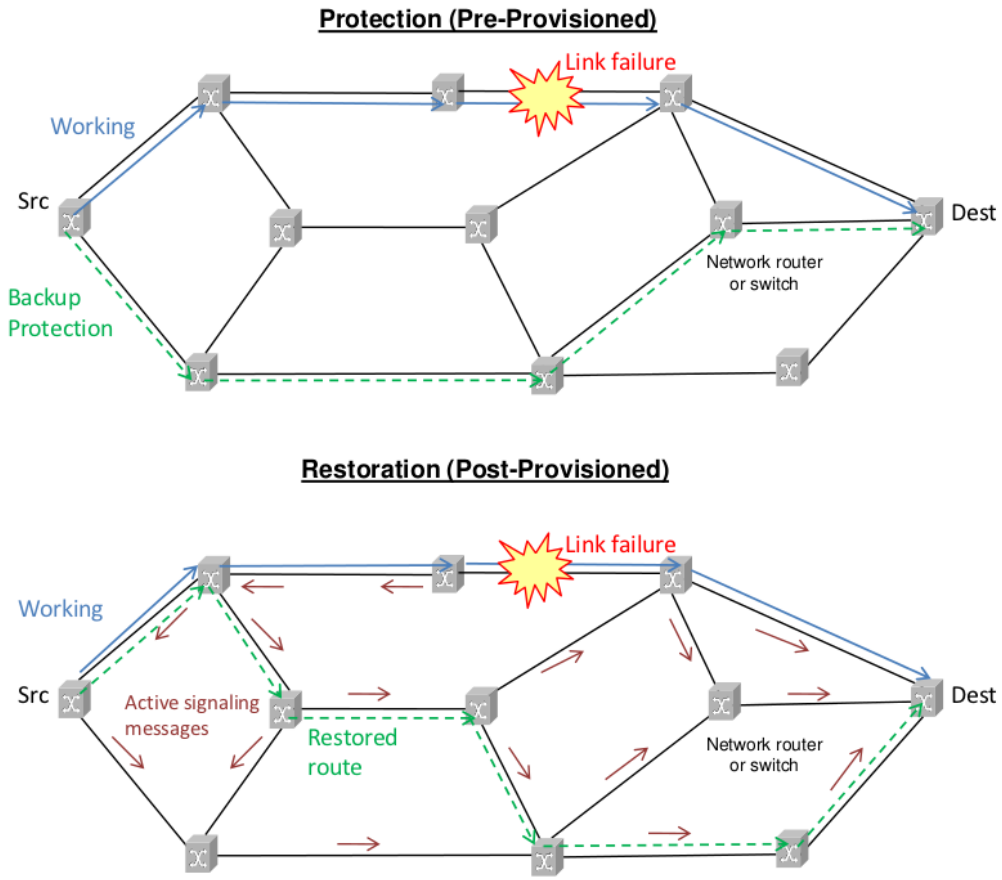

Figure 1.1: Protection and restoration strategies

be less efficient in regional multi-failure scenarios, particularly those with *correlated* and *cascading* failures. For example these failure conditions can arise during natural disasters (such as earthquakes, floods, hurricanes, tornadoes, etc), large scale power outages, and even malicious *weapon of massive destruction* (WMD) attacks. Point-to-point link protection and shared path protection strategies will be particularly vulnerable here. Hence in order to address these concerns, some recent studies have proposed newer protection strategies to handle randomized multi-failure scenarios. However, it is generally understood that pre-provisioned schemes cannot

guarantee recovery against all possible multi-failure events. As a result, researchers have assumed *probabilistic* link failure regions and used various algorithms to try to minimize the failure probability of end-to-end primary/backup path pairs [6],[7],[8]. Nevertheless, these approaches purely focus on "risk-minimization" and tend to yield longer routes with higher resource utilization. In turn, this leads to higher blocking and lower revenues for network carriers.

In light of the above, there is a clear need to develop further protection schemes to handle *multiple correlated* failures. Moreover, these new designs should also address *traffic engineering* (TE) resource efficiency concerns in order to deliver more meaningful solutions for networks carriers. Finally, these algorithms should also have acceptable computational complexity and be able to integrate with existing network provisioning frameworks, i.e., routing protocols, path computation engines, etc.

1.2 Motivations

Given the increased vulnerability of large network infrastructures to catastrophic events, there is a growing need to develop more effective network survivability schemes. Ideally, these solutions should make provisions for multiple correlated failures and incorporate probabilistic risk information into the recovery process. However, even though some related schemes have been proposed in [6], [7], [8], these strategies only consider risk minimization objectives. As a result, their TE performance is generally poor and can lead to high resource consumption and blocking. In order to address these crucial concerns, this thesis proposes a more balanced network protection solution which incorporates both failure recovery and TE needs.

1.3 Problem Statement

This thesis proposes a new multi-failure recovery scheme for high-speed networks. The solution is directly applicable to IP MPLS domains and also can be tailored (with minor modifications) for use in optical DWDM backbones as well. The performance of the proposed solution is evaluated and compared against various existing path protection strategies using network simulation.

1.4 Scope and Objectives

The main objective of this thesis is to design a robust *multi-failure* route protection solution. The work uses a graph-theoretic heuristics approach and assumes complete network resources knowledge at the provisioning entity, i.e., centralized control. A second objective is to also measure the effectiveness of the proposed solution versus that yielded by some existing schemes.

1.5 Research Approach

To achieve the above objectives, this thesis effort focuses on three core tasks. First, a detailed survey is conducted to review the latest research work in dual and multi-failure recovery. This framework is then used as a baseline to develop an improved (balanced) multi-failure protection scheme. Finally, detailed coding and simulation studies are done to evaluate the performance of the proposed solution in comparison to existing solutions. Specifically, all analyses are conduced using custom-developed models in the *OPNET Modeler*^{TM} simulation tool environment (i.e., C/C++ programming).

1.6 Thesis Outline

This thesis report is organized as follows. First, Chapter 2 presents a survey of the different approaches and latest techniques for dual-failure and multi-failure protection in high-speed networks. Next, Chapter 3 presents the proposed multi-failure protection solution, complete with detailed pseudocode listings. Chapter 4 then focuses on simulation design and defines the appropriate topologies used along with key performance evaluation metrics. Subsequently, Chapter 5 analyses the recovery and TE performances of the scheme for various failure scenarios. Finally, conclusions and further research directions are outlined in Chapter 6.

Chapter 2

Network Survivability: A Review

Network recovery from failed connections has always been a critical concern and many different strategies have been developed and studied over the years. In particular, most of these schemes are tailored to handle specific failure conditions. Along these lines, this chapter presents a brief survey of this topic area. Namely, the first part focuses on defining specific failure types and providing a general high-level taxonomy of different survivability schemes. Meanwhile the second part of the chapter focuses on specialized dual- and multi-failure recovery solutions, which are the main focus of this thesis study.

2.1 Failure Models and Survivability Mechanisms

According to [9], a *failure* is defined as a disruption in traffic caused by a malfunction of one or more components. Now, the wider literature lists many different types of failures that may occur in network settings. In turn, these failures have a strong impact on the type of recovery schemes developed. Overall, the most common failure types are :

- Link failures: These disruptions are caused by damage to an optical fiber or a link component. The common reasons here can include fiber cuts, malfunctioning port interfaces, and damaged amplifier and regenerators (in the case of optical DWDM networks). Overall a link failure is considered as an individual failure as it only affects a single component in a network.

- Node failures: These failures occur when a whole router or switching node goes down and are generally less common and more serious in nature. The common causes here include operator errors or power outages. Regardless, node failures are still considered as individual failures even though they may induce multiple link failures, i.e., nodes are also treated as individual components of a network.

- End-system failures: These failures affect the actual end-user systems generating data for transmission and reception, i.e., hardware/software components at the end of the communication routes (sender and receiver equipment). These faults can be caused by failure of a single communication channel, failure of a transceiver equipment, or even software faults.

- Path failures: These failures are more general and can be caused by one or more of the above faults, i.e., if one or more link(s) or node(s) fail. In this case, the whole end-to-end communication path will be affected.

Meanwhile *survivability* is defined as the ability of the network to maintain an acceptable level of service after failure(s) within the network [3]. Now depending upon the type of failures, various recovery strategies have been proposed. For example, some approaches may treat link or nodes failures as isolated single events and compute sub-routes to avoid the affected entities. Meanwhile other approaches may treat whole path failures and table solutions to compute new "end-to-end" source to destination routes.

Based upon the above, the literature lists two broad classes of survivability/recovery schemes, delineated by their mode of operation, i.e., pre-failure and post-failure.

- Network protection: This is a pro-active (pre-provisioning) mechanism in which the backup routes are computed at setup, i.e., when the working route is computed. Protection can be applied at the individual link, sub-route, or at full end-to-end path levels. Overall, this strategy requires pre-reservation of resources on backup routes and hence entails higher resource consumption, i.e., since the reserved resources are kept idle until the occurrence of a failure. To address this concern, improved *shared* protection strategies have also been defined to allow working routes to share backup resources in case of single failure events [3]. Overall, the main advantage of the protection approach is that it provides very fast and guaranteed recovery, minimizing the loss of data due to recover delays.

- Network restoration: This is a reactive (post-provisioning) mechanism in which backup routes (or sub-routes) are computed and selected *after* the failure of a route or link. Unlike protection, restoration does not reserve backup resources. Clearly, this leads to a better resource usage and allows for supporting more user communication requests. However restoration has the disadvantage of increased recovery times (higher data losses) since it does active re-computation of backup routes after the detection of the failure. Also, a more serious concern is that these strategies may not be able to provide recovery under all circumstances, i.e., even for single link failures.

In light of the above, this thesis focuses on protection design strategies for multiple failure scenarios. A more detailed survey of related schemes in this particular subarea is now presented.

2.2 Dual-Link Failure Survivability Schemes

Over the years several studies have looked at *dual* (independent) link failure recovery in networks. One of the key challenges here is to ensure that such failures do not impact *both* the working and the backup paths of a routed demand, as shown in Figure 2.1. For example [10] motivates the problem for dual independent failures (i.e., occurring in an unspecified arbitrary order) and details some real-world scenarios where these type of faults may occur. Subsequently, three solutions are proposed to re-route around failed links using pre-assigned capacity with the assumption that the network graph is also 3-connected, i.e., Menger's Theorem [11]. In particular, the first two methods require identification of the failed links and pre-compute two link-disjoint backup paths, $b_1(e)$ and $b_2(e)$, for each link in a working path, $w(e)$. Upon occurrence of the first link failure e, $b_1(e)$ is then used to reroute the traffic. Now a double-link failure scenarios appears when a second link f may also fail after link e fails. Here four possible cases are identified, depending upon the location of the second failure. Meanwhile the third solution only pre-computes a single backup path, $b(e)$, for each link e by using various algorithms, see [10] for details. Overall, these schemes are analyzed for three different network topologies, including ARPANET, *New Jersey LATA* (NJLATA), and a large national backbone. Results indicate that it is possible to achieve 100% recovery for dual-link failures with a modest increase in backup capacity. In particular, the first two methods show lower capacity usage and yield shorter hop-lengths, but also require considerably higher signaling overheads (and therefore longer restoration times). Meanwhile the third method yields faster recovery timescales.

Another dual failure recovery scheme is also proposed in [12], by combining both pre-fault protection and post-fault restorations strategies. Namely, protection is used to reserve backup capacity so as to ensure that the *majority* of affected (failed) demands can be restored via pre-computed backup routes upon the occurrence of

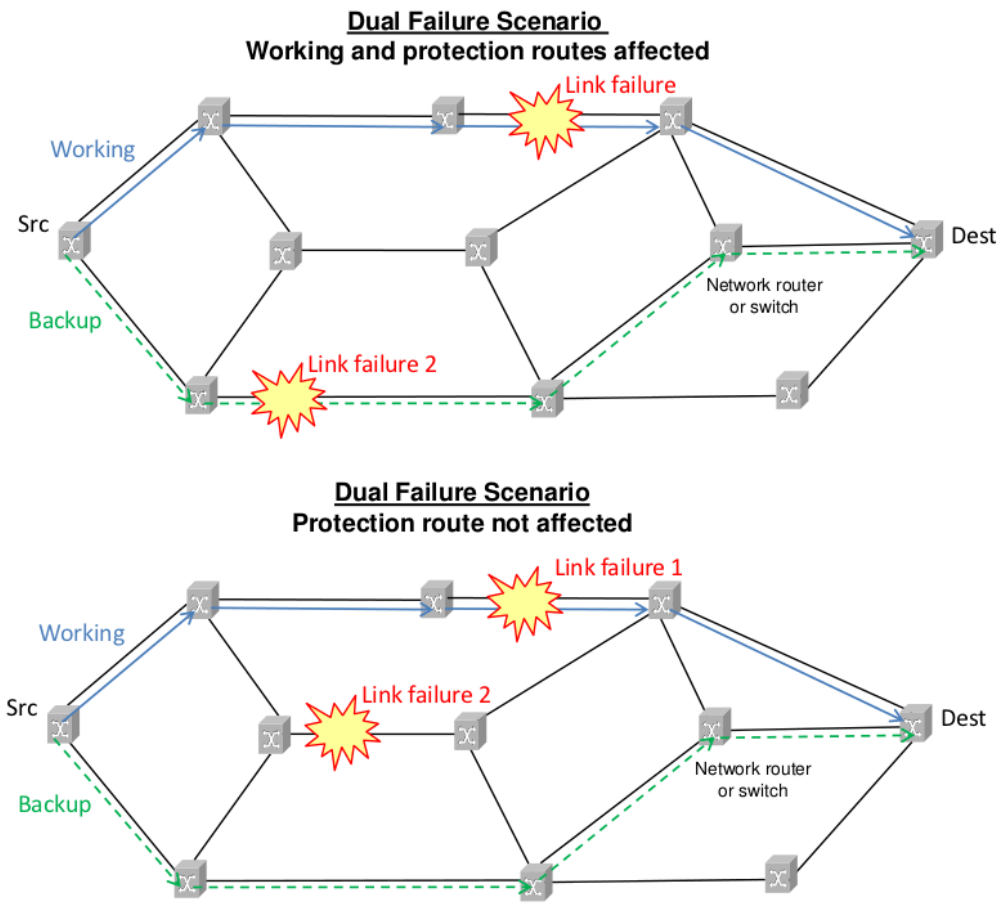

Figure 2.1: Scenarios of dual-link failure

dual-link failures. Meanwhile the restoration component is dynamically invoked to search for new backup routes for the remaining (unrestored) demands. Overall, the proposed protection algorithm uses Bhandari's scheme [13] to compute a working/backup path pair. Now, in case of a dual-link failure the scheme switches all affected demands into their respective backup routes. However, if the dual failure affects *both* the working and backup routes, then the scheme invokes dynamic post-fault restoration to try to re-compute a second backup route. This solution is tested using a network with 47 nodes and 98 links, and by varying demands from 500 to

2,500 in steps of 500. Overall, findings show that pre-planned protection can achieve 99.5% recovery of all demands against dual-link failures. Meanwhile the remaining 0.5% of demands can also be restored using dynamic restoration, i.e., close to 100% effectiveness.

Finally, [14] proposes a restoration-based scheme for dual-failures. Namely, this solution first tries to achieve full 100% *protection* restorability under single-link failures, while also maximizing coverage against any further dual-link failures. Specifically, when the first failure occurs, the restoration model attempts to dynamically compute a backup end-to-end link-disjoint route. This strategy basically uses sub-graph routing techniques to try to improve dual-link survivability, i.e., done by creating sub-graph of the base topology, each defined by removing a given link. Hence the scheme only accepts a request if it can be routed on all existing sub-graphs. Furthermore after the first link failure, all failed requests are switched to their respective backup routes. Following this, the scheme also evaluates if these demands can be routed in all of the subsequent sub-graphs of the network without the failed link. As a result, connections that can be successfully re-routed in all of these remaining sub-graphs will have full survivability against all dual-link failures as well. Conversely, connections that cannot be re-routed are only guaranteed survivability against the initial single-link failure. The overall scheme is tested for three different topologies, including the 14-node and 23-link NSFNET network, the 11-node and 22-link NJLATA network, and a standard 9-node and 18-link 3x3 mesh torus. Overall results show that the sub-graph routing technique can achieve complete dual-failure coverage for about 70 - 80% of all demands, i.e., in addition to full single-link failure protection.

2.3 Multi-Link Failure Survivability Schemes

Multi-link failure scenarios present a generalization of dual-link failures and usually pose much more complexity for network designers. In particular, these scenarios can arise during large-scale "catastrophic" events such as natural disasters (earthquakes, hurricanes, floods, etc), massive power outages, and malicious WMD attacks. Now a key challenge here is the fact that multiple failures make it very difficult (if not impossible) to pre-provision protection routes to recover from all possible link fault combinations, i.e., guarantee recovery. Hence, the use of *probabilistic* protection schemes and/or post-fault restoration is very appealing here. Along these lines some recent efforts that have tried to address these broader challenges and related strategies are now detailed.

One of the first studies on multi-failure recovery in (optical) networks is presented in [7]. Here the authors first try to compute maximally survivable end-to-end routes for *independent* link failure probabilities and define a detailed optimization problem, i.e., *maximum survivability under multiple failures* (MSMF) problem. This formulation tries to ensure that there are no overlaps between the backup sub-routes, but this problem is shown to be NP-complete. As a result the authors propose a modified heuristic scheme based upon Suurballe's disjoint path routing algorithm [7], denoted as *maximum survivability protection algorithm* (MSPA). This formulation assigns weights to each link, as defined by its respective failure probabilities, and then searches for two disjoint paths with the minimum total weight, see details in [7]. The scheme is tested using the 24-node NSFNET network, with an average node degree of 3.4, as well as a 14-node network with trap topology configuration and average node degree of 2.4. Tests are done for uniformly distributed link failure probabilities, and the average connections failure probability is used to evaluate the scheme. Overall, the MSPA heuristic compares well with some *integer linear program* (ILP) bounds (albeit not for the MSMF problem, which is not solvable). The

scheme also does well with trap topologies owing to its use of Suurballe algorithm.

Meanwhile, [15] presents another recovery scheme for large-regional failures. This effort is more realistic than that in [7], as it assumes dependency, i.e., *correlation*, between multiple failures. Namely, in most disasters scenarios, link/node faults are expected to have a very high degree of spatial, i.e. geographic, and temporal correlation [16]. Hence a failure model is defined by specifying an epicenter for each regional stressor along with an "impact range", given by R. Here it is assumed that any network component within this range will fail if the particular event transpires, i.e., each node and all of its attached links. Using this framework, a dynamic route restoration mechanism is utilized, although details on this algorithm are no presented in [15]. The scheme is then tested using the European COST-239 backbone with 11 nodes and 26 links, i.e., average nodal degree of 4.3 [17]. In addition, three different R ranges are evaluated (50, 150, and 200 km) by ensuring that no network disconnection occurs. Finally, the connection demands are varied up to a maximum value, D_{max}. Results for R values of 50 and 150 km show up to 15% traffic loss, and this figure increases to 28% for 250 km. Moreover, the findings also show that the average increase in network capacity requirements is constant for different request sizes, i.e., D_{max}, concluding that on average a 4% increase in network capacity is sufficient to re-route the remaining demands after a failure event.

Finally [6] presents a very comprehensive study of diverse route protection under probabilistic failure events. Here the concept of a *share risk link group* (SRLG) [18] is generalized to introduce a stochastic/random element, i.e., *probabilistic*-SRLG (p-SRLG). In particular, a p-SRLG is defined as a set of links which can fail with a non-zero failure probability given the occurrence of the SRLG event. However, it is important to note that this model assumes that all links within a given SRLG fail in an *independent* manner, i.e., no correlation. Leveraging this framework, the authors consider the case of computing two link-disjoint working/backup routes with min-

imum joint failure probability under the assumption that only one "catastrophic" SRLG event will occur at a given time, i.e., mutually-exclusive. This problem is shown to be NP-compete and a reduced ILP model is then developed along with a greedy heuristic. In particular, the latter scheme applies a two-step graph-theoretic approach to compute working/backup path routes. First, a modified graph is generated with link weights equal to the link failure probabilities averaged across all possible SRLG events. The shortest route on this graph is then chosen using Dijkstra's algorithm [19] and set as the the primary (working) route. Next, the primary route links are pruned and the remaining link weights are re-adjusted to compute a second path, i.e., backup route, again utilizing Dijkstra's algorithm. In particular, the link weights are re-assigned with the objective of reducing the joint path failure probability, see [6] for details. Note that the authors also define another more basic heuristic which does not modify the link weights after pruning the working route, i.e., termed as *shortest disjoint path* (SDP) scheme. These algorithms are then analyzed for randomly-generated network topologies with different node counts, where each topology has minimal 3-connectivity and node degree under 5. Most notably, the ILP defined above is also solved here for amenable network sizes using the CPLEX toolkit. Tests are done for 20 SRLG events and the overall findings show that the ILP scheme always yields the most reliable path pair, i.e., lowest joint path failure probability. However the greedy heuristic is also very close and consistently outperforms the SDP scheme.

2.4 Open Challenges in Network Survivability

Although the above-detailed studies on multi-failure recovery represent some key contributions in the field, there are still many open challenges. Foremost, the proposed schemes only focus on risk mitigation concerns, to the exclusion of other issues,

particularly TE resource efficiency. Hence there is a very high likelihood that these algorithms will yield overly lengthy (working, backup) routes as they implement detours around higher risk failure regions. In turn, this will increase resource consumption and lead to higher blocking rates and lower revenues for network operators. In addition, none of the proposed schemes have been tested for post-fault failure recovery, i.e., performance in terms of failed/survived connections. Instead, researchers have only measured the aggregate failure risks of the computed paths [6], which only provides a loose indicator of the relative protection capabilities, or analyzed the increased capacity requirements of re-routing once a failure event occurs [15]. Finally, the use of proactive post-fault restoration schemes has not been considered for multi-failure settings either, i.e., to try recover failed connection routes after a large-scale event. Indeed, restoration-based schemes can still provide a very credible "last-line-of-defense" against multiple failures, depending upon network load conditions. In light of the above, there is an urgent need to expand the work in the area of multi-failure recovery by addressing both survivability and TE resource efficiency concerns. A detailed heuristic solution is now presented to address these challenges.

Carefully note that some researchers have also studied the impact of multiple failures on physical networking topologies. Specifically, the aim here is to analyze the effect of these stressors on key metrics such as network connectivity or resilience. Overall, such analysis can provide some key insights for longer-term network planning and upgrade activities. Nevertheless these strategies are not really applicable to dynamic connection recovery scenarios, as is the focus of this thesis. Interested readers are referred to [16], [20] and [21] for more details.

Chapter 3

Joint Path Pair Survivability Scheme

A novel multi-failure path protection scheme is now proposed to jointly incorporate both resource efficiency and risk mitigation objectives. The solution embodies a heuristic strategy using graph-theoretic algorithms and is highly-amenable to real-world implementation. Now in order to broaden the scope of this effort, two different failure models are incorporated here, including the one presented recently in [6] as well as a more realistic rendition that incorporates failure correlation within a region. Full details on these failure models and path protection schemes are now presented, along with computational complexity analyses.

3.1 Notational Overview

Before presenting further details, the necessary notations are first introduced. In particular, consider a network topology operating in a generalized multi-stressor environment. This network can be modeled as a graph, $\mathbf{G}(\mathbf{V}, \mathbf{L})$, where $\mathbf{V} = \{v_1, v_2, \dots\}$

represents the set of switching/routing nodes and $\mathbf{L} = \{l_{ij}\}$ represents the set of bandwidth capacity links. In particular, l_{ij} represents a bi-directional link between nodes v_i and v_j with maximum capacity c_{ij}. The available free bandwidth on link l_{ij} is also given by b_{ij}. Furthermore, all incoming connection requests arrive in a random "on-line" manner and comprise of three key parameters, i.e., the source node, v_a, the destination node, v_b, and a desired capacity level, c. Hence a request can be summarized by the 3-tuple $\{v_a, v_b, c\}$.

Meanwhile the threat environment is assumed to comprise of N potential stressor events denoted by the set $\mathbf{F} = \{f_1, f_2, \ldots, f_N\}$. In particular, each event $f_i \in \mathbf{F}$ represents a SRLG and is comprised of a set of vulnerable links, $\mathbf{X}_i$, within a certain geographic region/proximity. Now given the severity of such events, it is reasonable to assume that only one particular stressor may transpire at a given time, i.e., mutually-exclusive stressors, as also assumed in [6]. This assumption precludes the need for more complicated inter-stressor dependency modeling and allows one to define a relative occurrence probability for each particular event, ϕ^i, where

$$\sum_{i=0}^{N} \phi^i = 1.$$

Carefully note that the SRLG is only characterized via a set of links, $\mathbf{X}_i$, and not nodes. However this representation is still sufficiently generic since susceptible nodes can also be captured by including all of their emanating links.

Now even though individual stressor events are treated as mutually-exclusive here, the same cannot be assumed of link (or node) failures within a SRLG region upon occurrence of a stressor. For example, in many cases link failure probabilities will be heavily-dependent on the proximity of the link from the epicenter of the stressor event. Along these lines, two different *probabilistic*-SRLG (p-SRLG) models are used

to represent link (node) vulnerability within a given region. These are now detailed further.

3.2 Multi-Failure Models

In general, pre-provisioned protection recovery schemes can benefit tremendously from prior stochastic knowledge of multi-failure stressor events. It is here that probabilistic models are of crucial importance, as they can help define the location and relationship between multiple failures. Further detailed are now presented.

3.2.1 Independent Intra-Region Failure

The first p-SRLG model is the same as that used in [6] and assumes fully-independent link failures within a given region, i.e., no link failure correlation. Namely, conditional to the occurrence of a particular event f_i, each link in the associated failure region, $l_{jk} \in \mathbf{X}_i$, is assumed to fail with probability p^i_{jk}, which is independent of the failure probabilities of other links in $\mathbf{X}_i$. In addition it is also assumed that all links $l_{jk} \notin \mathbf{X}_i$ have zero failure probability. This model is further illustrated in Figure 3.1 which depicts two potential stressors (SRLG regions) and their respective link failure probabilities.

Note that the independent link failure assumption simplifies network analysis as it allows one to use a simple dot product to compute the *conditional* failure probability of an end-to-end route (to a given stressor). Namely, the probability that a routed path $\mathbf{R}$ will fail upon occurrence of stressor f_i is given by:

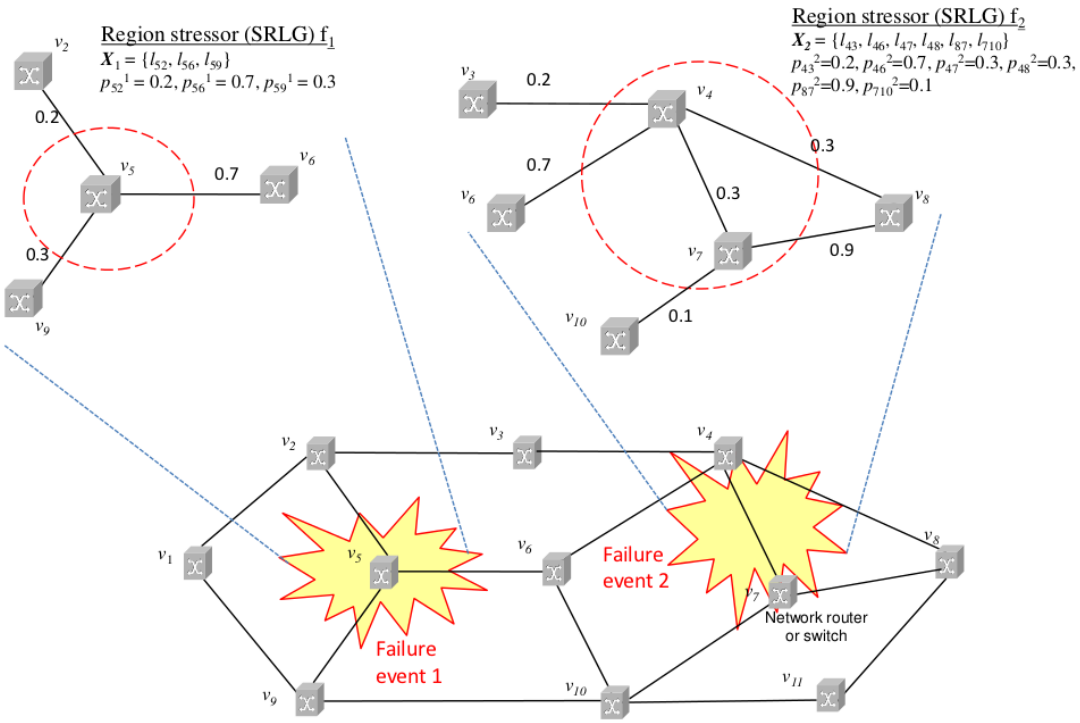

Figure 3.1: Representation of independent intra-region failures

$$\text{Prob(Route failure} \mid f_i) = 1 - \prod_{\forall\ l_{jk}\ \in\ \mathbf{R}} (1 - p^i_{l_{jk}}). \tag{3.1}$$

This formulation is leveraged further in Sections 3.3.2 and 3.3.3.

3.2.2 Dependent Intra-Region Failures

The p-SRLG model in Subsection 3.2.1 assumes no dependencies between link (node) failures within a common stressor region. However associated link failure probabilities are expected to vary according to their proximity from the epicenter of a stressor

event. Hence link failures within a SRLG region will likely exhibit a high degree of geographical correlation, effectively precluding independent failure assumptions as in [6]. To better reflected on this reality, and improved probabilistic model is proposed here based upon a two-dimensional Gaussian distribution. Specifically, a concentric *probability density function* (pdf) is defined for each SRLG event f_i as follows:

$$f(x) = \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-\frac{x^2}{2\sigma_i^2}}, \tag{3.2}$$

where x is the distance from the epicenter of f_i and σ_i is the standard deviation (radii) of the stressor. Using this definition, the conditional failure probability of a link l_{jk} in the share risk link group (SRLG) region is given as:

$$p_{l_{jk}}^i = \begin{cases} 1 - erf\left(\frac{d_{jk}^i}{\sqrt{2}\sigma_i}\right) & \text{if } l_{jk} \in \mathbf{X}_i \\ 0 & \text{otherwise,} \end{cases} \tag{3.3}$$

where $erf(\cdot)$ is the standard error function and d_{jk}^i is the closest distance of any point on link l_{jk} from the epicenter of f_i. Hence this implies that links closer to the epicenter will have higher failure routes, i.e., $p_{l_{jk}}^i \to 1$ as $d_{jk}^i \to 0$. Carefully note that the actual epicenter locations and radii (σ_i values) can be determined by conducting off-line threat analysis based upon a wide range of inputs, e.g., such as geographical locations, weather patterns, geopolitical constraints, etc. However, these broader considerations are clearly out of scope of the work herein and not considered further. A sample illustration of this dependent failure model is also shown in Figure 3.2 for a stressor event centered between two nodes.

Carefully note that link failure dependencies preclude the use of dot-product type computations of end-to-end route failure probabilities. Instead, the conditional route

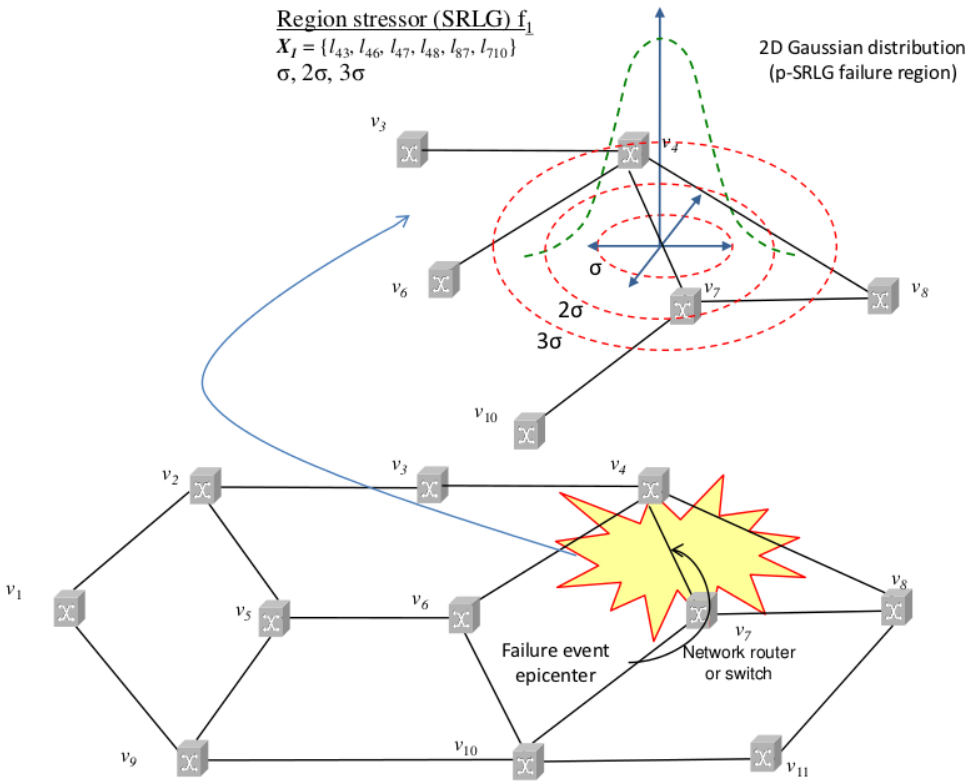

Figure 3.2: Representation of dependent intra-region failures

failure probability (to a given SRLG stressor event f_i) is now given by the maximum failure probability of any of the links along the path route R, i.e.,:

$$\text{Prob}(\text{Route failure} \mid f_i) = \max_{\forall\, l_{jk} \,\in\, \mathbf{R}} \{p^i_{l_{jk}}\}. \tag{3.4}$$

This formulation is also leveraged further in Section 3.3.3.

3.3 Path Protection Schemes

As noted earlier, the focus of this thesis is to improve recovery in stochastic multi-failure settings. In particular, end-to-end protection is used to route link-disjoint backup paths and provide rapid switchover recovery after failure events. However before detailing the proposed solution, some baseline algorithms are first detailed. These definitions are then leveraged in the new approach and also used for subsequent comparison analyses (Chapter 5).

3.3.1 Pure Traffic Engineering Strategies

A wide range of TE-based heuristics have been developed for network connections routing [22], [23], and these schemes can also be adapted and applied for computing diverse path pair routes. By and large, most solutions here pursue one of two objectives, *resource minimization* or *load balancing*. Along these lines, two baseline *traffic engineering* (TE) schemes are proposed here using *K-shortest path* (SP) routing [24] strategies:

Hop count routing (resource minimization): One of the most basic TE strategies is to pursue resource minimization by choosing routes with the lowest hop count. Now for the case of a single working route, this can be achieved by simply running Dijkstra's shortest path algorithm with unity link weights. However resource minimization for working/backup path pairs is slightly more complicated. For example, simply using a greedy "two-step" approach to first select the shortest working path (via Dijkstra's algorithm) followed by its shortest link-disjoint protection path may not yield the lowest aggregate hop count. This issue has been addressed earlier, and a modified K-SP approach has been pro-

posed to select from multiple link-disjoint path pair combinations [25]. Hence this "joint" selection approach is also re-used here, see Figure 3.3 for psuedocode listing. Specifically, a K-SP algorithm is first run to compute multiple working routes between v_a and v_b in the network graph $\mathbf{G}(\mathbf{V}, \mathbf{L})$, denoted by $\{\mathbf{wr}_i\}$, $1 \leq i \leq K$. Next, the associated protection routes, $\{\mathbf{br}_i\}$, for each of these working paths are computed by pruning their links and running Dijkstra's shortest-path algorithm. Based upon this, the working protection pair $\{\mathbf{wr}_i, \mathbf{br}_i\}$ with the minimum aggregate hop count is chosen, see Figure 3.3. Carefully note that the above computations are only done over *feasible* network links in order to help lower blocking probabilities, i.e., links with sufficient available capacity to provision the request. Namely the link weight for link l_{ij} in $\mathbf{G}(\mathbf{V}, \mathbf{L})$ is set to unity according to the following expression:

$$w_{ij} = \begin{cases} 1 & \text{if } b_{ij} \ \geq \ r \\ \infty & \text{otherwise.} \end{cases} \tag{3.5}$$

Although hop count routing can yield the most resource-efficient routes, it tends to overload certain network links under higher load situations, leading to increased request blocking rates, see [26].

Load balancing routing: This TE strategy improves upon hop count routing and tries to achieve more balanced traffic distribution across network links. In particular, this is done by simply adjusting the weights for feasible network links in a *dynamic* manner according to their congestion levels, i.e., and still re-using the same overall K-SP path pair selection approach detailed in Figure 3.3. Namely, an inverse-weighting approach is used here as follows:

$$w_{ij} = \begin{cases} \frac{c_{ij}}{c_{ij} - b_{ij} - \epsilon} & \text{if } b_{ij} \geq r \\ \infty & \text{otherwise,} \end{cases} \tag{3.6}$$

where c_{ij} represents the total capacity of link l_{ij}, b_{ij} represents the current available free capacity, and ϵ is a small value chosen to avoid division errors. Hence this algorithm chooses the working/backup path pair $\{\mathbf{wr}_i, \mathbf{br}_i\}$ with the minimum aggregate load-balancing cost. Overall, dynamic link weighting schemes have been shown to give notably better performances in terms of reduced blocking probabilities and higher network loads (revenues), see [26].

1. Given incoming request $\{v_a, v_b, r\}$.
2. Prune non-feasible links ($b_{ij} < r$) and generate $\mathbf{G}'(\mathbf{V}, \mathbf{L})$.
3. Assign link weights w_{ij} according to TE policies, i.e., hop count (Equation 3.5) or load balancing (Equation 3.6).
4. Compute K-shortest working paths $\{\mathbf{wr}_i\}$ from v_a to v_b in $\mathbf{G}'(\mathbf{V}, \mathbf{L})$.
5. For each $\{\mathbf{wr}_i\}$ prune links in $\mathbf{G}'(\mathbf{V}, \mathbf{L})$ and compute its link-disjoint backup path $\{\mathbf{br}_i\}$.
6. Select path pair with minimum total TE cost, i.e., $min(\sum_{l_{ij} \in \mathbf{wr}_m} w_{ij} + \sum_{l_{ij} \in \mathbf{br}_m} w_{ij})$, $\forall\ m \in [1, \ldots K]$.

Figure 3.3: Pseudocode for TE heuristic strategy: hop count and load balancing

3.3.2 Pure Risk Minimization Strategies.

As noted in Chapter 2, some recent studies in [6] and [15] have also proposed more specialized routing schemes to minimize the aggregate risk of working/backup path pairs. For example [6] defines a basic *shortest disjoint path* (SDP) scheme using a two-step Dijkstra's computation, as noted in Section 2.3, see Figure 3.4. This

algorithm follows a greedy approach where it first computes a working route **wr** over $\mathbf{G}(\mathbf{V}, \mathbf{L})$ with link weights set to the average link failure probabilities across all possible stressor events:

$$w_{ij} = \sum_{f_i \in \mathbf{F}} \phi^i p^i_{l_{ij}}, \forall\ l_{ij} \ \in \mathbf{L}. \tag{3.7}$$

After a working route is found for an incoming request, the algorithm then prunes all of its links on $\mathbf{G}(\mathbf{V}, \mathbf{L})$ and tries to compute a backup protection route, **br**, by utilizing same link weights from Equation 3.7. Next, [6] also presents an improved risk mitigation heuristic, termed here as *risk minimization* (RM), by deriving an approximation to an optimization problem. Akin to the *shortest disjoint path* (SDP) algorithm, this scheme also implements a greedy two-step approach but now performs link weight *re-adjustment* after computing the working route, **wr**. Namely, the working route links in $\mathbf{G}(\mathbf{V}, \mathbf{L})$ are first pruned and then the remaining link weights are re-computed, as shown in Figure 3.5.

1. Given incoming request $\{v_a, v_b, r\}$.
2. Prune non-feasible links ($b_{ij} < r$) and generate $\mathbf{G}'(\mathbf{V}, \mathbf{L})$.
3. Set link weights in $\mathbf{G}(\mathbf{V}, \mathbf{L})$, i.e., $w_{ij} = \sum_{f_i \in \mathbf{F}} \phi^i p^i_{l_{ij}}, \forall\ l_{ij} \ \in \mathbf{L}$.
4. Find shortest path, i.e., **wr**.
5. Prune all links in **wr** and create graph $\mathbf{G}'(\mathbf{V}, \mathbf{L})$.
6. Find shortest path, i.e., **br**.

Figure 3.4: Pseudocode of shortest disjoint path (SDP) heuristic strategy

3.3.3 Joint Path-Pair (JPP) Scheme

In general, TE routing schemes will likely yield higher failure rates at they do not incorporate a-priori risk state from p-SRLG models. Similarly, risk minimization

1. Given incoming request $\{v_a, v_b, r\}$.
2. Prune non-feasible links ($b_{ij} < r$) and generate $\mathbf{G}'(\mathbf{V}, \mathbf{L})$.
3. Set link weights in $\mathbf{G}(\mathbf{V}, \mathbf{L})$, i.e., $w_{ij} = \sum_{f_i \in \mathbf{F}} \phi^i p^i_{l_{ij}}, \forall\ l_{ij} \in \mathbf{L}$.
4. Find shortest path, i.e., **wr**.
5. Prune all links in **wr** and create graph $\mathbf{G}'(\mathbf{V}, \mathbf{L})$.
6. Set remaining link weights $w_{ij} = \sum_{(k,l)} wp_{l_{kl}} \sum_{f_i \in \mathbf{F}} \phi^i p^i_{l_{ij}} p^i_{l_{kl}}, \forall\ l_{ij} \in \mathbf{L}$.
7. Find shortest path, i.e., **br**.

Figure 3.5: Pseudocode of risk minimization (RM) heuristic strategy

schemes will likely yield higher resource usages as they try to achieve path detours around high risk regions. Hence the proposed joint strategy here tries to achieve a better balance by incorporating both objectives in the path pair selection process, as shown in Figure 3.6. In particular, as is done for the TE strategies, this algorithm first computes a set of working/backup path pairs $\{\mathbf{wr}_i, \mathbf{br}_i\}$ over the feasible network links by running the K-SP algorithm and K instances of Dijkstra's shortest path algorithm between the source and destination nodes, v_a and v_b. Note that path selection can either be done using hop count (Equation 3.5) or load balancing (Equation 3.6) link weighting. As such this yields two renditions of the scheme, JPP-HC and JPP-LB, respectively. Next, upon path-pair computation, the scheme assigns a route failure probability for each of the computed routes, as explained in Section 3.2. Finally, it chooses the path pair with the minimum *joint path failure probability*, computed as a vector dot product of the path failure probabilities across all failure events F, i.e., yielding a measure of correlation.

3.4 Complexity Analysis

It is also very important to classify the computational, i.e., run-time, complexity of the proposed schemes, as this will impact implementation in real-world network

1. Define incoming request $\{v_a, v_b, r\}$.
2. Prune non-feasible links ($b_{ij} < r$) and generate $\mathbf{G}'(\mathbf{V}, \mathbf{L})$.
3. Assign link weights w_{ij} according to TE policies, i.e., hop count (Equation 3.5) or load balancing (Equation 3.6).
4. Compute K-shortest working paths $\{\mathbf{wr}_i\}$ from v_a to v_b in $\mathbf{G}'(\mathbf{V}, \mathbf{L})$.
5. For each $\{\mathbf{wr}_i\}$ prune links in $\mathbf{G}'(\mathbf{V}, \mathbf{L})$ and compute its link-disjoint backup path $\{\mathbf{wr}_i\}$.
6. Assign route failure probabilities for each $\{\mathbf{wr}_i\}$ and $\{\mathbf{br}_i\}$, $\forall\ i\ \in [1, \ldots K]$ according to failure model, i.e., independent (Equation 3.1) or dependent (Equation 3.4).
7. Select path pair with minimum joint path failure probability, i.e., $min_{\forall\ i\ \in\ [1,\ldots K]}(f_{\mathbf{wr}_i} \cdot f_{\mathbf{br}_i})$.

Figure 3.6: Pseudocode for proposed joint path pair (JPP) heuristic strategy

provisioning systems. Now in general, the TE and joint provisioning schemes in Sections 3.3.1 and 3.3.3 are expected to have higher complexity as they use more complex K-shortest path (SP) algorithms. Consider the details here:

- Traffic engineering strategies: These schemes are detailed in Figure 3.3 and basically implement one K-SP computation followed by up to K Dijikstras shortest path computations. Now the overall run-time complexity of optimized variants of the K-SP algorithm is $O(K(|\mathbf{L}| + |\mathbf{V}|log|\mathbf{V}|))$, where $|\mathbf{V}|$ is the number of nodes and $|\mathbf{L}|$ is the number of links in network graph $\mathbf{G}(\mathbf{V}, \mathbf{L})$ [27]. Meanwhile Dijkstra's shortest path algorithm has $O(|\mathbf{V}|log|\mathbf{V}|)$ complexity. Clearly the former term is the dominant component here, yielding an overall computational complexity bound of $O(K(|\mathbf{L}| + |\mathbf{V}|log|\mathbf{V}|))$.

- Risk minimization strategies: These greedy schemes are detailed in Figures 3.4 and 3.5 and basically implement two Dijkstra's shortest path computations, i.e., $O(|\mathbf{V}|log|\mathbf{V}|)$ complexity. However, appropriate probabilistic link weights must also be computed here by looping over all possible stressor events, i.e.,

$O(N|\mathbf{L}|)$. As the number of nodes or stressors may vary under generalized scenarios, the aggregate computational complexity for these schemes is given as a sum of these two expressions, i.e., $O(|\mathbf{V}|log|\mathbf{V}| + N|\mathbf{L}|)$.

- Joint strategies: This joint path-pair scheme is shown in Figure 3.6 and essentially runs the same set of graph-theoretical computations as the TE heuristic in Figure 3.3. In particular, only the only difference here is in the choice of the final path pair after the K sets are computed. As this selection process is of minimal complexity compared to the graph-based operations, this scheme also has $O(K(|\mathbf{L}| + |\mathbf{V}|log|\mathbf{V}|))$ complexity.

Chapter 4

Simulation Design

The performance of the various multi-failure protection schemes in Chapter 3 is evaluated using *discrete event simulation* (DES). This method is a very powerful means of analyzing complex networking behaviors which cannot otherwise be modeled in a closed-form mathematical manner. In particular, DES simulates the operation of a system as a chronological sequence of events, triggered in response to a previous event. The network events here can be either connection requests, control messages, link failures, etc. Meanwhile the response to any of these events can also generate further new events and/or terminate or reschedule existing events.

Now a variety of DES tools are currently available, including *OPNET Modeler*™, *NS2/ NS3*, *OMNET++*, etc. All of these tools can be found in the public domain, with the exception of *OPNET Modeler*™, which is a commercial package that is available free of charge for university research only. Although each of these packages has its own benefits, this effort uses the *OPNET Modeler*™ tool owing to its complete set of features, i.e., *graphical user interface* (GUI) support and robust DES processing libraries. More importantly this tool also provides a full C/C++ interface for developing customized models. Overall, *OPNET Modeler*™ has gained very

strong traction with many users in the industrial and academic research sectors.

Given the choice of a simulation analysis tool, the next key step is to design some realistic network scenarios for evaluation purposes. Along these lines, this chapter details the network topology configurations used along with the specifies of the multi-failure stressors. The key performance evaluation metrics used to gauge the schemes are also presented here.

4.1 Network Topology and Failure Scenarios

Indeed, there are many different network topologies that can be used to evaluate protection schemes. However, in order to ensure realistic findings, it is very important to select examples which are reflective of real-world network designs, e.g., in terms of node count, links, connectivity, etc. In light of the above, a sample US IP backbone network is selected, see Figure 4.1, comprising of 24 nodes and 43 links. This network is also used in [6] and has an average node degree of 3.5 links per node, representative of the increased connectivity level of most commercial backbones. Furthermore, from Figure 4.1, it can be seen that the distances between nodes is generally less than a few hundred kilometers (assuming continental US overlay). Indeed, some nodes are much closer, implying that large scale stressors may impact multiple links.

Given that many of the schemes in Chapter 3 are focused on multi-failure recovery, it is also important to define a realistic set of regional stressor events. Clearly these definitions will have a direct impact upon the overall recovery performance. Now, as detailed in Chapter 3, this thesis utilizes two different failure models within a SRLG region: dependent and independent failures. As a result, two classes of failures are also defined for the network in Figure 4.1. Namely the first specifies $N = 8$ regionalized stressors events, whereas the second defines $N = 5$ events. Both models assume that the stressors are mutually-exclusive and are chosen via a random uniform

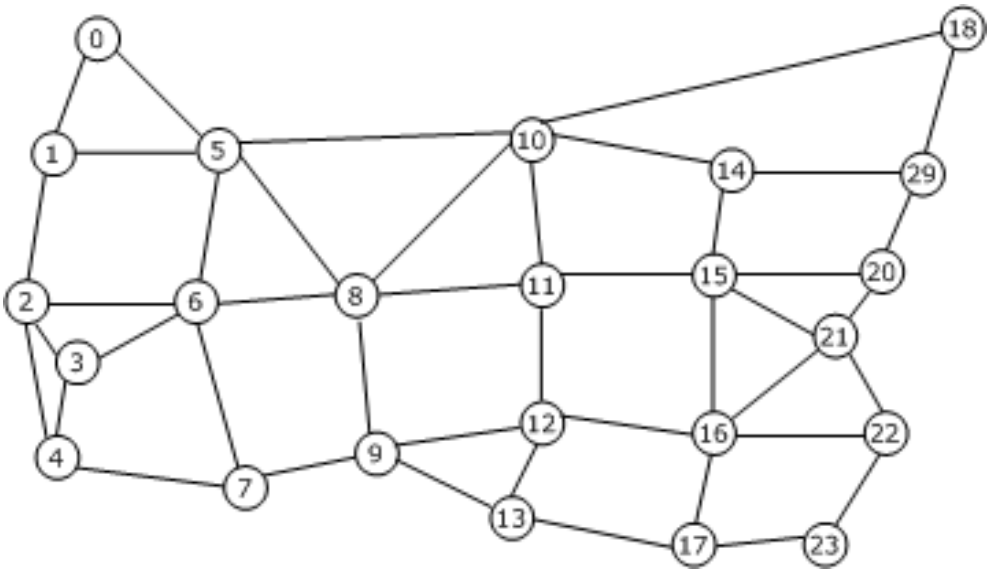

Figure 4.1: US IP backbone

distribution. Further details are now presented.

1. Independent intra-region link failures: This scenario defines $N = 8$ different (SRLG) stressor events, with each characterized by a pre-defined static failure region. These regions may or may not be circular in nature. Here all links within a given stressor region are deemed to fail in an *independent* manner, as per the p-SRLG model in Section 3.2.1, and without loss of generality, regions are centered at network nodes, see Figure 4.2.

2. Dependent intra-region link failures: This scenario defines $N = 5$ failure regions with dependent intra-region link failures, as per the p-SRLG model described in Section 3.2.2. Namely, each region is characterized by a circular Gaussian failure distribution with varying radii (standard deviation σ), as shown in Figure 4.3. Specifically, this diagram shows up to 3σ of coverage, which effectively spans all potential links impacted by the given stressor.

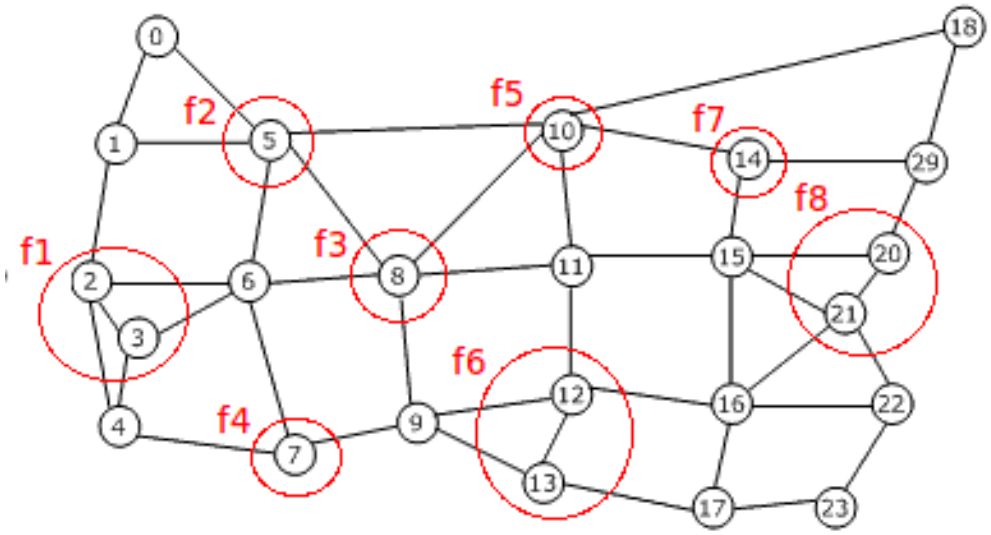

Figure 4.2: Independent failure scenario with $N = 8$

4.2 Performance Evaluation

In order to effectively gauge and quantify the performance of the proposed protection schemes, a variety of metrics are used. The overall goal here is to choose those measures which help provide a clear understanding of the scheme's behavior, and also help differentiate their relative performances.

1. Failure Rate (FR): This metric measures the recovery effectiveness of a scheme after a (multiple) failure event. Specifically, consider a network with M_{active} active connections immediately prior to a stressor occurrence. After this event transpires, up to W_{fail} working routes and P_{fail} protection routes can fail. However, depending upon the severity/scale of the failure event, some connections may experience *both* working and backup route failure, i.e., B_{fail} joint path pairs may also fail. Based upon these values, the FR is given by:

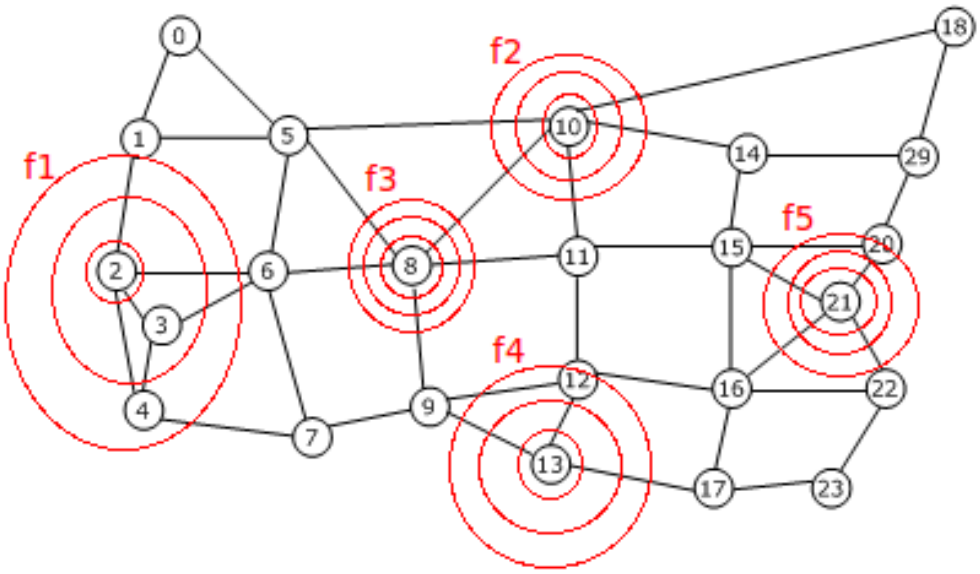

Figure 4.3: Dependent failure scenario with $N = 5$ and different radii (σ)

$$FR = \frac{B_{fail}}{M_{active}}, \tag{4.1}$$

i.e., the total number of affected demands is less than $W_{fail} + P_{fail}$. Carefully note that B_{fail} does not count any connections whose source and/or destination nodes fail. Clearly service cannot be restored in such cases using any form of path protection.

2. Protection Fail Rate (PFR): This metric measures the fraction of non-affected working demands whose protection routes are impacted by a failure event. It is important to measure this value as these working paths will essentially lose their backup recovery, i.e., as opposed to working paths which fail but still have unaffected (i.e., non-failed) backup routes. As per the notation used to define FR metric above, the PFR is given by:

$$PFR = \frac{P_{fail} - B_{fail}}{M_{active}}. \tag{4.2}$$

Carefully note that the above expression discounts the case of joint working/protection route failures. In addition, any connections with source and/or destination node failures are not included in the P_{fail} count, i.e., akin to B_{fail}.

3. Bandwidth Blocking Rate (*BBR*): This measure is used to quantify user request rejection rates. Specifically the BBR is computed as the ratio between the aggregate bandwidth requested but not provisioned to the total amount of requested bandwidth. Namely, consider a total of M connections requests arriving at the network, with the i-th request having capacity r_i. Hence the total requested bandwidth $B_{request}$ is given by:

$$B_{request} = \sum_{i=1}^{M} r_i.$$

Now further assume that P of these incoming connection requests cannot be provisioned by the setup phase. Hence the total rejected bandwidth B_{reject} is:

$$B_{reject} = \sum_{i=1}^{P} r_i$$

and, hence the resultant BBR is defined as:

$$BBR = \frac{B_{reject}}{B_{request}}. \tag{4.3}$$

4. Average Route Length (*ARL*): This last metric measures average per-connection resource utilization and is computed as the average hop count across

all successful path pairs, i.e., both working and backup routes. Specifically, consider a network with $M_{success}$ total successful connections, where W_{length} is the cumulative length of all working routes and B_{length} is the cumulative length of all backup routes. Hence the ARL is defined as:

$$ARL = \frac{W_{length} + B_{length}}{M_{success}}. \tag{4.4}$$

Overall, this metric delivers important information regarding the usage of the network resources. Namely, a larger average ARL may indicate inefficient resource utilization.

Chapter 5

Results and Analysis

The performance of the proposed joint routing scheme in Chapter 3 is now tested using specially-developed models using *OPNET Modeler*™. As detailed in Chapter 4, these tests are done using the network shown in Figure 4.1 and two different failure models are also treated. In addition, comparative evaluations are performed using the hop count and load balancing TE strategies detailed in Section 3.3.1, i.e., denoted here as TE-HC and TE-LB, respectively. Furthermore, the pure risk minimization schemes in [6] are also tested, albeit for the independent link failure scenario only, i.e., SDP as well as the improved RM approach. Finally, the proposed *joint path pair* (JPP) scheme can be applied using either hop count or load balancing link weight metrics, i.e., as per Section 3.3.3. However, this evaluation study only considers the latter variant, i.e., JPP-LB, since it is shown to yield improved blocking and recovery performance in some preliminary tests.

Overall, each simulation run is tested for $2,500,000$ random connections with scaled holding times of mean 600 seconds, i.e., exponentially-distributed. Meanwhile the actual network input loads are varied by changing the mean request *inter arrival time* (IAT) values, also distributed in an exponential manner. All network link

capacities are set to 10 Gb/s and user bandwidth requests, r, are uniformly varied between 1 Gb/s and 200 Mb/s in 200 Mb/s increments, i.e., to model fractional Ethernet demands. Furthermore, path pair computation for the TE-HC, TE-LB, and JPP-LB schemes is done using a value of $K = 7$. Finally all failure metrics are averaged by initiating 500 stressors events at random intervals in the simulation. The results for the two different failure models are now presented.

5.1 Independent Link Failure Scenarios

The various schemes are initially tested for independent link failure scenarios. In particular, these tests are done for both even (equiprobable) and non-even stressor event distributions. Consider the findings.

5.1.1 Even (Equiprobable) Stressor Event Distribution

The first set of tests assume that all stressors (SRLG regions) are equiprobable and occur with the same probability, i.e., $\phi_i = \frac{1}{8}$, as per 8 SRLG regions shown in Figure 4.2. To start out, simulation runs are done to measure the recovery effectiveness of the various schemes, averaged over 500 randomly-generated stressors (with each following the independent intra-region link failure model of Section 3.2.1). Results are shown in Figure 5.1, which plots the associated FR values (Equation 4.1) for varying input load regimes, i.e., both working and protection path failures. Here it is clear that the proposed JPP-LB and optimized risk-only RM schemes give the best reliability and consistently outperform the more basic SDP risk reduction approach, i.e., slightly over 1% connection failures at very high loads. Also, RM outperforms the JPP-LB solution at very low loads, yielding close to zero failures. By contrast the TE-based strategies yield much lower recovery, with close to twice the number

of failures versus the JPP-LB and RM solutions. Furthermore, the load balancing TE-LB strategy does slightly better than the resource minimization TE-HC scheme.

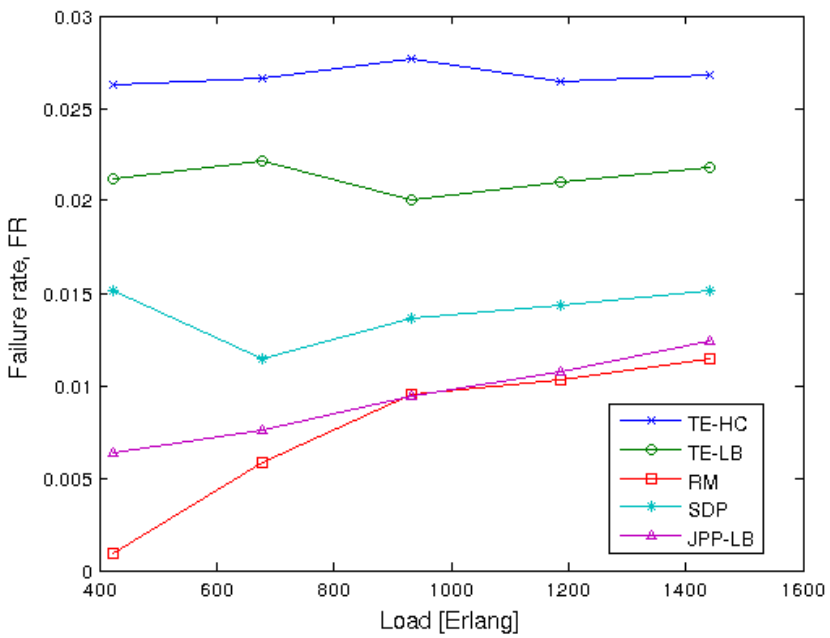

Figure 5.1: Failure rate for both working and protection routes (FR)

As noted in Section 4, it is also important to gauge the impact of multi-failure stressors on the reliability of protection (backup) routes for non-failed working connections. Along these lines, Figure 5.2 plots the PFR results (Equation 4.2) for the various schemes for different input load ranges and reveals some very interesting findings. Foremost, the advanced RM risk-based scheme gives the highest failure rates for backup routes, averaging almost three times more than those for TE-based strategies. By contrast the joint JPP-LB solution is much more effective here, and closely tracks the improved performance of the TE-based solutions. Meanwhile, even though the SDP risk-based solution does better than the RM variant, it still yields higher protection route failures than the JPP-LB scheme. Overall, these findings confirm that the JPP-LB scheme provides much better post-fault reliability for non-

failed connections. In turn this will simplify operational concerns for network carriers by reducing the need to re-compute and re-provision failed protection routes after massive network disruptions.

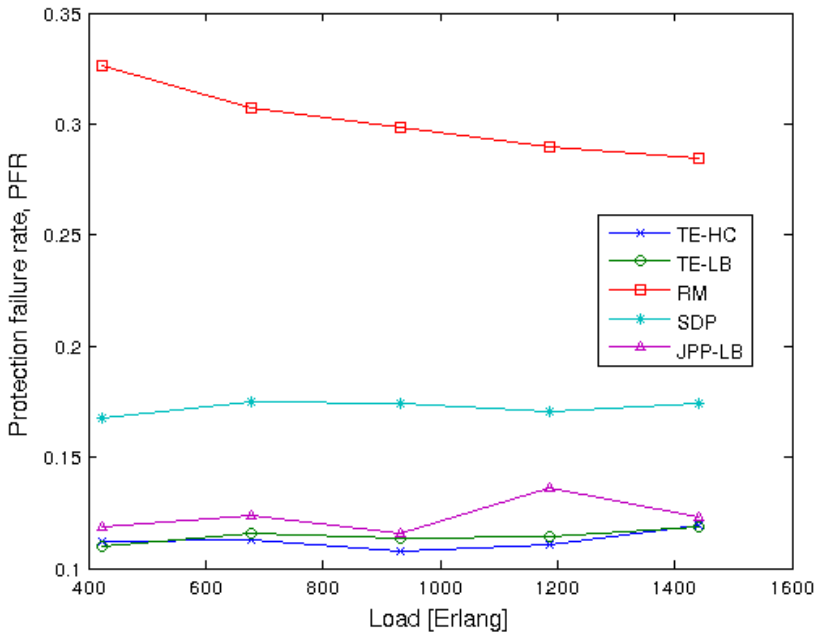

Figure 5.2: Failure rate for protection routes (PFR)

The results in Figures 5.1 and 5.2 strictly focus on connection reliability analysis. Hence additional metrics are now used to gauge overall TE performance as well. In particular, the average blocking performances of the four schemes (BBR values, Equation 4.3) are plotted in Figure 5.3 using a log-scale to cover a wide range of operational input load regimes. These results clearly show major shortcomings with the pure risk-based strategies in [6], with neither scheme (RM, SDP) giving under 1% blocking for almost all input traffic loads. By contrast, the pure TE and proposed JPP-LB algorithms are much more effective here, yielding many orders of magnitude lower demand blocking at low-to-medium load ranges. In particular, the JPP-LB scheme is very competitive here as it consistently outperforms the TE-HC heuristic

and closely tracks the TE-LB scheme (which gives the lowest blocking). As per these findings, it is likely that the SDP and RM schemes may see limited deployment in real-world settings as their increased blocking rates will yield much lower levels of carried load, i.e., reduced revenues, for network carriers. This is a key finding of this work.

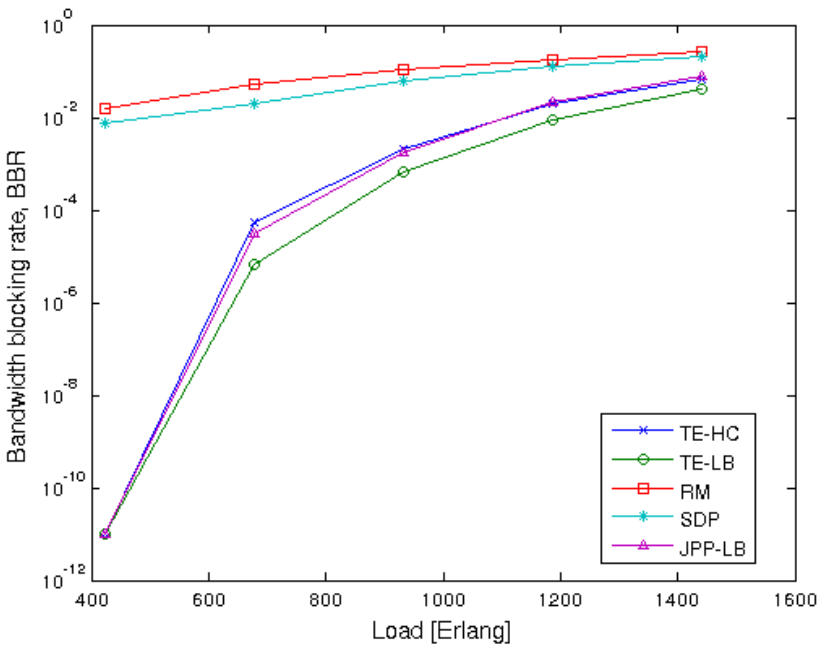

Figure 5.3: User request blocking rate (BBR)

Finally, Figure 5.4 plots the average resource utilization of the working/protection path pairs using the ARL metric (Equation 4.4). As expected, the TE-HC scheme gives the lowest usage and closely followed by the TE-LB algorithm. By contrast, the risk-based strategies are much more inefficient as they generate longer route pairs to detour around higher-risk regions. For example, the RM scheme in [6] gives almost two times longer routes than the TE-LB solution. Meanwhile the JPP-LB solution achieves a very good balance here, with average route lengths much closer to the TE-based schemes, i.e., about 15% higher usages. Overall, these findings also

corroborate the relative blocking rate performances for the schemes, as observed in Figure 5.3. Finally, it is also noted that higher input loads lead to a slight increase in resource utilization, i.e., as higher blocking values lead to longer route selection for all schemes.

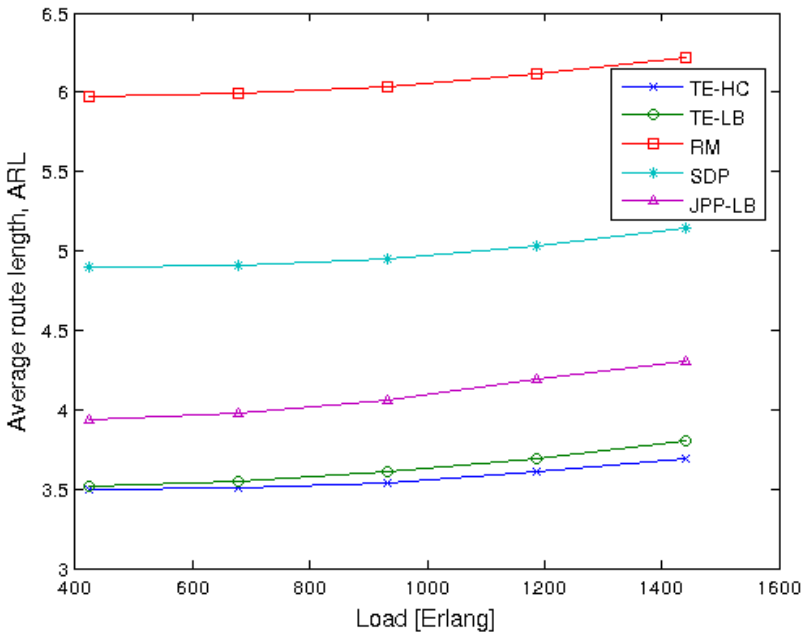

Figure 5.4: Average connection resource utilization (ARL)

5.1.2 Non-Even Stressor Event Distribution

The case of non-even stressor occurrences is tested next in order to model regions with higher/lower risk exposures. In particular, 3 of the SRLG regions in Figure 4.2 are given increased occurrence probabilities, i.e., $\phi_1 = 0.2$, and $\phi_1 = 0.3$. Meanwhile all other occurrence probabilities are fixed to $\phi_3 = 0.1$. As per Section 5.1.1, simulations are first done to measure the FR values, i.e., for both working and protection path failures (Equation 4.1). The related results are plotted in Figure 5.5 and follow closely

along the lines of those for equiprobable stressor events, i.e., Figure 5.1. In particular the RM scheme gives the best recovery performance, followed by the proposed JPP-LB scheme (which this time is closer to the SDP heuristic). However, the TE-based strategies still give very low recovery performance with the load-balancing option (TE-LB) outperforming the hop count option (TE-HC). The impact of non-even failure events on backup protection routes is also gauged by plotting the PFR values (Equation 4.2) in Figure 5.6. Overall these results are very similar to those in Figure 5.2 for equiprobable stressor event distributions. Namely, the risk-based RM heuristic again gives the highest failure rates for backup routes, i.e., by almost a factor of two. By contrast, the proposed JP-LBB scheme provides much better post-fault reliability for non-failed (working) connections and closely matches the TE-based strategies.

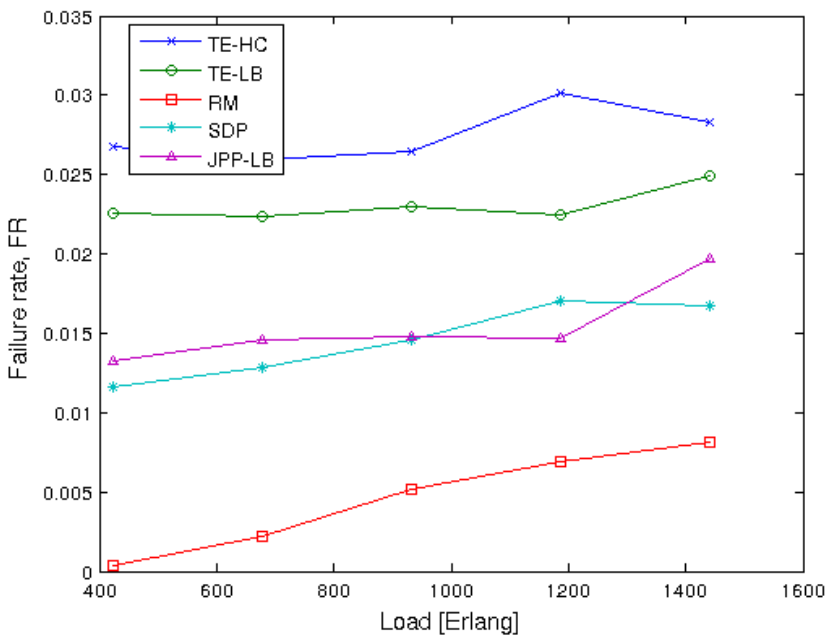

Figure 5.5: Failure rate for both working and protection routes (FR)

Simulation tests are also done to measure TE performance. Namely Figure 5.7 plots the overall request blocking rates using the BBR metric (Equation 4.3). Again,

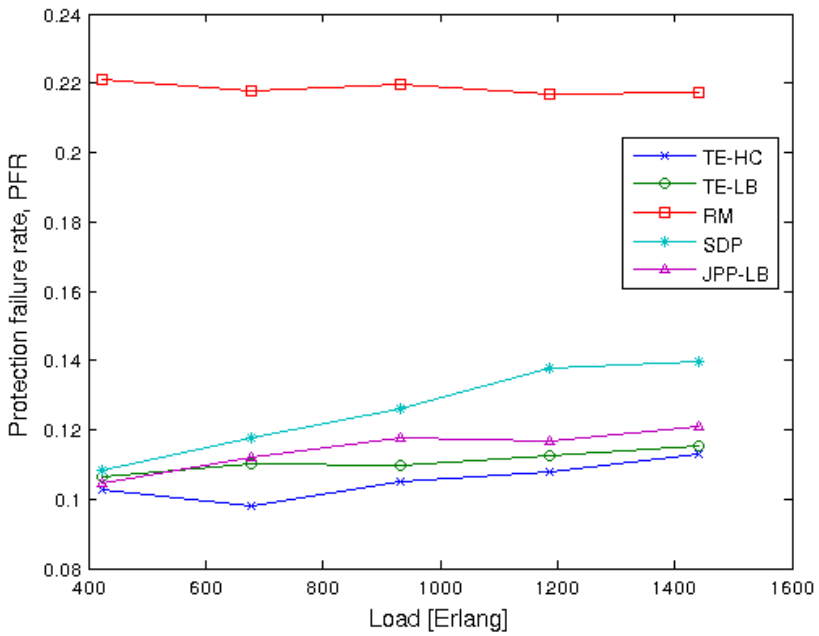

Figure 5.6: Failure rate for protection routes (PFR)

these findings show excessive demand rejection ratios with the risk-based heuristic strategies. For example the more advanced RM scheme is unable to achieve under 1% blocking even for very low input load regimes. By contrast the proposed JPP-LB scheme is more effective and yields blocking rates which are much closer to the TE-based strategies. Nevertheless, the TE-LB scheme does slightly better than the JPP-LB scheme for the case of non-even failure distributions, i.e., see increased separation of curves in Figure 5.7 as compared to Figure 5.3. Finally, average resource utilizations are also compared in Figure 5.8 using the ARL metric (Equation 4.4). Overall these results closely follow along the lines of those in Figure 5.4 for the case of equiprobable stressor distributions. In particular, the risk-based RM and SDP strategies are very resource-heavy and yield notably longer route lengths. For example, the route lengths with the RM scheme are almost twice as long as those with the TE-based schemes. By contrast, the JPP-LB scheme is much more efficient

and yields route lengths that are much closer to the TE-based approaches.

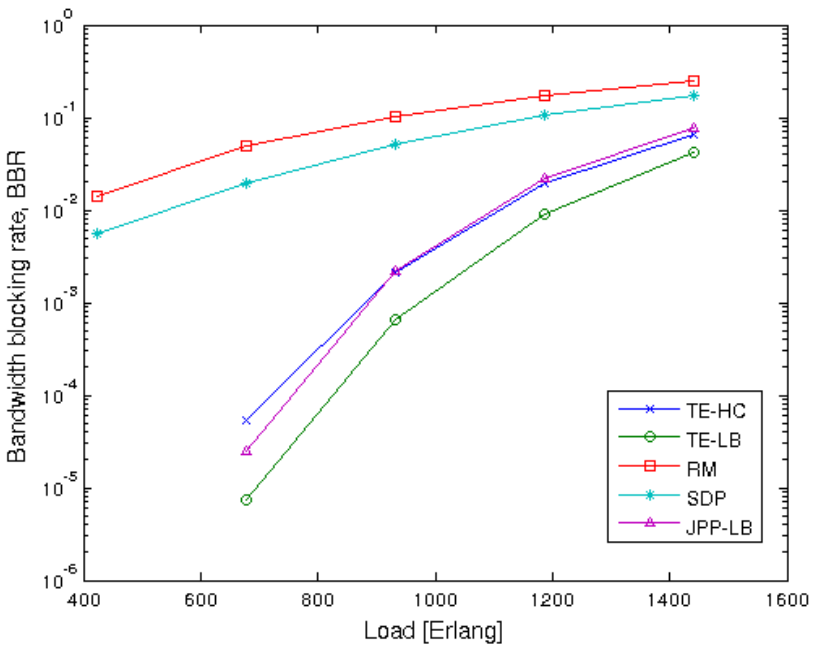

Figure 5.7: User request blocking rate (BBR)

5.2 Dependent Link Failure Scenarios

Performance evaluation is also conducted for the more realistic SRLG failure model developed in Section 3.4 (for correlated intra-regional link failures). However, as noted earlier in Section 3.3, the risk-based RM and SDP algorithms in [6] cannot handle these cases and are only applicable to independent link failure scenarios. As a result, this section only presents results for the JPP-LB and TE-based heuristic schemes. Furthermore, only equiprobable stressor occurrences are tested here, i.e., $\phi_i = \frac{1}{5}$ in Figure 4.3.

First, the reliability of working and protection routes is measured using the FR metric and the results plotted in Figure 5.9. Here the JPP-LB scheme gives the

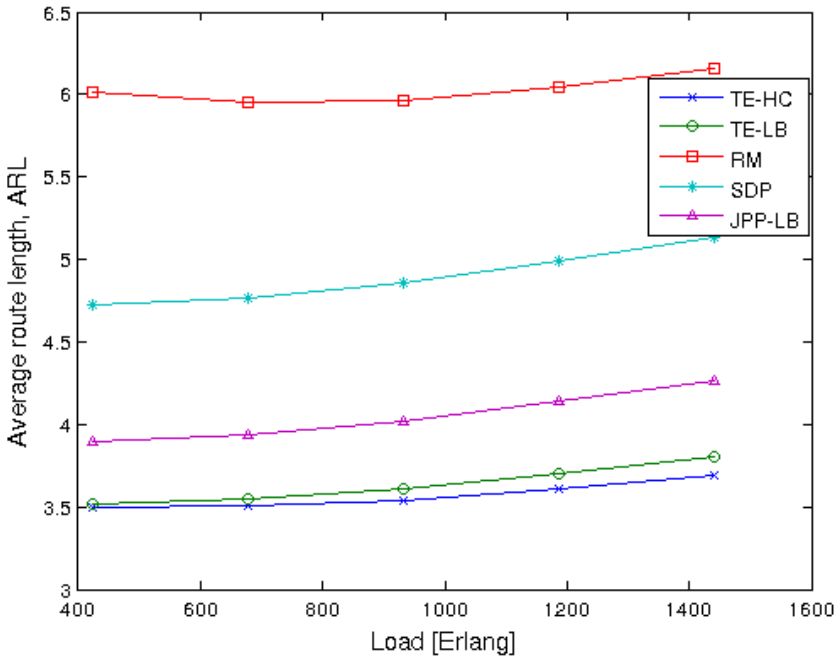

Figure 5.8: Average connection resource utilization (ARL)

lowest failure rates, averaging almost three times lower than the TE-based metrics, i.e., increased separation as compared to independent link failures, see Figure 5.1. Carefully note that increased input loads also tend to drive up the failure rate for the JPP-LB scheme, i.e., almost 5% higher at heavy loads. Meanwhile, the corresponding protection route failure rates are also shown in Figure 5.10 and indicate minimal separation (under 2%) between the proposed JPP-LB solution and the TE heuristic strategies. These results mirror earlier findings for independent link failures, Section 5.1. However, carefully note that dependent intra-SRLG link failures also yield slightly higher failure rates (versus independent link failures) for the three schemes under equivalent input loads. For example, this can be observed by comparing the FR results in Figures 5.1 and 5.9 and well as PFR results in Figures 5.2 and 5.10. This observation shows the importance of using more realistic failure models when trying to gauge multi-failure recovery performances.

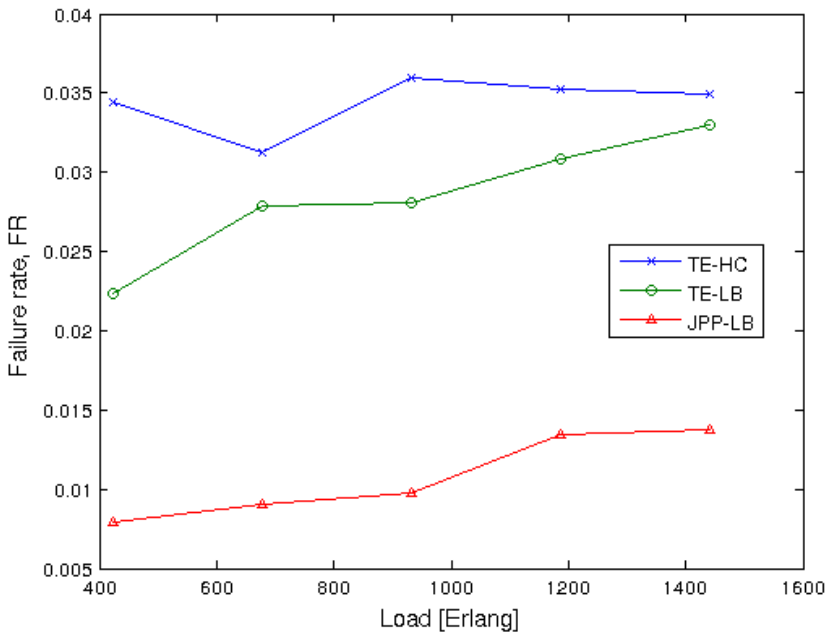

Figure 5.9: Failure rate for both working and protection routes (FR)

Finally, overall TE efficiencies are measured for dependent failures by plotting request blocking (BBR values) and resource utilization (ARL values) results in Figures 5.11 and 5.12, respectively. Again, these findings confirm that the JPP-LB solution gives very competitive blocking performance, which is consistently lower than the TE-HC heuristic and very close to the TE-LB scheme, see Figure 5.11. Meanwhile, the resource utilization results are also very impressive, averaging between $10-14\%$ higher usages. Overall, these findings show that the joint TE and risk minimization schemes developed in this thesis can allow operators to achieve much better service reliability without having to sacrifice their service revenues.

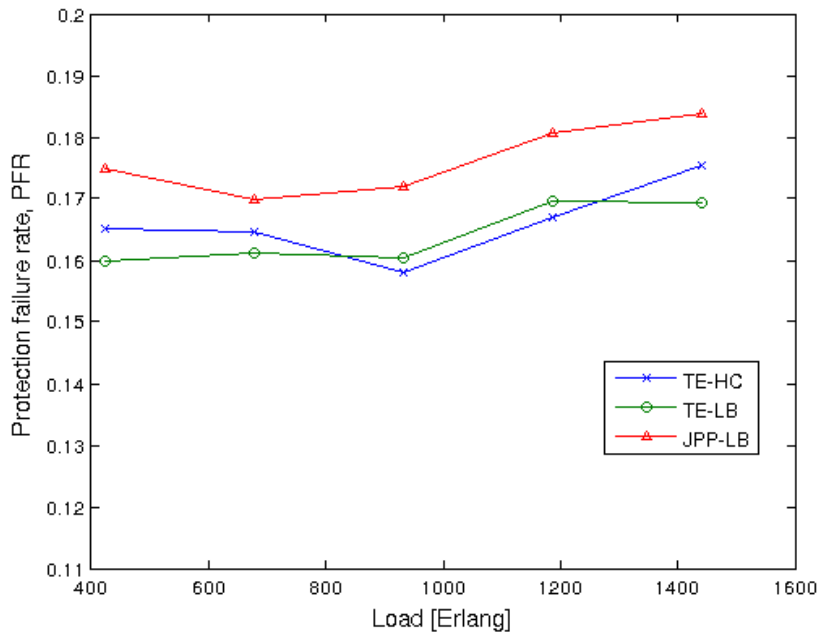

Figure 5.10: Failure rate for protection routes (PFR)

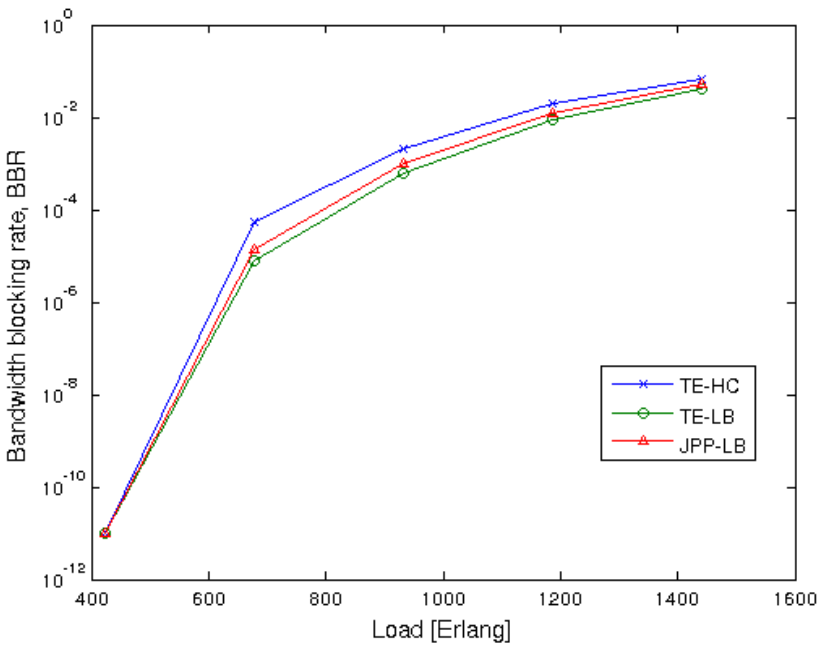

Figure 5.11: User request blocking rate (BBR)

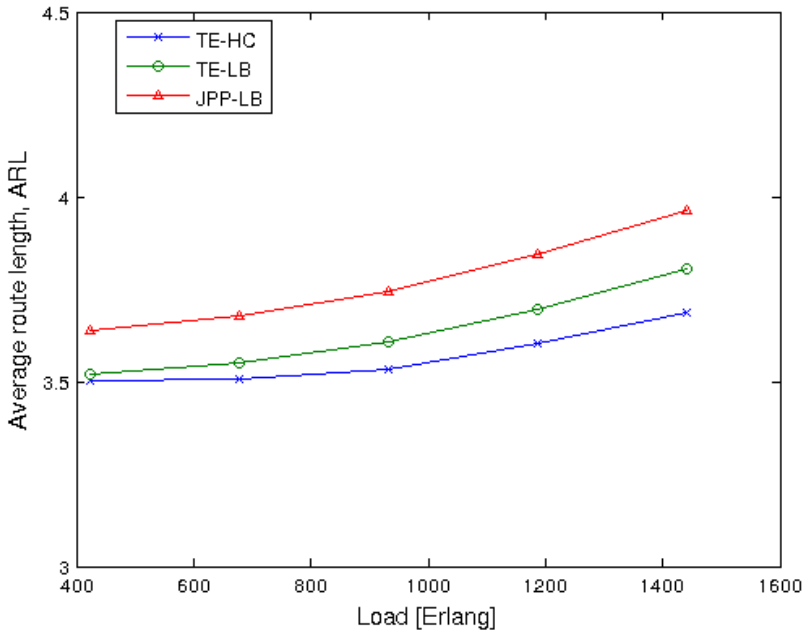

Figure 5.12: Average connection resource utilization (ARL)

Chapter 6

Conclusions and Further Work

Multi-failure network recovery is a very challenging topic and has become a key focus for many organizations today. By and large, most existing survivability schemes are simply not capable of handling massive correlated failures, particularly those arising during natural disasters, large-scale power outages, or malicious WMD attacks. As it is very difficult to pre-provision protection resources against all possible multi-failure node/link combinations, the use of probabilistic recovery schemes has been proposed to mitigate the risks associated with multiple network failures. However, there are only a handful of studies in this particular sub-topic area and most related solutions yield rather inefficient resource performances. Hence there is much room for improvement to develop new schemes for multi-failure network recovery.

Along these lines, this thesis focuses on building and testing an improved recovery scheme to handle correlated failures. This novel solution leverages graph-theoretic heuristics and implements a pre-provisioned protection approach to compute backup protection routes for working connections. A key innovation here is the design of a joint path pair selection scheme which takes into account both resource efficiency and risk minimization concerns. The overall performance of this solution is evalu-

ated using discrete event simulation for a sample representative national backbone network. Detailed comparisons are also made against an array of existing solutions. The overall findings and conclusions from this effort are now presented.

6.1 Conclusions

This research has proposed and tested an effective pre-protection solution to recover from multiple correlated network failures. The key findings from this effort include:

- The TE-based protection strategies performed as expected and delivered very high resource efficiency. This conclusion is based upon the ARL metric measurements. Moreover, the TE-LB scheme also gave the best blocking performance. However, these strategies yielded the lowest protection recovery rates, i.e., worst reliability. In particular, the TE-LB did worse that the TE-HC scheme owing to its selection of longer routes.

- The risk minimization schemes presented earlier in [6] did achieve high reliability, with the improved RM scheme doing notably better than the SDP scheme. However, both of these strategies yielded very high resource utilization and blocking behaviors, never falling below even 1%. As such, this will pose notably increased costs for network operators seeking to deploy such solutions.

- The proposed solution JPP achieves very good overall results. From a reliability perspective, it outperformed the basic SDP scheme and in many cases, closely matched the FR of the improved RM scheme. At the same time, it also yielded much lower blocking performances, almost on par with the TE-LB scheme. As such, this solution will provide a much more feasible alternative for network operators looking to improve service reliability with sacrificing their bottom line.

- The correlated failure model proposed in this work (Section 3.2.2) is a good representation of real-world scenarios, and the proposed JPP scheme can effectively apply it to improve service reliability. By contrast, the alternate RM and SDP schemes in [6] are not designed to handle more realistic correlated failure scenarios.

6.2 Further Research Directions

This effort has proposed one of the first network survivability schemes to jointly incorporate both resource efficiency and service reliability (risk mitigation) concerns. As such, this solution provides a very strong foundation from which to expand and develop more capable algorithms for the challenging multi-failure problem. For example, the further use of post-fault restoration schemes is an area that can be investigated to help increase survivability rates for user demands experiencing failures of both working and protection routes. Furthermore, detailed optimization formulations can also be developed to minimize resource overheads and risk exposure costs based upon the proposed joint heuristic strategy. These formulations can then be solved (or relaxed and solved) to yield theoretical performance bounds against which to further gauge the effectiveness of the proposed scheme. Finally, additional studies can be done to validate the heuristic strategy in distributed routing networks operating with delayed inaccurate routing state, i.e., network topology and resource information. This will provide detailed insights into the effectiveness of the solution under real-world conditions.

References

[1] Bruce Davie and Yakov Rekhter. *MPLS: Technology and Applications*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2000.

[2] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An architecture for differentiated service, 1998.

[3] P. Cholda etc al. A survey of resilience differentiation frameworks in communication networks. *IEEE Communication Surveys and Tutorials*, 2007.

[4] Kang Xi and Chao H. Jonathan. IP fast rerouting for single-link/node failure recovery. September 2007.

[5] Hungjen Wang, Eytan Modiano, and Muriel Medard. Partial path protection for wdm networks: End-to-end recovery using local failure information. In *IEEE ISCC02*, July 2002.

[6] Hyang-Won Lee, E. Modiano, and Kayi Lee. Diverse routing in networks with probabilistic failures. *Networking, IEEE/ACM Transactions on*, 18(6):1895 – 1907, December 2010.

[7] Qingya She, Xiaodong Huang, and J.P. Jue. Maximum survivability under multiple failures. March 2006.

[8] S. Stefanakos. Reliable routings in networks with generalized link failure events. *Networking, IEEE/ACM Transactions on*, 16(6):1331 –1339, December 2008.

[9] Mahesh Sivakumar, Rama K. Shenai, and Krishna M. Sivalingam. A survey of survivability techniques for optical wdm networks. In Krishna M. Sivalingam and Suresh Subramaniam, editors, *Emerging Optical Network Technologies*. Springer US, 2005. 10.1007/0-387-22584-6_13.

[10] Hongsik Choi, Suresh Subramaniam, and Hyeong ah Choi. On double-link failure recovery in wdm optical networks. In *Proceedings of IEEE INFOCOM*, March 2002.

[11] Bondy J. A. and Murty U. S. R. *Graph Theory With Applications*. Macmillan Press, London, 1976.

[12] Lu Ruan and Taiming Feng. A hybrid protection/restoration scheme for two-link failure in wdm mesh networks. In *GLOBECOM 2010, 2010 IEEE Global Telecommunications Conference*, December 2010.

[13] Ramesh Bhandari. *Survivable Networks: Algorithms for Diverse Routing*. Kluwer Academic Publishers, Norwell, MA, USA, 1998.

[14] M.T. Frederick, P. Datta, and A.K. Somani. Evaluating dual-failure restorability in mesh-restorable wdm optical networks. In *Computer Communications and Networks, 2004. ICCCN 2004. Proceedings. 13th International Conference on*, October 2004.

[15] Bijan Bassiri and Shahram Shah Heydari. Network survivability in large-scale regional failure scenarios. In *Proceedings of the 2nd Canadian Conference on Computer Science and Software Engineering*, C3S2E '09, New York, NY, USA, 2009.

[16] M. Rahnamay-Naeini, J.E. Pezoa, G. Azar, N. Ghani, and M.M. Hayat. Modeling stochastic correlated failures and their effects on network reliability. In *Computer Communications and Networks (ICCCN), 2011 Proceedings of 20th International Conference on*, August 2011.

[17] M. Menth and R. Martin. Network resilience through multi-topology routing. In *Design of Reliable Communication Networks, 2005. (DRCN 2005). Proceedings.5th International Workshop on*, October 2005.

[18] Greg Bernstein, Bala Rajagopalan, and Debanjan Saha. *Optical Network Control: Architecture, Protocols, and Standards*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2003.

[19] O.J.S. Parra, C. Manta, and G. Lopez Rubio. Dijkstra's algorithm model over mpls / gmpls. In *Wireless Communications, Networking and Mobile Computing (WiCOM), 2011 7th International Conference on*, September 2011.

[20] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano. Assessing the vulnerability of the fiber infrastructure to disasters. In *INFOCOM 2009, IEEE*, April 2009.

[21] P.K. Agarwal, A. Efrat, S. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman. The resilience of wdm networks to probabilistic geographical failures. In *INFOCOM, 2011 Proceedings IEEE*, April 2011.

[22] D. Medhi. Quality of service (qos) routing computation with path caching: A framework and network performance. *IEEE Communication Magazine*, 40(12):106–113, December 2002.

[23] D.O. Awduche. Mpls and traffic engineering in IP networks. *Communications Magazine, IEEE*, 37(12):42 –47, December 1999.

[24] David Eppstein. Finding the k shortest paths, 1997.

[25] C. Xin, Y. Ye, S. Dixit, and C. Qiao. A joint lightpath routing approach in survivable optical networks. *Optical Networks Magazine*, 3:13–20, 2002.

[26] N. Ghani, S. Park, A. Shami, C. Assi, K. Atthuru, and B.J. Ayeleso. Multi-tiered services in next-generation sonet/sdh networks. In *Communications, 2006. ICC '06. IEEE International Conference on*, Istanbul, Turkey, June 2006.

[27] John Hershberger, Matthew Maxel, and Subhash Suri. Finding the k shortest simple paths: A new algorithm and its implementation. *ACM Trans. Algorithms*, 3(4), November 2007.

[28] Qingya She, Xiaodong Huang, and J.P. Jue. Maximum survivability using two disjoint paths under multiple failures in mesh networks. In *Global Telecommunications Conference, 2006. GLOBECOM '06. IEEE*, Sun Francisco, CA, December 2006.

[29] M. Esmaeili, M. Peng, S. Khan, J. Finochietto, Y. Jin, and N. Ghani. Multi-domain dwdm network provisioning for correlated failures. In *Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2011 and the National Fiber Optic Engineers Conference*, March 2011.