9-15-2014

Optimization Design for Multi-Domain Optical Network Provisioning and Survivability

Liang Kaile

Kaile Liang

Candidate

Electrical and Computer Engineering

Department

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

Approved by the Dissertation Committee:

Nasir Ghani , Chairperson

Wei Shu

Marios S. Pattichis

Patrick Bridges

Optimization Design for Multi-Domain Optical Network Provisioning and Survivability

by

Kaile Liang

B.S., Beijing University of Posts and Telecommunications, 2009
M.S., Beijing University of Posts and Telecommunications, 2012

DISSERTATION

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy
Engineering

The University of New Mexico

Albuquerque, New Mexico

July, 2013

Dedication

To my parents, my advisor, and all my friends...

Acknowledgments

I would like to extend my sincere gratitude to my advisor, Dr. Nasir Ghani, for his encouragement, advisement, and continuous help throughout my studies and research, and particularly in the completion of this dissertation.

I would also like to thank my dissertation committee members, Dr. Wei Shu, Dr. Marios S. Pattichis, and Dr. Patrick Bridges, for their valuable comments and advice on my research.

Next, I would like to thank my research colleagues and friends, including Dr. Feng Xu, Dr. Feng Gu, Mr. Hao Bai, Mr. Juan Antonio Elices and Mr. Zhenjie Chen, for their help and invaluable friendship.

Finally, I would also like to thank my parents, for giving me life, freedom, and endless love.

Optimization Design for Multi-Domain Optical Network Provisioning and Survivability

by

Kaile Liang

B.S., Beijing University of Posts and Telecommunications, 2009
M.S., Beijing University of Posts and Telecommunications, 2012

Ph.D, Engineering, University of New Mexico, 2013

Abstract

User bandwidth demands have continued to expand at an unprecedented rate over the past two decades. This growth has been driven by the emergence of many new applications across a wide range of sectors, including commercial, private, and scientific computing. As a result, network carriers and operators have deployed a wide range of high-speed technologies to meet their growing needs. In particular, these solutions include state-of-the-art *Internet Protocol* (IP) and Ethernet systems as well as optical *wavelength division multiplexing* (WDM) networking platforms. In particular, the latter solutions provide unmatched terabit-per-second speeds and are commonly used to provide underlying "lightpath" connectivity between IP and Ethernet devices.

Now researchers have developed a range of schemes for lightpath provisioning and survivability in WDM networks. Most notably these solutions include optimization and heuristic-based strategies to solve the *routing and wavelength assignment* (RWA)

problem. However, as WDM deployments have expanded, there is a further need to provision lightpath connections across *multiple* network domains. For example, these domains can be delineated in a variety of manners, including administrative ownership (intra- and inter-carrier), vendor or technology type, geographic, etc. Given the above, multi-domain (optical) network provisioning and survivability has become a key focus area. Indeed, this is a rather challenging problem as scalability and privacy concerns limit the amount and type of information that can be shared across domain boundaries, i.e., particularly in inter-carrier settings. Hence researchers have developed various solutions, with most using distributed graph-based heuristics to resolve connection routes with partial (dated, inaccurate) network state.

Nevertheless, it is well-understood that heuristic schemes are sub-optimal in nature and cannot provide any bounds on network performance. As a result, most multi-domain studies have used other heuristics for comparison purposes. In light of this, it is very difficult for network carriers to gauge the true achievable performance of their multi-domain networking setups. However, optimization-based methods offer a very effective means of formally analyzing network performance under idealized conditions with full a-priori knowledge of user demands. Moreover, these schemes have been widely-used to bound lightpath RWA performance in single-domain settings. Nevertheless, the further application of such methods in multi-domain network settings has not yet been considered.

To address these concerns, this dissertation presents a comprehensive optimization-based study of lightpath routing and survivability in multi-domain optical networks. First, a novel (two-stage) hierarchical model is introduced to optimize lightpath routes over inter-/intra-domain topologies pursuant to several *traffic engineering* (TE) objectives, i.e., including throughput maximization, resource minimization, and load balancing. Next, this model is extended to implement lightpath protection recovery for single-link failures by adding link-disjointness constraints at both

the intra- and inter-domain levels. Finally, a novel optimization formulation is also developed to implement *probabilistic* lightpath protection for multiple correlated failures, i.e., as occurring during large-scale disaster events. The performance of these differing schemes is also tested for several multi-domain network configurations and compared against some advanced heuristic strategies, i.e., including regular (working-mode) provisioning, single-link protection, and probabilistic protection. Overall, the detailed findings from this effort show that the new optimization schemes give significantly better results, thereby providing an invaluable benchmark reference from which to develop improved solutions.

Contents

List of Figures

List of Tables

Glossary

ASTN	Automatic Switched Transport Network
BGP	Border Gateway Protocol
BRPC	Backward Recursive Path Computation
CIDA	Capacitated Iterative Design Algorithm
COPRP	Constraint-based Optical Path-vector Routing Protocol
DDRP	Domain-to-Domain Routing Protocol
DD	Domain Disjoint
ESnet	Energy Science Network
ER	Explicit Route
FF	First Fit
FIPP	Failure Independent Path Protection
GMPLS	Generalized Multi-Protocol Label Switching
GLPK	GNU Linear Programming kit
HT	Hold-Down Timer

IETF	Internet Engineering Task Force
ILP	Integer Linear Programming
INLP	Integer Non-Linear Programming
IP	Internet Protocol
IT	Information Technology
ITU-T	International Telecommunications Union
LR	Loose Route
LD	Link Disjoint
NNI	Netowrk-to-Network Interface
OADM	Optical Add-Drop Multiplexers
OBGP	Optical BGP
OC-n	Optical Carrier
OIF	Optical Internetworking Forum
OSPF	Open Shortest Path First
OTU-n	Optical Transport Unit
OXC	Optical Cross-Connects
PCC	Path Computation Client
PCE	Path Computation Element
PCEP	PCE Communication Protocol

p-cycle	Pre-Configured Cycle
p-SRLG	Probabilistic Shared Risk Link Group
Qos	Quality of Service
RC	Routing Controller
RSVP-TE	Resource Reservation Protocol-Traffic Eigineering
RWA	Routing and Wavelength Assignmentt
SCF	Signigicance Change Factor
SLA	Service Level Agreement
SMF	Single Mode Fiber
SONET	Synchronous Optical Network
SRLG	Shared Risk Link Group
TE	Traffic Engineering
TEDB	Traffic Engineering Databases
UNI	User-Network Interface
VSPT	Virtual Span Path Tree
WDM	Wavelength Division Multiplexing

Chapter 1

Introduction

This thesis dissertation presents a detailed optimization-based study of multi-domain optical network provisioning and survivability. In order to properly introduce the work, this initial chapter reviews the main research in multi-domain networking, which is largely comprised of heuristics-based strategies. The key motivations for the dissertation are then presented, along with the core contributions of the research.

1.1 Background Overview

The past two decades have seen massive growth in networking bandwidth demands. This expansion has been fueled by the continual emergence (and flux) of a wide range commercial, scientific, and personal user applications and services. Most notably, the scientific community has emerged as one of the heaviest consumers of high-bandwidth services. For example, many modern research projects involve collaborations between dispersed teams and require transfer/processing of increasingly large data-sets, e.g., in areas such as high-energy physics, astrophysics, genetics, climate change modeling, etc [1]. An example of this growth is presented in Figure 1.1, which shows that the

aggregate transfer demands over the Department of *Energy Science Network* (ESnet) backbone have increased by an order of magnitude every 4 years. Moreover current projections in the scientific community are pointing towards *exascale-level* needs within the 5-10 years [2]. At the same time, commercial and private bandwidth consumption has also expanded manifold, driven by the increasing popularity of content-sharing sites (such as Youtube, Facebook, etc) and an array of recent cloud-based service offerings.

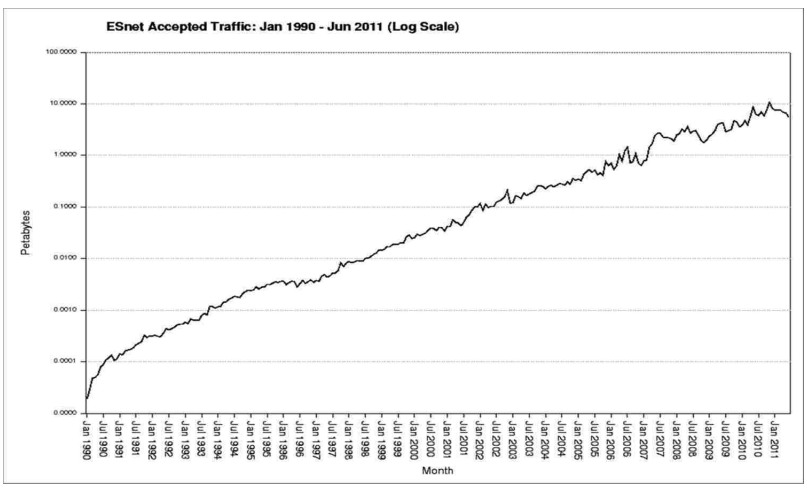

Figure 1.1: Accepted traffic demand on ESnet backbone, from [3]

Now in order to meet these growing needs, equipment designers have developed a wide range of *wireline* networking technologies and architectures. For example, core *Internet Protocol* (IP) routing platforms have evolved to support multi-gigabit port interfaces and advanced flow-level *quality of service* (QoS) features. Ubiquitous Ethernet technologies have also undergone a full metamorphosis and can now deliver high-speed connectivity across large metro and core distances, i.e., via new developments in Carrier Ethernet standards [4]. Finally, legacy *synchronous optical network* (SONET) technologies have been phased out at the underlying fiber-optic layer and replaced by optical *wavelength division multiplexing* (WDM) paradigms [5]. In particular, WDM systems use multiple optical frequencies (wavelengths) to send

data over *single mode fiber* (SMF), and current systems can readily support over 100 wavelengths/fiber at speeds of 10-40 Gbps each, i.e., tributaries such as 10/40 Gigabit Ethernet, SONET *optical carrier* (OC-n), and *optical transport unit* (OTU-n). As such, these capabilities offer unprecedented terabit-level throughputs over a single strand of fiber. Moreover, new WDM sub-system technologies have allowed designers to build specialized optical platforms to support spatial wavelength switching across arbitrary fiber-plants, e.g., such as *optical cross-connects* (OXC), *optical add-drop multiplexers* (OADM), etc [6]. In turn these advances have enabled high-bandwidth wavelength "lightpath" connection routing over fiber substrates, greatly reducing (even eliminating) the need for costly electronic data buffering and processing, see Figure 1.2. Arguably, these advances have been the most critical enablers for network scalability in recent years.

Overall, network carriers have extensively deployed new IP, Ethernet, and WDM system technologies to provision scalable connectivity across their wide-area cores. In particular, WDM platforms are being used to support dynamic "on-demand" optical lightpath connectivity between "higher-layer" IP and Ethernet systems, i.e., two-layer network architecture, as shown in Figure 1.2. To streamline these capabilities, a range of control plane solutions have also been standardized. Most notably, the *Internet Engineering Task Force* (IETF) has evolved its *generalized multi-protocol label switching* (GMPLS) framework to support routing and signaling setup across multiple network technology layers [5]. The related *path computation element* (PCE) framework [7] has also been introduced to help improve *traffic engineering* (TE) path computation. Finally, other organizations, such as the *Optical Internetworking Forum* (OIF), have also developed standards to interconnect optical networking domains, i.e., protocols for routing and signaling [5, 8].

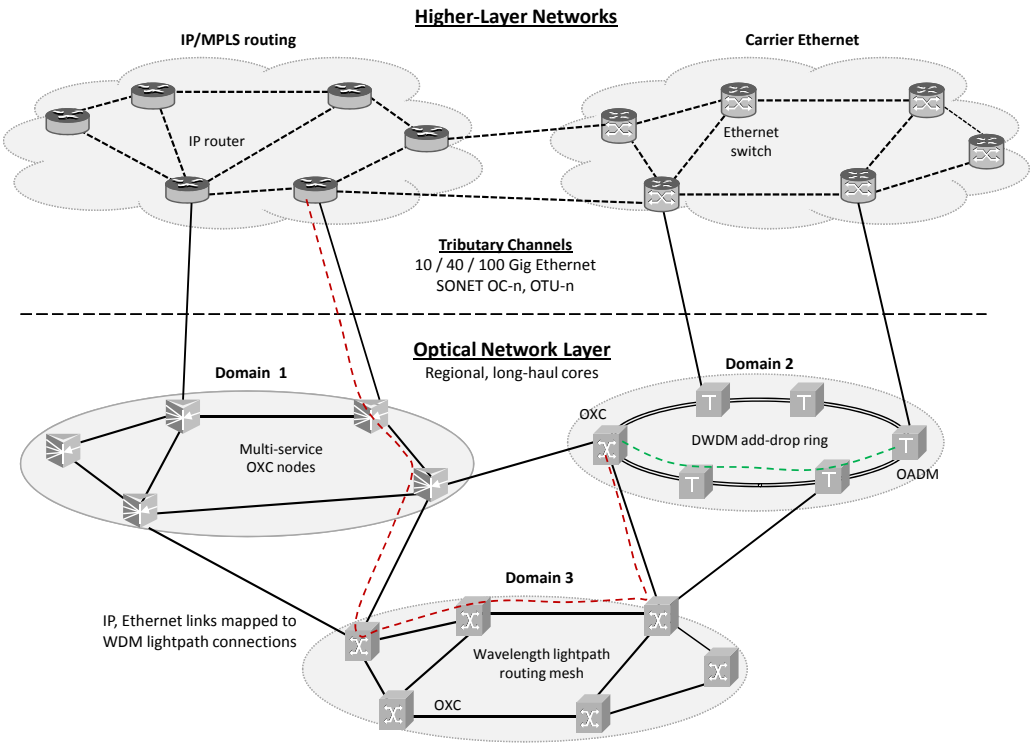

Figure 1.2: Multi-domain optical network overview

1.2 Motivations

As WDM technologies have become increasing prevalent, there is a growing need to provision optical lightpath connectivity across *multiple* network domains. In particular, these domains can be delineated in a variety of manners, either based upon vendor technology type, geographic distance/coverage, administrative carrier ownership, etc [8, 9]. For example, modern scientific collaborations are increasingly dependent upon gigabit-level bandwidth interconnection between research teams located on different continents, e.g., high energy physics researchers transferring data between Europe and the United States [10]. Typically, these interconnections must

be routed over multiple public, and even private, network domains, each operated by different entities and located in different countries. The continued growth in commercial (cloud-based) services is also driving similar bandwidth needs within the commercial sector.

However, multi-domain network provisioning is a rather challenging problem area, and one that is receiving increased focus in recent years, see [8, 9]. In particular, privacy concerns prevent most operators from exercising complete routing state exchange across domain boundaries, especially in inter-administrative/inter-carrier settings. Moreover, in cases where privacy may not be that important, scalability issues still inhibit complete state exchange, e.g., single carrier settings with multiple administrative/regional domains. Hence due to these restrictions, multi-domain lightpath provisioning has to be done using limited partial knowledge of global network state, i.e., in terms of static topological interconnectivity and dynamic resource usage. In light of the above, a range of decentralized multi-domain (optical) network provisioning strategies have been developed, with many drawing from earlier solutions for inter-area connectivity in packet-switching IP networks [8]. Consider some further details here.

One of the main challenges in multi-domain provisioning is to achieve a level of "global" state visibility to support "end-to-end" constraint-based path provisioning. Along these lines, various hierarchical routing protocols have been developed to disseminate a reduced amount of information between select domain nodes, e.g., typically inter-gateway border nodes, see generic architecture in Figure 1.3. Generally speaking, these designs can be classified as either path/distance-vector routing or link-state routing designs, see [8]. For example, several researchers have looked at link-state routing in multi-domain optical network settings and applied topology abstraction techniques [11] to reduce domain-internal state, see [12–16]. This condensed information is then used by domain-level path computation systems to

compute skeleton inter-domain domain routes using graph-based heuristics. Finally, these routes are expanded using distributed signaling procedures to resolve full intra-domain node/link sequences. Nevertheless, link-state routing designs can impose notable messaging overheads, especially when using complex topology abstraction schemes. Hence alternate path-vector routing strategies have also been developed for multi-domain optical lightpath provisioning, see [13, 17–19]. Finally, recent studies have outlined "per-domain" provisioning strategies [20, 21] based upon the IETF PCE framework [7]. In general, these schemes do not require hierarchical routing support and instead assume the availability of pre-computed, i.e., static, inter-domain route sequences. Hence the focus here is on selecting ingress and egress border gateway nodes and using graph-based algorithms to compute local intra-domain traversing sub-path segments. As such, per-domain strategies tend to yield sub-optimal resource efficiency under dynamic traffic conditions.

As user demands scale and more and more organizations move their "mission-critical" applications over expanding network infrastructures, multi-domain survivability is also becoming a major concern. Now given the high bandwidth scalability of underlying fiber-optic substrates, even isolated single link failures can cause extreme service disruption for many users. Moreover these concerns are further exacerbated in more complex failure scenarios, particularly those arising during catastrophic disaster events. For example, these stressors can include natural events (such as earthquakes, floods, hurricanes, etc) as well as malicious or inadvertent man-made disasters (such as weapons of mass destruction attacks, power outages, etc). In general, these occurrences will yield large numbers of node and link failures with very high-levels of spatial and temporal correlation, see [22]. In turn, these faults can easily overwhelm existing (mostly single failure) recovery mechanisms, and related network recovery concerns are only now starting to be addressed. Consider some further details here.

Broadly speaking, network survivability algorithms can be classified into two

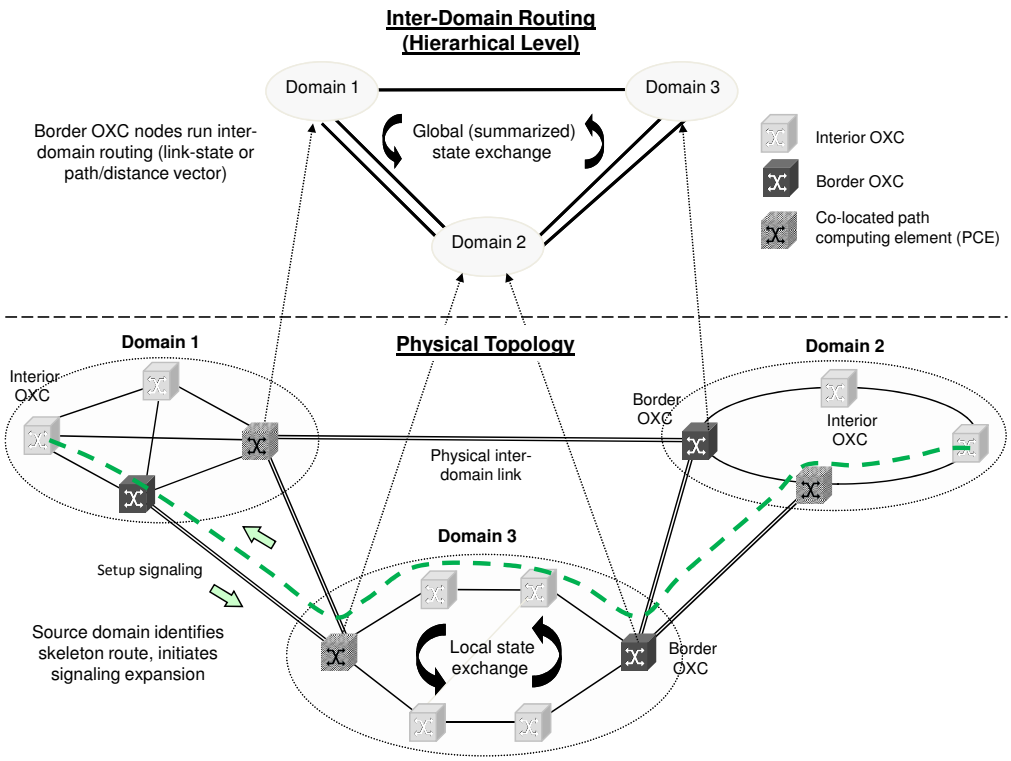

Figure 1.3: Overview of distributed hierarchical multi-domain architecture

main categories, i.e., pre-provisioned *protection* and post-fault *restoration*, both of which can also be applied in multi-domain settings [8]. Namely, protection schemes are "proactive" and pre-provision bandwidth resources at setup time, e.g., disjoint backup paths, sub-paths, or links. As a result, these schemes can provide very fast recovery by simply switching to pre-allocated backup resources upon failure detection along primary resource routes, i.e., milliseconds range. However protection is also more complex and costly to implement. By contrast, restoration schemes do not reserve any backup resources and simply re-compute affected working routes after a failure event. However, these schemes cannot guarantee recovery, even for single link failures, and yield longer recovery times, i.e., seconds range [5,23]. Now consider

these strategies within the multi-domain context.

In general, multi-domain survivability is a very challenging problem since designers must contend with limited knowledge of intra- and inter-domain topological diversity. Now most related solutions have proposed protection designs based upon dedicated (even shared) resource pre-provisioning. For example, several solutions have built augmented topology abstraction schemes for hierarchical routing to carry domain-level path diversity information [24–26]. This state is then used by modified graph-based algorithms to compute/expand disjoint primary/backup path-pairs across multiple domains. However, these strategies entail very high routing overheads (due to specialized topology abstractions) and can only guarantee recovery from single link failures. Meanwhile, other efforts have proposed extensions to per-domain provisioning strategies to compute primary/backup path-pairs along the same (fixed, pre-specified) sequence of domains [20,27,28], However, these schemes are largely ineffective against multi-failure stressors since both primary and backup routes can be affected. As a result more recent work in [29] has proposed one of the first solutions for multi-domain protection under disaster conditions. Specifically, large stressors are fist modeled as probabilistic events centered about a given geographic region, and diverse path-pair computation schemes are then developed to minimize risk exposures (joint failure probabilities) in conjunction with TE concerns. Finally, several multi-domain post-fault restoration schemes have also been studied using signaling crankback (re-try) methodologies, i.e., end-to-end and intermediate [30]. These schemes can also be coupled with protection strategies to provide improved multi-tiered recovery, particularly if exact failure patterns are not known [31].

Nevertheless, most multi-domain provisioning and survivability schemes are heuristic-based in nature and rely upon graph-theoretic algorithms to compute routes over intra- and inter-domain graph topologies. Moreover, these solutions also operate with delayed/inaccurate hierarchical routing state [8], resulting in reduced effective-

ness, e.g., in terms of blocking performance and resource efficiency. Hence it is very difficult to ascertain the ideal achievable performance of such schemes. By contrast, optimization-based methods provide a very powerful means of analyzing network performance. Moreover, these techniques have been widely used for optical network *routing and wavelength assignment* (RWA) with a-priori knowledge of user demands, see [32–34]. Nevertheless, these optimization-based studies only consider single-domain network settings with complete global network topology and resource state. Indeed, the further application of such algorithms in *multi-domain* settings operating with partial state information has not been considered and remains an open issue. Perhaps the only related works here are some studies on multi-domain *protection-cycle* (p-cycle) design [35–37]. However, these strategies only implement static link-level protection and are not suitable for protecting individual demands.

1.3 Problem Statement

This dissertation is motivated by the need to develop more formalized analytical models for multi-domain optical network RWA provisioning and survivability. As most efforts in this field have focused on "adhoc" heuristic strategies, these schemes cannot provide any type of optimal bounds per say. Hence it is very difficult for network operators to gauge the true achievable performance of their multi-domain setups. In addition, few efforts have studied multi-domain network survivability under highly-challenging multi-failure disaster conditions, especially from an analytical perspective.

To address these concerns, this dissertation develops some novel, scalable optimization models for multi-domain optical network provisioning, i.e., for both regular working-mode and survivability scenarios (single and multi-failure). A key goal here is to capture the hierarchical nature of multi-domain networks operating with partial

global state information. These detailed formulations are then solved using advanced programming packages, and their results compared against state-of-the-art heuristic strategies (implemented via network simulation). As such, this work can be used bound performance in multi-domain settings and also improve existing heuristic-based designs.

Note that this dissertation only considers optimization in multi-domain optical network settings, i.e., with wavelength resource link constraints. Although further extensions can be considered for more general networks with link-level bandwidth *and* delay constraints, e.g., such as IP/MPLS and Ethernet domains, this is a broader problem that is left for future study.

1.4 Thesis Outline

This dissertation addresses a range of open challenges in the area of multi-domain optical network provisioning, with a focus on developing new optimization-based strategies. Namely, the first part looks at optimization models for regular (non-survivable) provisioning in the presence of reduced (hierarchical) global routing state information. Meanwhile, the next part extends this work to consider multi-domain protection optimization for basic single-failure scenarios. Finally, the last part studies optimization design for more challenging multi-failure disaster conditions. In summary, the following topics are addressed:

(1) Hierarchical two-stage optimization model for lightpath provisioning in multi-domain optical networks

(2) Hierarchical two-stage optimization model for link- and domain-disjoint protection against single link failures in multi-domain optical networks

(3) Optimization models for "risk-aware" lightpath protection in multi-domain networks in the presence of multiple (probabilistic) correlated failures

Overall, this thesis report is organized as follows. First, Chapter 2 presents a survey of the various strategies and latest techniques for provisioning and survivability in multi-domain (optical) networking environments. Next, Chapter 3 presents a novel formulation for lightpath optimization in multi-domain networks and also compares its performance with an existing heuristic strategy. Chapter 4 then expands this formulation to handle singe link failures via the addition of new disjointness constraints and also conducts a detailed performance evaluation study. Subsequently, Chapter 5 treats optimization design for lightpath protection in multi-domain settings with multiple probabilistic failures. Finally, conclusions and further research directions are presented in Chapter 6.

Chapter 2

Background and Related Work

The topic of (optical) multi-domain network provisioning and survivability has seen increased focus and attention in recent years. Overall, the majority of solutions presented in this area have been heuristics-based strategies that operate with a reduced, i.e., partial global state visibility. Along these lines, this chapter categorizes and reviews of some of the key contributions in this space, including standards, research studies, and recent work on multi-failure recovery.

2.1 Multi-Domain Architectures & Standards

Multi-domain networking is an important area for carriers, and a variety of protocol architectures have emerged to support routing exchange and provisioning/setup between domains. These standards are briefly reviewed here as many of them are also leveraged by related multi-domain heuristic schemes, Sections 2.2 and 2.3. Interested readers are also referred to [5] and [8], and the related references therein for more detailed overviews of some of these standards.

The IETF has developed some of the most widely-used standards for multi-

domain networking. Broadly speaking, these offerings include protocols for routing, signaling, and path computation. Foremost, the IETF's legacy *border gateway protocol* (BGP) is now the defacto standard for address prefix reachability exchange across Internet domains. This solution uses a hierarchical routing approach in which border gateway router nodes exchange path-vector state information, i.e., by specifying complete domain sequences to IP prefix ranges. However, "vanilla" BGP does not carry explicit topological state information and inherently couples routing exchange with path computation [8]. As a result, it is rather difficult to implement advanced TE support or survivability with this protocol.

The IETF also offers another routing solution via its *open-shortest path first* (OSPF) protocol [38], albeit primarily for intra-domain support. In particular, this protocol implements detailed link-state routing exchange between nodes and provides complete topology and active resource exchange. Furthermore, OSPF also supports a higher (second) level of routing in which select nodes disseminate condensed state information, and in turn, this can be leveraged for inter-domain provisioning. In addition, a range of OSPF-TE-based extensions have also been standardized to support link-state exchange for a full range of "non-packet switching" network link types, e.g., such as optical WDM and SONET time-division-multiplexing, see [5]. Note that these additions are part of the broader GMPLS protocol framework [8, 39].

Resource reservation is also a critical requirement in multi-domain settings. Now here the IETF *resource reservation protocol-traffic engineering* (RSVP-TE) already presents a very well-established signaling capability using a backwards reservation approach [40]. Moreover this protocol provides some inherent functions for multi-domain setup as well. For example, the RSVP-TE *loose route* (LR) feature allows operators to compute skeleton inter-domain routes and then expand them using *explicit route* (ER) signaling procedures. Newer crankback signaling extensions have also been proposed for RSVP-TE [41] and are very useful when trying to establish

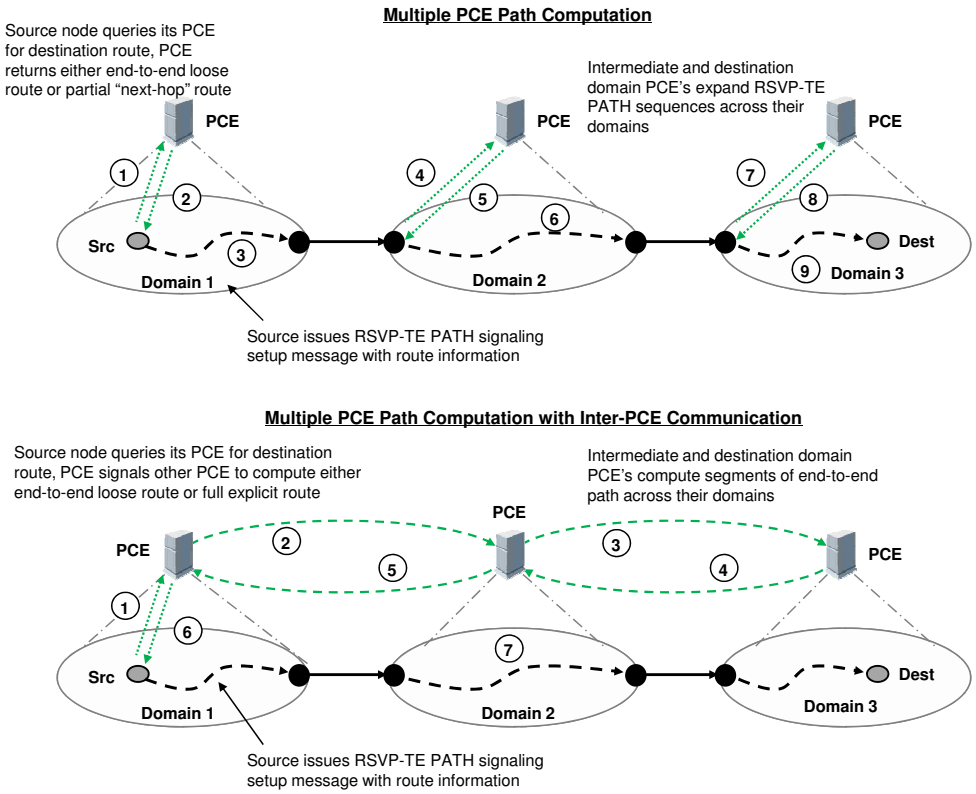

Figure 2.1: PCE-based multi-domain path computation schemes, from [8]

multi-domain routes using limited inaccurate state information.

Finally, the IETF has also standardized a framework for constraint-based path computation across domains. Namely, the PCE [7] standard defines a path computation entity for each domain (area) to compute local and inter-domain route sequences for requesting *path computation client* (PCC) nodes. Here a PCE can either be located at a stand-alone server or be co-located with a network routing/switching node. Now at the intra-domain level, a PCE is assumed to have complete local domain state knowledge, i.e., via access to a distributed (OSFP-TE) routing database or a centralized controller repository. Hence local-domain routes can be resolved in a standalone

manner. However at the inter-domain level, PCE devices will likely have varying degrees of global visibility, i.e., depending upon the type of hierarchical routing setup in place, if any. To handle this ambiguity, the PCE standard proposes a more generalized *distributed* multi-domain computation approach in which PCE systems interact with each other by using a *PCE communication protocol* (PCEP) [42] to compute end-to-end routes. Specifically, two methodologies are proposed here, *per-domain* and *inter-PCE signaling*, as shown Figure 2.1. The former performs path computation in conjunction with inter-domain signaling setup for a connection route, i.e., joint computation and signaling. Conversely, the latter first resolves the complete end-to-end path between the PCE systems and then initiates signaling setup, i.e., separate computation and signaling.

Some other organizations have also developed their own multi-domain networking standards. For example, the OIF has designed protocols for interfacing with multi-vendor networking devices/domains. This work includes physical-layer standards, as well as client-to-carrier *user network interface* (UNI) and carrier-to-carrier *network-to-network interface* (NNI) protocols [5]. For example, the OIF UNI provides bandwidth signaling for client devices to request/release connections from optical networking domains. Meanwhile the NNI standard provides more extensive capabilities for reachability/resource exchange and setup signaling. This standard also specifies two variants for intra- and inter-carrier operation, respectively [8]. Finally, a hierarchical *domain-to-domain routing protocol* (DDRP) [43] is also included and leverages the OSPF-TE link-state routing design. Overall, the core UNI and NNI functions have been demonstrated in several successful multi-vendor trials.

The *International Telecommunications Union* (ITU-T) has also matured an abstract control framework for generalized multi-domain operation, i.e., termed as the *automatic switched transport network* (ASTN) standard (G.8080) [44]. This solution defines a hierarchical setup with multiple routing layers, with each comprising

of multiple domains. Aggregation is then used to combine resources at lower layers/domains and define logical nodes and links at higher layers. For example, the lowest routing layer usually comprises of multiple physical nodes/links arrayed in different domains. Domain entities are also defined to setup, maintain, and release connections, termed as *routing controllers* (RC). Using this framework, the ASTN standard outlines the requirements for all key inter-domain functions, including auto-discovery, auto-provisioning, and auto-restoration. However, this solution only details architectural/interface requirements and does not specify any protocol solutions per say. As a result, other standards bodies (IETF, OIF) have used these guidelines to build ASTN-compliant routing and signaling protocols, see references in [8].

2.2 Multi-Domain Provisioning

To date, a wide range of schemes have been proposed for multi-domain network provisioning under regular (working-mode) conditions. Now most of these solutions have been developed for IP/MPLS bandwidth-routing networks, and only a subset of them have treated more specialized optical WDM networks. Furthermore, many of these strategies are heuristic-based and use graph-theoretic algorithms to implement distributed "on-demand" provisioning. As such, these solutions can readily be adapted in real-world settings by using the existing multi-domain protocol frameworks outlined in Section 2.1.

Overall, multi-domain provisioning requires some level of "global" state visibility to implement (lightpath) route computation. However, as noted in Chapter 1, it is very difficult to achieve full topology and resource exchange across domains due various reasons, i.e., including scalability, inter-carrier privacy and competitiveness, and policy constraints. In light of this, most existing multi-domain schemes implement some type of inter-domain state "peering" exchange. For example, many

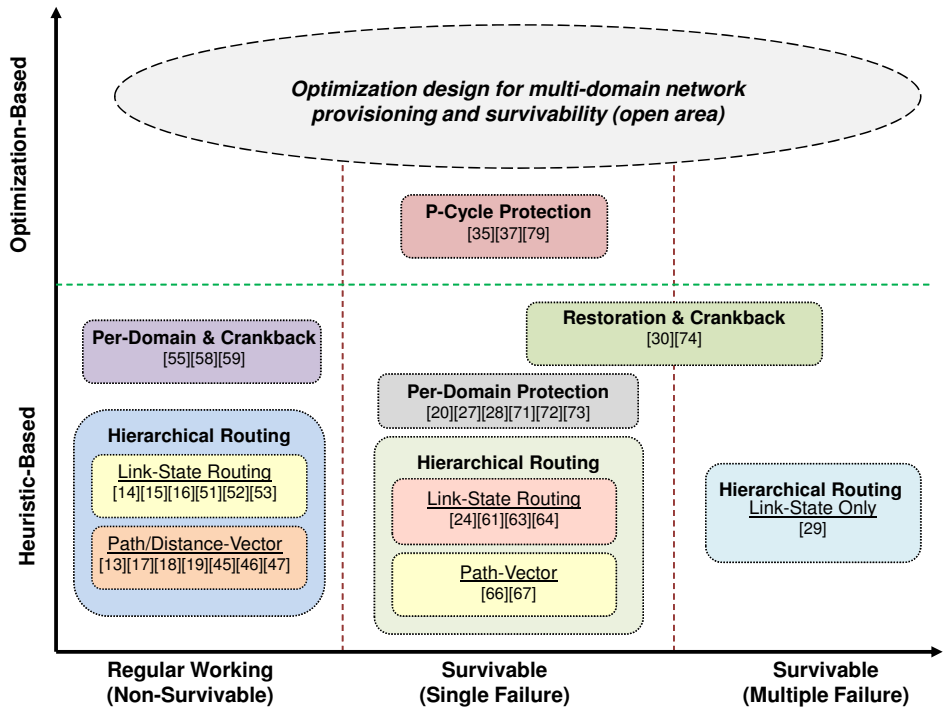

Figure 2.2: Summary of multi-domain provisioning and survivability research

solutions leverage *hierarchical* routing protocols to distribute dynamic domain state, i.e., path/distance-vector and link-state routing setups [23]. These solutions basically try to achieve a tradeoff between information (topology, resource state) accuracy and scalability/privacy concerns. Meanwhile other solutions have also proposed iterative "domain-to-domain" computation techniques to minimize the need for dynamic routing exchange. These main categories are now reviewed, with a particular focus on multi-domain optical WDM networks, see also summary in Figure 2.2.

2.2.1 Hierarchical Distance/Path Vector Routing

As mentioned in Section 2.1, BGP is one of the most widely-used inter-domain routing protocols today. However, basic "vanilla" BGP only propagates best route information and does not provide any *traffic engineering* (TE) capabilities, i.e., QoS support. To address these shortcomings, a host of proposals have outlined extensions to advertise *multiple* path routes to a destination domain, i.e., contingent to specific metrics/constraints such as bandwidth, delay, cost, etc. For the most part these contributions are strictly focused on IP/MPLS packet-routing networks, and not optical WDM networks per say, see references in [8] for further details. However, some further efforts have tried to extend path and distance-vector routing designs to multi-domain optical network settings. Consider some details here.

The work in [13] presents one of the first proposals for *optical-BGP* (O-BGP) to support lightpath RWA across domains. Namely, this is done by augmenting BGP update messages to carry additional wavelength state information for path routes. However, since BGP does not perform any resource reservation, new signaling messages are also introduced to reserve wavelength resources. A similar approach in [17] also outlines a path-vector solution to advertise multiple destination domain routes, i.e., termed as the *constraint-based optical path-vector routing protocol* (COPRP). Separate signaling and reservation protocols are also proposed here, but performance evaluation studies are not conducted.

Some additional studies have also analyzed the performance of path-vector routing for multi-domain RWA. For example, the scheme in [18, 19] advertises multiple destination routes and uses adaptive filtering to predict the number of available wavelengths on each route. Results show very good blocking reduction versus shortest-path O-BGP. Nevertheless, disseminating multiple BGP routes can yield notably higher messaging overheads, and this may pose additional scalability concerns,

see [45, 46]. Finally, an alternate (non-BGP) distance-vector routing scheme is also presented in [47] to support multi-domain RWA. Here border OXC nodes maintain alternate "next-hop" route information and use distributed signaling procedures to concatenate and build end-to-end lightpath sequences. This solution also incorporates physical layer impairment concerns in the provisioning process. The scheme is then analyzed using simulation, but detailed comparisons with other (path-vector) strategies are not presented.

2.2.2 Hierarchical Link-State Routing

Hierarchical link-state routing has also been used to disseminate more complete global topological and resource views and help improve multi-domain provisioning. In particular, this approach condenses domain-level link-state information, and then propagates it across domain boundaries to build "skeleton" abstract views. The overall aim here is to improve scalability and also preserve intra-domain (inter-carrier) privacy. Now earlier studies have proposed several hierarchical link-state routing designs for multi-domain IP/MPLS networks. Most of these solutions use some form of *topology abstraction* to reduce domains into various standardized configurations, i.e., such as simple node, star, tree, and full-mesh [11, 48–50]. For example, simple node abstraction condenses a domain into a single node and provides no intra-domain visibility. By contrast, the other schemes provide better visibility by reducing domains to their border nodes inter-connected by an appropriate set of abstract links. In particular, full-mesh abstraction reduces a domain to a mesh of "abstract" links between all of its border nodes, i.e., $O(N^2)$ overheads for N border nodes.

Building on the above, researchers have also used topology abstraction to condense wavelength and converter state in optical WDM networks. For example, [14] develops an aggregation scheme to summarize wavelength and delay information on

(domain-internal) lightpath routes between border nodes. Next, "skeleton" loose route computation algorithms are outlined to use this information to compute inter-domain sequences, with added bypass paths for handling state inaccuracies. Detailed simulations show very good blocking performance, even when compared with "flat" global routing. Related work in [15] also develops simple node and full-mesh topology abstraction schemes for all-optical WDM domains and outlines improved relative-change triggering policies for inter-domain dissemination, i.e., to build active inter-domain link-state routing databases. Novel graph-based RWA schemes are then introduced to compute "skeleton" inter-domain routes in a load-balancing manner, and these routes are expanded using RSVP-TE signaling. Overall findings show that full-mesh abstraction gives the lowest blocking, albeit routing overheads are much higher due to the dependency on the square of the number of border nodes. Further work in [16] extends the above to treat more realistic settings with intra/inter-domain wavelength conversion, i.e., by tracking the most congested sub-paths between border node pairs in full-mesh abstraction. Again, results show improved blocking with full-mesh abstraction, i.e., versus basic simple node abstraction.

Finally, some studies have also looked at state dissemination scalability for multi-domain link-state routing. For example, [51] evaluates RWA performance for varying update intervals, and results show that full-mesh abstraction is much more sensitive to larger update intervals. Meanwhile, [52] addresses scalability concerns for full-mesh abstraction and proposes an improved gradient-based triggering policy to extract links with the maximum relative change (for the most-used wavelengths). Results have show very good reduction in routing overheads with almost negligible increases in blocking. In a similar theme [53] uses link load information to monitor a subset of domain-internal links when generating abstract link updates. Again, results show much lower overheads with minimal impact on blocking.

2.2.3 Per-Domain and Crankback Strategies

As noted in Section 2.2, the IETF PCE standard [54] provides a broad framework for inter-domain path computation. Leveraging this, researchers have developed related "per-domain" strategies to compute end-to-end routes in an incremental manner. Namely, domain PCE systems compute their local domain traversing routes and then forward the setup request to the "next-hop" downstream domain PCE (and so on and so forth until the destination domain is reached). However, most of these solutions assume that "next-hop" domains are pre-specified, i.e., fixed, in order to simplify operation. Furthermore, since there can be multiple node/link interconnections between two domains, more elaborate inter-PCE computation schemes have also been proposed using a *backward recursive path computation* (BRPC) approach [21, 55]. These solutions basically use inter-PCE signaling to build a *virtual span path tree* (VSPT) from the destination to the source domains (including border nodes). The source domain PCE then chooses the best "skeleton" route, i.e., sequence of border nodes, from the VSPT based upon a particular metric. Overall, PCE BRPC-based schemes have been well-studied for IP/MPLS networks, see survey in [8].

Leveraging the above, further efforts have adapted the BRPC scheme for optical lightpath provisioning. For example, [56] adds wavelength information to the VSPT to facilitate RWA across multiple all-optical WDM domains. Overall findings show the lowest blocking with *first-fit* (FF) wavelength selection. The work in [57] also presents a similar solution, and results show notably lower lightpath blocking versus the basic BRPC scheme, i.e., without any wavelength state information (albeit with longer setup delays). Carefully note these existing BRPC-based RWA solutions do not address realistic network settings with wavelength conversion at border gateway nodes, i.e., likely needed for *service level agreement* (SLA) enforcement at domain boundaries.

However, one of the key drawbacks with "per-domain" (and BRPC-based) computation is the assumption of pre-specified domain sequences, i.e., versus dynamic domain sequences for the case of hierarchical routing (Sections 2.2.1, 2.2.2). Clearly, this choice will yield sub-optimal inter-domain routes and lead to increased request blocking. As a result, some new BRPC-based schemes have tried to incorporate dynamic (hierarchical routing) information to compute end-to-end domain sequences in IP/MPLS networks, see [8]. Alternatively, others have used multi-domain *crankback* signaling to re-try multiple routes and improve setup success. Namely, crankback schemes rely upon resource setup failure notifications from downstream nodes (domains) to re-direct setup attempts to different upstream nodes (domains). More recently, the RSVP-TE protocol has also been extended to support such functionality [41]. Consider some further details here.

Overall, several studies have looked at crankback signaling operation in multi-domain settings. For example, the scheme in [55] uses crankback in IP/MPLS bandwidth-provisioning networks and checks all egress border nodes in a failed domain before signaling the upstream domain. However, since this strategy implements an exhaustive re-try process, it gives lower success rates and excessive path lengths. Improving upon this, [58] presents another solution that limits the number of intra- and inter-domain crankback attempts and also leverages basic hierarchical state information to drive the crankback process along progressively longer inter-domain routes, i.e., as extracted from BGP routing tables. Furthermore, [59] extends this work to support optical lightpath RWA, and overall findings show much lower blocking rates versus no crankback, i.e., at times even matching or exceeding that of hierarchical link-state routing (albeit with increased setup delays, as expected).

2.3 Multi-Domain Survivability

Failure recovery across multiple network domains is also a major concern for operators. Now in general, network survivability schemes can be classified into two broad categories, i.e., pre-configured protection (dedicated, shared) and post-fault restoration [60]. Namely, the former *proactive* strategies achieve rapid recovery by pre-computing and pre-reserving backup resources at various levels, e.g., link, end-to-end path, or even sub-path/path. However, since pre-allocation gives increased cost and resource usage, the latter schemes simply re-compute alternative paths after failure events, i.e., *reactive* approach. Overall, both of these methodologies are applicable in multi-domain settings and various of related solutions have been studied here, see [8]. Some of these key contributions are now reviewed, see also summary in Figure 2.2.

2.3.1 End-to-End Path Protection

A full range of multi-domain protection strategies have been developed using hierarchical routing to compute diverse end-to-end working (primary) and protection (backup) paths/lightpaths. These schemes basically pre-provision backup resources to achieve fast switchovers and provide guaranteed recovery against single link failures. For example, [61] considers different levels of global state visibility and outlines three dedicated multi-domain protection schemes. The first approach assumes full visibility and is used as an idealized (unrealistic) reference. Meanwhile the second strategy uses simple-node abstraction and computes domain-disjoint routes. Finally the third strategy routes end-to-end primary and backup paths along the same domain sequence and handles intra-domain link failures via localized recovery methods, i.e., no inter-domain link failure recovery. Overall results show that the second domain-disjoint strategy achieves a median performance between the other two

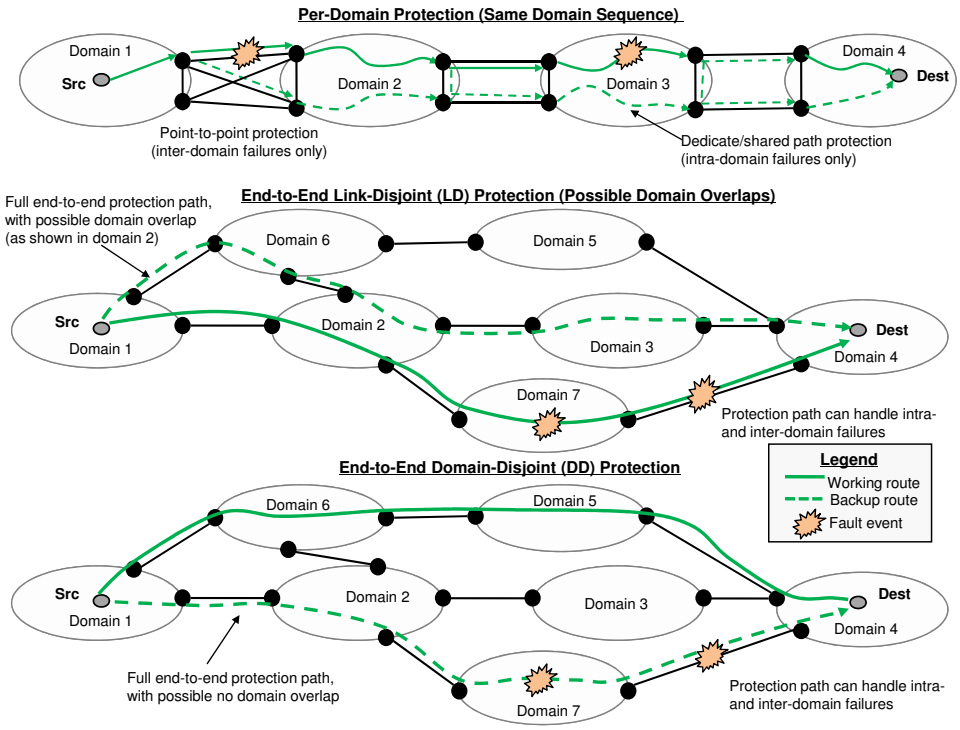

Figure 2.3: Multi-domain disjoint path-pair protection strategies, from [8]

strategies in terms of blocking. However, no details are presented for the disjoint path-pair computation algorithms in [61], i.e., over the global "flat" topology or the abstract topology.

Meanwhile, [24] presents a more advanced dedicated protection scheme that uses a specialized full-mesh topology abstraction scheme (for hierarchical link-state routing). Namely, this approach extracts intra-domain path-level diversity information (between border node pairs) using Suurballe's algorithm [62]. A distributed algorithm is then developed to simultaneously compute link-disjoint path-pairs, i.e., parallel setup. However, no performance evaluation studies are presented here, and the associated topology abstraction routing overheads are prohibitive, on the order of $O(N^4)$, for N border nodes. Hence to better address these scalability concerns, [63]

develops an alternate solution that uses regular simple node and full-mesh topology abstraction (no diversity state) and instead relies upon *sequential* signaling to achieve end-to-end primary/backup path-pair diversity. Specifically, two dedicated loose route protection schemes are proposed at the inter-domain level, i.e., *link-disjoint* (LD) and *domain-disjoint* (DD), see Figure 2.3. Suurballe's algorithm is also applied here to ensure path-pair diversity during intra-domain signaling expansion. As expected, results show much lower blocking with the LD scheme, i.e., since enforcing complete domain separation between routes is very restrictive.

Meanwhile, other studies have also looked at shared protection for multi-domain optical networks running hierarchical link-state routing. The objective here is to improve resource efficiency between backup routes on both intra- and inter-domain links. For example [64] presents novel full-mesh topology abstraction schemes to aggregate specialized domain-internal state relating to residual capacity, backup capacity, etc. This information is then used to develop new sequential and parallel (joint) loose route computation schemes. Overall findings show improved blocking performance with the joint computation approach. Further studies have also extended this work to implement more complex shared sub-path protection, see [65] for complete details. Nevertheless, it may be difficult to achieve resource sharing in practical multi-carrier settings owing to various concerns. Foremost, policy limitations may restrict the level of sharing on inter- and intra-domain links. Moreover, many carriers may be reluctant to run the required topology abstraction schemes, as they may reveal too much information about their internal networks.

Note that some studies have also considered multi-domain path-vector protection [66, 67]. For the most part, these schemes compute and disseminate multiple "diverse" alternate routes to destination domains, i.e., node-disjoint and policy-compliant routes. However these offerings only focus on IP/MPLS packet-forwarding networks and require further extensions to operate in optical WDM settings, see dis-

cussions in [8].

As noted in Chapter 1, multi-failure disaster recovery is also a major concern for network operators. In particular, expansive multi-domain infrastructures are particularly vulnerable to catastrophic stressor events such as earthquakes, floods, cascading power outages, and even malicious destructive attacks. Now it is well understood that such occurrences can yield multiple link and node failures with relatively high levels of spatial and temporal correlation [22, 68]. By contrast, most existing (multi-domain) network recovery schemes are only designed to handle isolated single failures, and hence will be largely ineffective under such catastrophic conditions. Moreover, it is very difficult (and overly costly) to pre-provision protection resources to try to guarantee recovery against multiple randomized failures. As a result, recent efforts have started to focus on more advanced *probabilistic* protection schemes. Consider the details here.

Most probabilistic network recovery schemes have only looked at single domain networks. For example, the work in [69] presents a heuristic scheme to minimize the joint path-pair failure probability of primary/backup routes for networks with multiple independent link failures. However, this simplistic approach does not capture the high levels of spatial correlation between link failures under disaster conditions. Alternatively, [68] develops a new *probabilistic shared risk link group* (p-SRLG) model to capture the spatial failure correlation between links, i.e., as induced by large-scale stressors. Namely, this model augments the well-known *shared risk link group* (SRLG) [70] concept and adds several new parameters, i.e., including an occurrence probability for each SRLG region and conditional failure probabilities for each node and link within a given region. Leveraging this framework, the authors then develop some novel optimization- and heuristic-based schemes to minimize the joint failure probabilities of primary/backup path-pairs (see also Section 2.3.5). Carefully note that this work assumes that all p-SRLG events are mutually-exclusive–a realistic

assumption given the rare nature of most disaster events.

Finally, [29] extends the p-SRLG model to multi-domain settings and presents one of the first solutions for probabilistic multi-failure protection across domain boundaries. In particular, this solution assumes a hierarchical link-state routing design, and first outlines a novel "risk-aware" topology abstraction scheme to extract critical domain-level link vulnerability state. This information is then used to develop two "risk-based" loose route path-pair computation schemes. In particular, the first scheme strictly focuses on path-pair risk minimization but yields overly lengthy (resource intensive) routes. Conversely, the second scheme achieves a better tradeoff by jointly incorporating both TE load-balancing and risk minimization concerns. Overall results show that the joint scheme achieves very good disaster recovery performance, closely tracking the reliability of the pure risk minimization heuristic. At the same time, this scheme also gives very competitive request blocking and resource utilization performance, i.e., closely matching pure TE load-balancing strategies.

2.3.2 Per-Domain Protection

Per-domain protection schemes have also been proposed to implement path protection along the same domain sequence as primary routes, i.e., extending upon per-domain path computation strategies (Section 2.2.3). However, the focus here is only on inter-domain link failure recovery as all domain-internal faults are assumed to be handled at the intra-domain level to achieve rapid recovery (usually under 100 ms). Overall this approach is somewhat flexible since operators can use different recovery methods within their own domains and thereby avoid any inter-domain (end-to-end) signaling requirements across administrative boundaries. However, since domain sequences are fixed here, per-domain protection strategies will need adequate levels of domain-to-domain connectivity, e.g., pre-engineered dual/multi-homing setups.

Clearly, this will pose increased costs for network operators. In addition, these schemes are less resource-efficient as they ignore load information on inter-domain links and simply relegate primary/backup routes to the same fixed domain sequences. Consider some recent contributions in this space.

The work in [27] presents a dedicated per-domain protection scheme for IP/MPLS networks based upon various inter-domain link interconnection strategies. Furthermore, [28] extends this solution by using 1:1 and full-mesh interconnection between domains, and proposes two different signaling strategies to setup end-to-end primary/backup routes. Namely, the first scheme uses sequential signaling to avoid overlaps on back routes, whereas the second scheme implements parallel signaling expansion to avoid intra-domain trap topology concerns (using Suurballe's algorithm). Expectedly, the latter scheme gives better setup success rates and lower overall costs.

Meanwhile, [71] extends the well-studied BRPC scheme (Section 2.2.3) to compute disjoint virtual spanning trees between the source and destination domains. However the performance of this scheme is not analyzed here. Meanwhile, [20] proposes several distributed setup schemes for per-domain protection in multi-domain GMPLS networks. Specifically, two different sequential signaling strategies are presented here based upon the RSVP-TE and PCEP protocols, respectively. Namely the PCEP-based solution extends the BRPC-based approach to reverse compute a protection VSPT. The appropriate border nodes for the protection path are then selected from this tree. However, owing to trap topology concerns, sequential signaling schemes may not find a diverse path-pair even if one exists. Hence, more complex parallel setup schemes are also proposed using inter-PCE path computation to jointly expand primary/backup routes, see [20] for details. These schemes are then analyzed using simulation, and results indicate the lowest blocking with parallel setup, especially for smaller numbers of inter-connecting border nodes.

Finally, some studies have tried to incorporate shared protection with per-domain

protection, albeit at the intra-domain level. For example, [72] assumes that all domains have at least two connecting border nodes and that all inter-domain routes are pre-computed using *integer linear programming* (ILP) optimization. Shared protection is then implemented to improve resource efficiency between the primary/backup routes within a domain, i.e., by solving an ILP formulation. As expected, the results show much better intra-domain resource efficiency with ILP-based shared protection. Meanwhile, [73] presents another protection strategy for optical WDM networks with full wavelength conversion and redundant inter-domain link connectivity. Namely, a load-balancing heuristic is used to compute shared intra-domain protection routes, and results indicate very good resource efficiency. Note that further extensions for resource sharing on inter-domain links can also be considered here, i.e., 1:1 type protection schemes.

2.3.3 Post-Fault Restoration

In general, protection schemes pose sizable (routing, computational) complexities in multi-domain settings and also entail increased service costs due to backup pre-provisioning. By contrast, post-fault restoration schemes are much more efficient as they simply try to re-compute affected routes (or parts of affected routes) after failure events. However, these schemes are more latent and cannot provide any recovery guarantees, e.g., as per single failure protection schemes. Nevertheless, restoration schemes are still very appealing for mid-tier services and can also be effective for handling multiple random failures. As a result, various studies have looked at multi-domain post-fault restoration using crankback signaling recovery, and some of these contributions are now reviewed.

The authors in [55] outline a simple restoration scheme for multi-domain IP/MPLS networks that uses crankback operation to test random upstream border nodes and

initiate failed route recovery, i.e., intermediate restoration. However, simulation results for isolated single link failures show lower restoration success rates versus a modified "per-domain" protection approach. To resolve these concerns, [30] proposes an improved solution (for optical WDM networks) that combines intra- and inter-domain crankback recovery using standardized RSVP-TE signaling extensions [41]. In particular, two strategies are proposed here based upon intermediate and end-to-end crankback (at the inter-domain level). Namely, the former approach tries to re-use non-failed portions of affected routes to speed up recovery, whereas the latter approach tries to improve resource efficiency by re-computing new routes. Additional features are also added to track failed crankback links, and overall results for isolated single-link failures show better recovery with the end-to-end crankback scheme. Finally, [74] adapts this solution for multi-domain IP/MPLS networks and tests its performance under more challenging regional/disaster scenarios, i.e., multiple correlated failures. Again, the findings show improved recovery with end-to-end restoration, albeit recovery rates decline with increasing stressor sizes.

2.3.4 p-Cycle Protection

The *pre-configured cycle* (p-cycle) concept was originally proposed in [75] to provide rapid "ring-like" recovery over generalized optical mesh topologies. This approach works by providing a restoration path for every span on a cycle as well as every span straddling a cycle, i.e., end-points on different cycles. Now a wide range of studies have looked at extending and improving the basic p-cycle concept, see survey in [76]. In fact, some researchers have even applied this approach for multi-domain protection. Consider the details here.

The work in [35] presents one of the first studies on multi-domain p-cycle protection and decomposes the problem into two inter-dependent sub-problems, i.e.,

intra-domain and inter-domain p-cycle design. In particular, the latter uses simple-node topology aggregation to summarize the global network and then routes p-cycles over the abstract topology to recover from inter-domain link failures. Next, the required abstract link (domain-traversing) connections are identified and provisioned (protected) by routing p-cycles over the individual domain topologies. In particular, the *capacitated iterative design algorithm* (CIDA) [77] is used to resolve the p-cycles at both the inter-/intra-levels and improve sharing on backup routes. Overall results show that p-cycle protection improves availability but also yields higher resource consumption than traditional protection strategies. Meanwhile [37] develops another two-level hybrid solution for multi-domain p-cycle setup. Namely, an ILP approach is used to minimize spare capacity usage on inter-domain links using full-mesh topology abstraction. Next, individual domain-traversing sub-paths are protected using the *failure independent path protecting* (FIPP) p-cycle approach in [78]. The proposed ILP model is then mapped to a linear relaxation problem, which is then solved using column generation techniques to reduce the number of variables. In addition, the work in [36] also proposes some practical signaling extensions to the RSVP-TE and PCEP protocols to implement p-cycle setup and recovery.

In general, optimizing p-cycle routes over full global (abstract) topologies entails higher computational complexity and can yield overly-lengthy routes. In turn, this approach can lead to increased resource inefficiency and higher signal impairments across optical WDM domains. To address this challenge, the authors in [79] use spectral clustering techniques to partition domains into separate clusters, i.e., network decomposition into multiple smaller "sub-multi-domain" networks. Localized p-cycles protection techniques are then applied over these smaller networks to improve computational efficiency. Overall findings here show much lower computational times along with a slight increase in redundancy, i.e., 2-4% range.

Nevertheless, p-cycle protection schemes mostly focus on link-level recovery and

embody a pre-computed (off-line) solution. As such, this approach offers very little demand selectivity as all flows on a p-cycle span end up receiving full protection. This limitation contrasts with end-to-end connection protection and/or restoration strategies which can be applied to individual user demands. Furthermore, p-cycle design in multi-domain settings also requires detailed inter-domain skeletal state as well as full domain-internal topological state. Clearly, this information will be difficult to obtain in generalized multi-carrier settings, further limiting the applicability of these solutions. Finally, nearly all studies on p-cycle recovery have only looked at single-link failures. As such, these designs are very vulnerable to larger disaster scenarios–a key deficiency that is also noted in [80].

2.3.5 Optimization Design Strategies

Overall, most studies on optical RWA optimization design have only looked at single-domain networks. For example, [81] presents an ILP formulation to maximize the number of lightpaths established in wavelength-routed optical networks with a-priori demands. This model is then solved using *linear programming* (LP) relaxation and rounding techniques. Recent work in [82] also presents a more advanced *multi-objective* ILP formulation, i.e., to jointly improve throughput, decrease resource utilization, and achieve load-balancing. This model is then solved for the NSF network topology, and its results compared with a heuristic solution using the Ford-Fulkerson widest path algorithm. Overall findings here show improved load-balancing and reduced resource consumption with the ILP-based approach.

Meanwhile, other studies have also looked at optimization design for link-disjoint path protection in single-domain settings. For example, [83] shows that lightpath protection with wavelength continuity constraints is an NP-complete problem and then proceeds to develop an ILP formulation. However, this optimization model is

not solved, and two heuristic strategies are proposed instead, i.e., with one computing link-disjoint path-pairs and then searching for available wavelengths and the other scanning wavelengths for link-disjoint path-pairs. Meanwhile, [84] studies path protection for more generalized SRLG scenarios. Here a SRLG represents any set of links with a common vulnerability [70], and this definition can be used to capture regional disaster-type events as well. Overall [84] shows that the SRLG-disjoint path-pair routing problem is also NP-complete and related ILP formulations are largely intractable for most medium-large size networks. As a result, some alternate heuristic solutions are proposed to help avoid trap topologies which are shown to be much more prevalent in SRLG routing scenarios.

Finally, as mentioned in Section 2.3.1, [68] studies diverse path-pair routing with probabilistic regional failures based upon a *probabilistic SRLG* (p-SRLG) concept. In particular, an *integer non-linear programming* (INLP) optimization model is proposed to try to minimize the joint failure probability of primary/backup path-pairs. However, owing to extreme INLP computation complexity, an alternate ILP approximation is also proposed by simplifying/omitting high order terms and using variable substitution. The results from this approximation are then compared against some heuristic strategies developed by the authors in [68], with the former yielding lower path-pair failure probabilities.

2.4 Open Challenges

In summary, multi-domain provisioning and survivability is a very challenging problem area, and one that has attracted significant attention to date. Nevertheless, even though a wide range of solutions have been proposed here, for the most part these schemes use sub-optimal heuristics-based methodologies. Hence it is very difficult to gauge the ideal achievable performance in multi-domain settings, and most

studies simply use other heuristic schemes for comparison purposes. In light of the above, optimization-based techniques provide a very promising avenue for expanding the work in this area and developing some performance bounds. Furthermore, the related findings from such efforts can also be used to develop improved practical solutions, i.e., either by developing new heuristic strategies or adapting/applying optimization methods themselves.

However, most existing studies on (optical) network provisioning optimization have only looked at single-domain settings. Although a few efforts have applied such techniques for multi-domain p-cycle design, these schemes are not very selective and cannot handle multiple link failures. As a result, *multi-domain* optimization remains an open problem area with much scope for new work. Foremost, there is a pressing need to develop realistic models for (non-survivable) working-mode and (survivable) protection-mode provisioning to effectively capture multi-domain visibility concerns. Furthermore, there is a growing need to model and analyze large-scale disaster scenarios as well, given the increased vulnerability of expansive multi-domain infrastructures to such events. Overall, these challenges form the main motivations for this dissertation research.

Chapter 3

Mutli-Domain Lightpath Provisioning Optimization

As surveyed Chapter 2, most existing studies on multi-domain network provisioning use graph-based heuristics and analyze their performance against other heuristics via network simulation. Although these efforts represent some important contributions in their own right, they still do not provide any detailed analytical treatments per say. However many operators are interested in bounding the "ideal" achievable performance in multi-domain settings in order to gauge the effectiveness of various heuristics-based strategies. It is here that optimization-based solutions offer much potential, and these strategies have already been applied in single-domain networking studies (see discussions in Section 2.3.5). However, the further extension of these formulations to multi-domain settings has not yet been considered.

Now simply re-applying single-domain optimization models across multi-domain topologies is not very feasible for several key reasons. Foremost, multi-domain networks have much higher node and link counts, and this will pose immense computational scalability challenges, i.e., formulations become intractable for even moderate

numbers of domains. Moreover, optimizing lightpath routes over "flat" multi-domain topologies implies global network state, an unrealistic assumption that is not reflective of real-world settings.

Hence in order to address these shortcomings, this chapter presents a novel optimization-based formulation for multi-domain lightpath provisioning under regular working (non-failure) conditions. The key goal here is to bound network performance in terms of blocking rates and resource utilization, and to also provide a benchmark against which to test and improve heuristics-based strategies. Overall, the proposed framework models realistic multi-domain environments in which the provisioning entities only have partial skeleton views of global network state, i.e., to reflect practical hierarchical routing designs. Based upon this, the formulation pursues a "two-stage" optimization strategy to resolve routes at the inter- and intra-domain levels pursuant to multiple TE objectives. Note that this solution is also used as a basis to develop further optimization models for multi-domain survivability in Chapters 4 and 5. The requisite notation is now introduced followed by a detailed presentation and analysis of the proposed optimization scheme.

3.1 Optimization Formulation

A novel ILP optimization model is now presented for multi-domain lightpath provisioning. This formulation assumes full *a-priori* knowledge of inter-domain user requests and mimics the overall hierarchical routing and provisioning process via a two-stage optimization process, see Figure 3.1. Namely, skeleton path sequences are first derived for all requests over a global abstract topology subject to specific TE objectives and constraints. In particular, this abstract topology uses full-mesh abstraction to represent domains as a mesh of links between their respective border nodes. Next, the individual domain-traversing lightpath segments (sub-paths) are

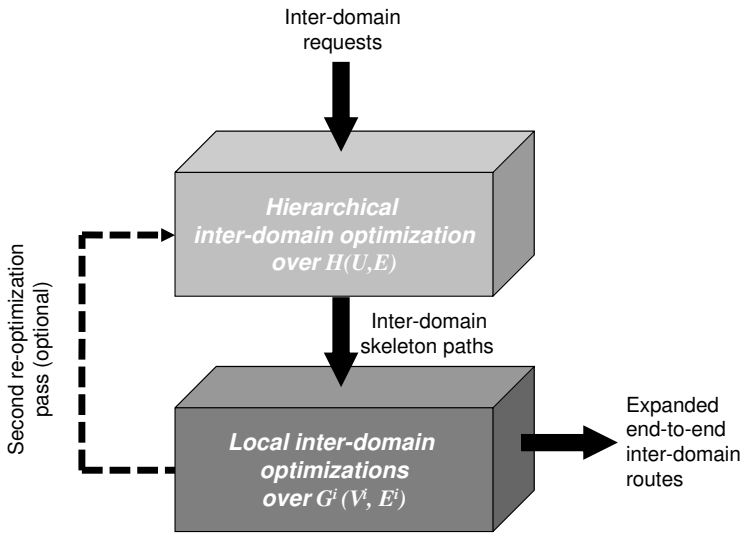

Figure 3.1: Two-stage multi-domain optimization solution

extracted from the skeleton paths and then optimized over their respective local domain topologies. Overall, this two-stage optimization approach is much more feasible as it is difficult (impossible) to solve a single ILP formulation over a larger "flat" single domain network comprising of all domain nodes and links, i.e., unrealistic case of global state.

Now in order to model realistic environments, it is assumed that all domains are internally-transparent but support full opto-electronic wavelength conversion at their boundaries. As a result intra-domain OXC nodes have all-optical switching fabrics and do not use any converters, see Figure 3.2. Conversely, border nodes have added opto-electronic conversion stages which are only used for inter-domain lightpath requests, as shown in Figure 3.3. Indeed, similar considerations have also been made in other studies [16, 85], owing to the increased geographic distances (signal degradation) across multiple domains and the need to provide bit-level SLA

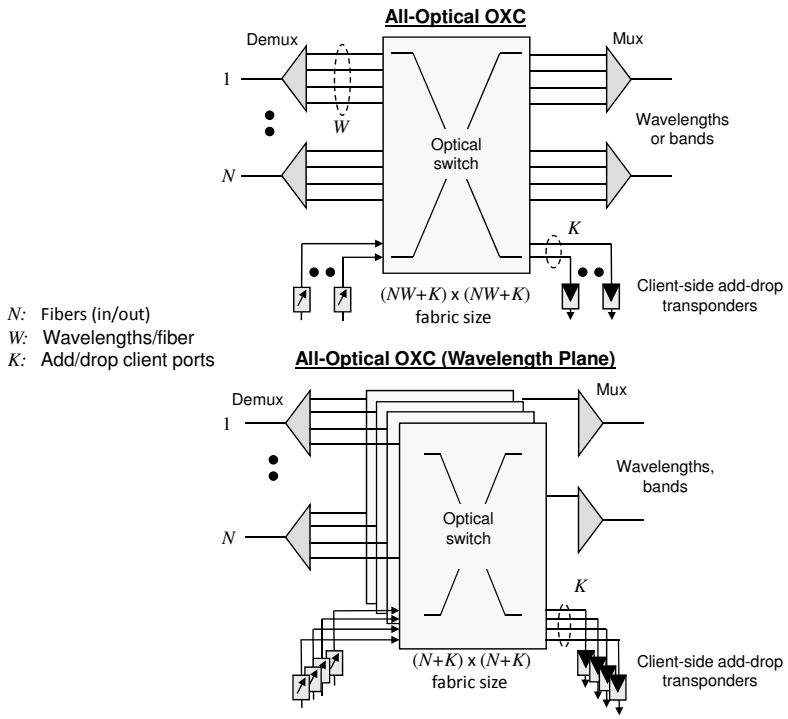

Figure 3.2: Overview of all-optical optical cross-connect (OXC) node

monitoring at boundary nodes. The requisite notation is now introduced followed by the detailed optimization objectives and constraints.

3.1.1 Notation Overview

Before presenting the multi-domain optimization model, the requisite notation is first introduced. Consider a multi-domain WDM network with D domains, with the i-th domain having b^i border nodes, $1 \leq i \leq D$. Here each i-th domain is represented by a sub-graph, $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$, where $\boldsymbol{V}^i = \{v_1^i, v_2^i, ...\}$ are the physical domain OXC nodes and $\boldsymbol{L}^i = \{l_{km}^{ii}\}$ are the physical intra-domain links, i.e., l_{km}^{ii} is the intra-domain link between OXC nodes v_k^i and v_m^i $(1 \leq i \leq D, 1 \leq k, m \leq n^i)$. Without loss of

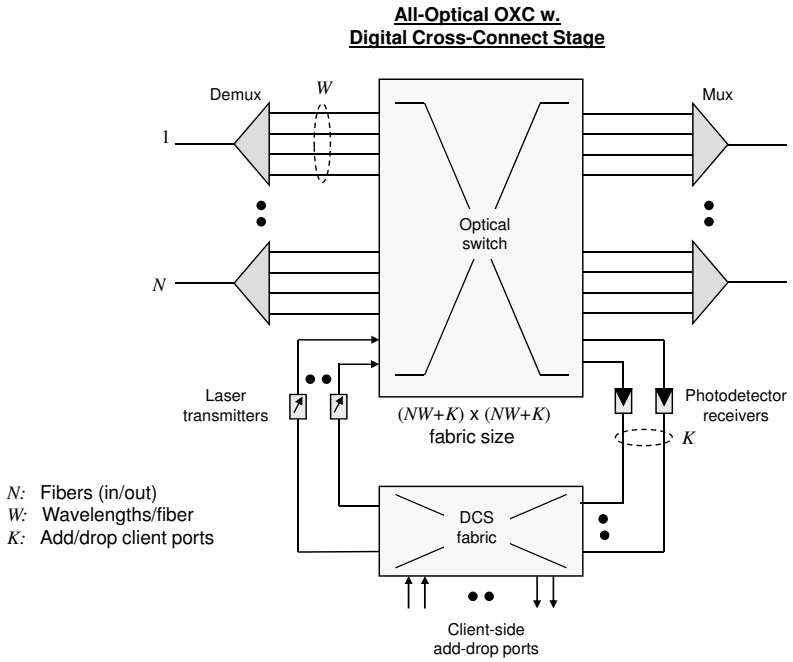

Figure 3.3: Overview of opto-electronic OXC node w. digital cross-connect stage

generality, all intra-domain links are assumed to be bi-directional with maximum wavelength capacity C_1. In addition the border OXC nodes in domain i are also denoted by the set $\boldsymbol{B}^i \subseteq \boldsymbol{V}^i$, i.e., where $|\boldsymbol{B}^i| = b^i$.

Now assuming hierarchical multi-domain routing with full-mesh abstraction, an associated "higher-level" *abstract* topology is also defined comprising of all border OXC nodes and their interconnecting physical (inter-domain) and abstract (intra-domain) links. Specifically, this topology is represented by the graph $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$, where $\boldsymbol{U} = \sum_i \{\boldsymbol{B}^i\}$ is the global set of border nodes and $\boldsymbol{E}$ is the global set of links, i.e., $\boldsymbol{E} = \boldsymbol{E}_{phy} \cup \sum_i \boldsymbol{E}^i_{mesh}$, where $\boldsymbol{E}_{phy}$ is the set of physical inter-domain links and $\boldsymbol{E}^i_{mesh}$ is the set of abstract intra-domain links for domain i ($1 \leq i \leq D$). Namely, $\boldsymbol{E}_{phy} = \{e^{ij}_{km}\}$ where e^{ij}_{km} is the physical link interconnecting OXC node v^i_k in domain

i with OXC node v_m^j in domain j ($1 \leq i, j \leq D$, $1 \leq k \leq b^i, 1 \leq m \leq b^j$). Without loss of generality, it is also assumed that all physical inter-domain links in $\boldsymbol{E}_{phy}$ have maximum wavelength capacity C_2. Meanwhile, the abstract link set for domain i is given by $\boldsymbol{E}^i_{mesh} = \{e^{ii}_{jk}\}$, where e^{ii}_{jk} is the abstract link between border nodes v^i_j and v^i_k and $|\boldsymbol{E}^i_{mesh}| = b^i(b^i - 1)$ ($1 \leq i \leq D$, $1 \leq j, k \leq b^i$).

To further illustrate this notation a sample abstract graph, $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$, is shown in Figure 3.4 for a 3-domain topology. For example, here topology abstraction reduces the physical topology for domain 2, $\boldsymbol{G}^2(\boldsymbol{V}^2, \boldsymbol{L}^2)$, to a full-mesh sub-graph consisting of 3 border nodes (v_1^2, v_2^2, v_3^2) and 6 abstract links (e_{12}^{22}, e_{21}^{22}, e_{13}^{22}, e_{31}^{22}, e_{23}^{22}, e_{32}^{22}).

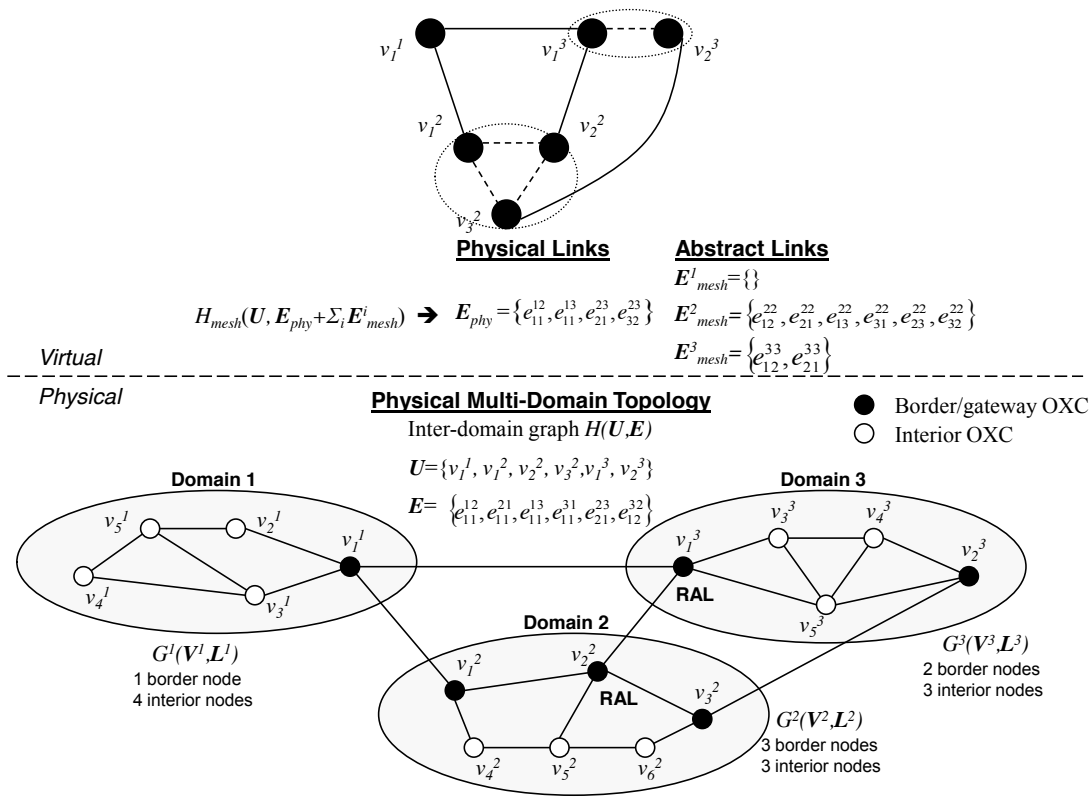

Figure 3.4: Full-mesh topology abstraction and notation overview

3.1.2 Constraints and objectives

Based on the above notation, an ILP formulation is now presented for the multi-domain lightpath provisioning problem. Now as mentioned earlier, the proposed solution (Figure 3.1) applies optimization at both the inter- and intra-domain topology levels, i.e., over the graphs $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ and $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$, respectively. Consider the inter-domain level first. Here, the assumption of full opto-electronic conversion at border OXC nodes clearly obviates the need for wavelength selection on skeleton routes, i.e., only path selection is required. However, wavelength selection is required at the intra-domain level, and hence the proposed formulation decouples lightpath route selection from wavelength assignment. Namely, an ILP optimization is first used to compute the intra-domain routes, and then an appropriate wavelength selection policy is used to assign the wavelengths, i.e., akin to other single-domain lightpath RWA optimization studies in [32, 33, 82, 86, 87]. Overall, this approach allows similar TE objectives to be pursued at each level, i.e., global inter-domain and local intra-domain. As a result only the inter-domain optimization model is presented here for brevity's sake, i.e., for abstract graph $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$.

Now consider an a-priori set of inter-domain lightpath demands, denoted by the set of 3-tuples $\{(s_n, d_n, r_n)\}$, where n represents the request index ($1 \leq n \leq N$), s_n is the source OXC node, d_n is the destination OXC node, and r_n is the number of required wavelengths, i.e., $s_n \in \boldsymbol{V}^i$, $d_n \in \boldsymbol{V}^j$, $i \neq j$ and $r_n \leq C_1, C_2$. Also, let f_n denote the number of wavelengths allocated to request n, x_{km}^{nij} denote the number of wavelengths routed over link e_{km}^{ij} for request n, and α denote the *maximum link utilization* (MLU) allowed, i.e., in order to prevent link saturation. Furthermore, without loss of generality, assume that all abstract links in $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ also have capacity C_2. Using these variables, a multi-objective function, F, is defined to pursue

several objectives as follows:

$$\text{Max } F = w_1 \sum_{n \in \boldsymbol{N}} f_n - w_2 \sum_{n \in \boldsymbol{N}} \sum_{e_{km}^{ij} \in \boldsymbol{E}} x_{km}^{nij} - w_3 \alpha = w_1 F_1 + w_2 F_2 + w_3 F_3 \tag{3.1}$$

where

$$F_1 = \sum_{n \in \boldsymbol{N}} f_n \tag{3.2a}$$

$$F_2 = - \sum_{n \in \boldsymbol{N}} \sum_{e_{km}^{ij} \in \boldsymbol{E}} x_{ij}^{n} \tag{3.2b}$$

$$F_3 = -\alpha \tag{3.2c}$$

and w_1, w_2, and w_3 are fractional weighting factors that sum to unity, i.e., $w_1 + w_2 + w_3 = 1$, and subject to the following constraints:

$$\sum_{(j,m):e_{km}^{ij} \in \boldsymbol{E}} x_{km}^{nij} - \sum_{(j,m):e_{mk}^{ji} \in \boldsymbol{E}} x_{mk}^{nji} = \begin{cases} f_n;\ if\ v_k^i = s_n \\ -f_n;\ if\ v_k^i = d_n \\ 0; otherwise \end{cases} \tag{3.3}$$

$$\sum_{n \in \boldsymbol{N}} x_{km}^{nij} \leq \alpha C_2; e_{km}^{ij} \in \boldsymbol{E} \tag{3.4}$$

$$f_n \leq r_n; n \in \boldsymbol{N} \tag{3.5}$$

$$x_{km}^{nij} \in \{0, 1, 2, ...\}; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{3.6}$$

$$f_n \in \{0, 1, 2, ...\}; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{3.7}$$

$$0 \leq \alpha \leq 1 \tag{3.8}$$

Overall, F comprises of three parts and is similar to the objective function used in [82] for single-domain optical WDM networks (shown to give improved load balancing and decreased resource consumption). Overall, the key aim here is to achieve a weighted balance between maximizing the aggregate throughput, F_1, minimizing resource consumption, F_2, and achieving load balancing, F_3. Namely, Eq. (3.2a) represents the total throughput for the given requests and network topology, Eq. (3.2b)

represents the negative of total resource consumption, and Eq. (3.2c) is the negative of MLU. Meanwhile Eqs. (3.3)-(3.8) specify the ILP model contraints. Specifically, Eq. (3.3) represents the flow conservation constraint between the incoming and outgoing flows at each (border) node in the abstract graph. Meanwhile, Eq. (3.4) restricts the total *relative* traffic carried on a link to below the MLU value, i.e., less than αC_2. Also, Eq. (3.5) ensures that the number of allocated wavelengths is less than the number of requested wavelengths for each request. Finally, Eqs. (3.6) and (3.7) represent integrality constraints, and Eq. (3.8) restricts the MLU for all links to be below α.

Overall, the above optimization formulation has a total of $(N + N|\boldsymbol{E}|)$ variables and $O(N(|\boldsymbol{U}| + |\boldsymbol{E}|))$ equations, see Table 3.1. Namely, each request has one flow variable, f_n, as well as other flow-specific link variables, i.e., x_{km}^{nij}. Now the proposed formulation has to be solved for the inter-domain abstract topology, $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$, as well as all intra-domain topologies, $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$. Hence the overall computational complexity here will be dominated by the topology with the maximum variable/equation count, as given by $\max_i\{N(|\boldsymbol{U}| + |\boldsymbol{E}|), N_i(|\boldsymbol{V}^i| + |\boldsymbol{L}^i|)\}$, where N_i is the number of local requests for domain i. Carefully note that if the above formulation is applied over a "flat" multi-domain network comprising of all intra-/inter-domain node and links, i.e., unrealistic case of global network state, the resultant variable count will be much higher (intractable).

3.2 Solution Approach

As mentioned in Section 3.1, the ILP solution follows a "top-down" strategy and first optimizes skeleton inter-domain routes for all requests over $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$. These routes are then used to identify the traversed set of abstract links, from which the required "all-optical" intra-domain requests are generated to drive the second optimization

Equation	$\boldsymbol{H}(\boldsymbol{U},\boldsymbol{E})$	$\boldsymbol{G}^i(\boldsymbol{V}^i,\boldsymbol{L}^i)$
Eq.(3.3)	$N\lvert\boldsymbol{U}\rvert$	$N_i(\lvert\boldsymbol{V}^i\rvert)$
Eq.(3.4)	$\lvert\boldsymbol{E}\rvert$	$\lvert\boldsymbol{L}^i\rvert$
Eq.(3.5)	N	$\boldsymbol{N}_i$
Eq.(3.6)	$N\lvert\boldsymbol{E}\rvert$	$N\lvert\boldsymbol{L}^i\rvert$
Eq.(3.7)	N	N_i
Eq.(3.8)	1	1
total	$O(N(\lvert\boldsymbol{U}\rvert+\lvert\boldsymbol{E}\rvert)$	$O(N_i(\lvert\boldsymbol{V}^i\rvert+\lvert\boldsymbol{L}^i\rvert))$

Table 3.1: Number of equations

stage, i.e., solving ILP formulations over the individual sub-graphs $\boldsymbol{G}^i(\boldsymbol{V}^i,\boldsymbol{L}^i)$, Section 3.2.2. Finally the entire end-to-end lightpath sequences are resolved by concatenating all of the intra-domain segments (with the same flow index) with their respective inter-domain links in $\boldsymbol{H}(\boldsymbol{U},\boldsymbol{E})$. Without loss of generality, it is also assumed that all multi-domain lightpath requests are for a single wavelength, i.e., $r_n = 1$, to simplify local domain request indexing.

3.2.1 Hierarchical Inter-Domain Optimization

Consider the case of skeleton route optimization over the global abstract topology, $\boldsymbol{H}(\boldsymbol{U},\boldsymbol{E})$. Since this topology represents a compact "domain-level" summary of the entire network, it is reasonable to assume that this problem can be be solved in reasonable time using a standard workstation computer. Now given a valid solution here, the computed number of unit-flows (wavelengths) routed along on each *physical* inter-domain link and abstract intra-domain link can be determined. Specifically, let the outputted (computed) values for x_{km}^{nij} be denoted by $X_{km}^{nij} = \{0,1\}$, i.e., due to single-wavelength request assumption $r_n = 1$. Using these values, the number of requests routed on any given (inter-domain physical or intra-domain abstract) link can be determined by simply summing up the associated X_{km}^{nij} values over all requests, i.e., $X_{km}^{ij} = \sum_{n\in N} X_{km}^{nij}$.

Now consider an abstract link e_{km}^{ii} in domain i. Here the total number of wavelength routed over this link will be equal to the number of local intra-domain sub-path requests that must be set up between its respective ingress/egress border nodes, i.e., v_k^i and v_m^i. Hence these local domain requests can be further grouped into the set $\{(v_k^i, v_m^i, 1)\}$ and then used to drive the second stage of the optimization, detailed next in Section 3.2.2. Carefully note that above numbering scheme assigns the same request indices to the local domain requests as those for their corresponding inter-domain requests, i.e., essentially matching which "end-to-end" lightpath each sub-path request belongs to. Indeed, this is possible due to the assumption of single-wavelength requests, i.e., as the optimization solution may otherwise yield multiple routes for multi-wavelengths requests.

An example of hierarchical skeleton path computation and sub-path request generation is shown for a sample 3-domain network in Figure 3.5. In particular, here there are two multi-domain lightpath requests which need to be routed, request 1 $(v_2^1, v_3^3, 1)$ and request 3 $(v_2^2, v_3^3, 1)$. Now assume that the ILP solution returns valid skeleton paths over $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ for both requests, as shown in Figure 3.5 via the dashed colored lines, i.e., $e_{21}^{12} - e_{13}^{22} - e_{31}^{23} - e_{13}^{33}$ (for request 1) and $e_{23}^{22} - e_{31}^{23} - e_{12}^{33}$ (for request 2). From these skeleton route sequences, the associated intra-domain sub-path requests can be generated for each domain, as shown in Table 3.2.

	domain 1	domain 2	domain 3
request 1	-	$(v_1^2, v_3^2, 1)$	$(v_1^3, v_3^3, 1)$
request 2	-	$(v_2^2, v_3^2, 1)$	$(v_1^3, v_3^3, 1)$

Table 3.2: Intra-domain requests

3.2.2 Local Intra-Domain Optimization

Meanwhile, the second optimization stage takes the sub-path requests generated from the skeleton paths and attempts to provision (domain-traversing) lightpath segments for them over the individual domain topologies. In particular, intra-domain route selection is done using the same optimization formulation in Section 3.1, i.e., by applying it over the local domain topology, $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$. Once these local sub-path routes have been resolved, wavelength selection is then done in order to ensure wavelength continuity across the all-optical domains. In particular, the *most-used* (MU) assignment strategy is used here as it shown to give good results in both single and multi-domain network settings, see [15, 32]. If, however, an available wavelength cannot be found for a particular sub-path, then the lightpath request for that particular index is dropped (failed). Finally, once all intra-domain wavelengths have been assigned, the complete "end-to-end" multi-domain lightpath route is generated. Namely, this is done for each request index by concatenating the physical links on the inter-domain skeleton route with the expanded physical links for the intra-domain sub-paths (with the same request index). As wavelength selection on inter-domain links is not an issue here, i.e., due to full opto-electronic conversion at gateway OXC nodes, any available wavelength can be selected. The overall psuedocode description of this two-stage ILP solution is also presented in Figure 3.7.

3.2.3 Optional Re-Optimization Pass

Now special cases may arise if requests traversing two (or more) common domains fail optimization setup at *different* domains. In particular, these scenarios can result in multiple requests being denied even if there are available resources to support at least some of them. Hence in order to improve setup success rates here, a second optional "re-optimization" pass is also added, i.e., depicted via the dashed line in

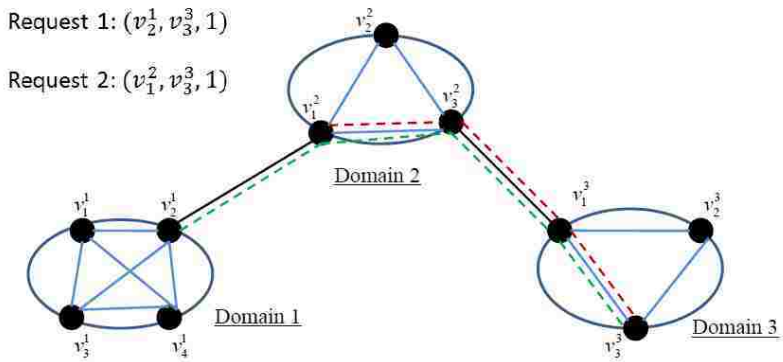

Figure 3.5: Inter-domain skeleton path routing for sample 3-domain network

Figure 3.1. Namely, the goal here is to prune (i.e., deliberately fail) some of the "overlapping" requests in order to allow others to be successfull. This is perhaps best shown via an example in Figure 3.6, in which request 1 (from v_2^1 to v_3^3) overlaps with request 2 (from v_1^2 to v_3^3) at domains 2 and 3. Hence both requests experience blocking, with request 1 failing at at domain 3 and request 2 failing at domain 2, i.e., intra-domain optimizations cannot find wavelength-continuous sub-paths. However, slight resource re-assignment can be done to allow at least one of these requests to be successful. Specifically, if the two failed domain sub-paths in Figure 3.6 are both assigned to request 1, then request 2 can be routed across domain 2, i.e., *without* changing the total number of wavelengths used in the original assignment.

The proposed re-optimization procedure is formally detailed now and tries to re-allocate resources between failed requests traversing one or more common domains (and experiencing blocking at different domains). Carefully note that multiple overlaps can occur between the skeleton routes and this can result in many potential combinations for pruning failed skeleton routes. In fact this pruning problem itself can be treated as an optimization formulation. However in order to simplify matters here, a different strategy is used. Namely, the two-stage optimization is re-run with the capacities of the abstract links in $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ set to the number of *successfully*

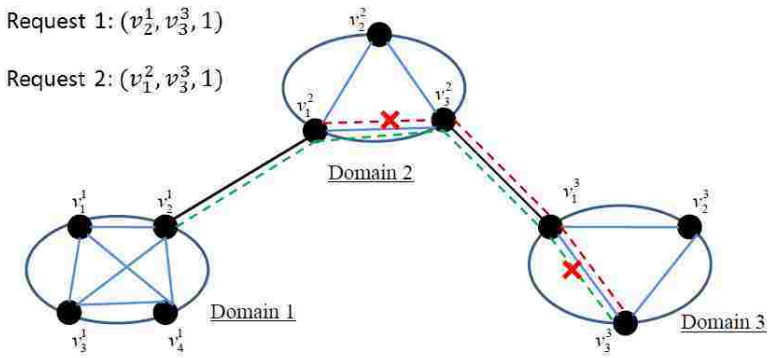

Figure 3.6: Overlapping inter-domain requests failing setup at different domains

routed intra-domain routes between the respective border nodes, i.e., X_{km}^{ij}, as obtained from the first-pass of the optimization (Section 3.2.1). This approach ensures that the number of wavelength resources used at the inter-domain level is the same as that computed in the initial optimization pass.

1: Given domain-level sub-graphs, $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$, intra-domain link sets, $\boldsymbol{E}_{phy}$, and a-priori multi-domain lightpath request set, $\{(s_n, d_n, r_n)\}$.
2: Generate full-mesh topology abstractions for each domain and construct abstract graph $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$.
3: Run first-level ILP optimization to compute skeleton loose routes over $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ to maximize multi-objective function F, Eq. (1).
4: Extract intra-domain traversing segments on successful skeleton paths (above) and generate intra-domain sub-path request sets for each domain i, $\{(v_k^i, v_m^i, 1)\}$.
5: Run second-level ILP optimizations over all domain sub-graphs, $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$, to compute domain-traversing segments. Perform MU wavelength selection for successful sub-paths.
6: Concatenate successful domain-traversing segments to generate complete "end-to-end" lightpath sequences.

Figure 3.7: Pseudocode for two-stage ILP solution (single-pass only)

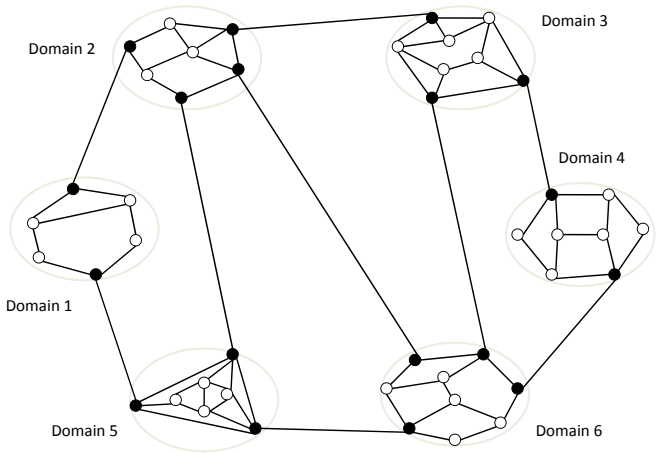

Figure 3.8: 6-domain network topology

3.3 Performance Evaluation

The multi-domain optimization solution is tested and compared against a distributed heuristic scheme in [16]. In particular this heuristic is detailed in Appendix A.1 and uses hierarchical routing and graph-theoretic algorithms to compute and expand skeleton inter-domain lightpath routes (in a load balancing manner). Overall, the evaluations are done using two different network topologies, including a smaller-sized 6-domain network and a larger 16-domain network (reflective of a national backbone). The former topology is shown in Figure 3.8 and has an average domain size of 7 nodes and 9 bi-directional inter-domain links. Meanwhile, the latter topology is a modification of the ubiquitous NSFNET backbone network and is built by replacing all nodes with domains, see Figure 3.9. Here the individual domain sizes are varied between 7-10 nodes and there are a total of 25 bi-directional inter-domain links.

From a programming standpoint, the proposed ILP formulation in Section 3.1 is specified using the PuLP package, which is a Python-based linear programming

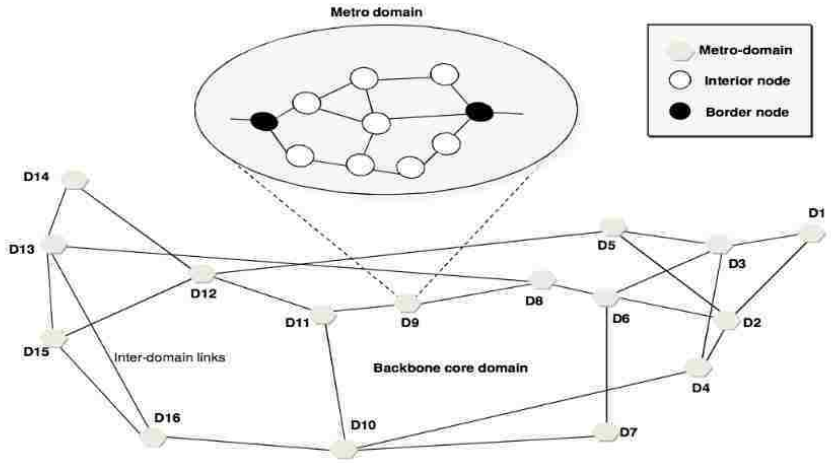

Figure 3.9: Modified 16-domain NSFNET topology

modeler. This formulation is then solved hierarchically (as per Section 3.2) using the *GNU Linear Programming Kit* (GLPK) for a randomly-generated set of multi-domain requests, i.e., uniform random selection of source and destination domains followed by uniform random selection of nodes within these domains. Furthermore, the weights for the objective function in Eq. (3.1) are set to $\{w_1, w_2, w_3\} = \{0.9, 0.05, 0.05\}$. In general, these values emphasize throughput maximization but also try to reduce resource consumption by assigning a non-zero weight to F_2, i.e., to account for the generally higher cost of bandwidth usage on inter-domain links. Load balancing is also done by assigning a non-zero weight to F_3.

Meanwhile the multi-domain heuristic scheme in [16] (Appendix A.1) is analyzed using custom-developed discrete-event simulation models in *OPNET Modeler*^{TM}. Here, the inter-domain routing update thresholds are set to 1% (SCF=0.01) and the corresponding routing *hold-down timers* (HT) are set to 120 seconds. Furthermore, in order to ensure a fair comparison between the optimization and heuristics-based strategies, the same randomly-generated set of lightpath requests are tested for each particular input "load" point, i.e., measured by the number of requests. Furthermore,

all of these requests are assumed to have infinite holding times in order to compare the findings with those from the hierarchical ILP solution, i.e., no departures. Note that the arrival times of these requests are also staggered in multiples of 1,000 seconds in order to allow the inter- and intra-domain link-state routing updates to propagate and associated routing databases to stabilize. Finally, since the order of these input requests can affect the resulting route selection and success rates, these requests are randomly shuffled to generate 10 different variations, and the best results taken for comparison with the ILP model. The detailed findings are now presented.

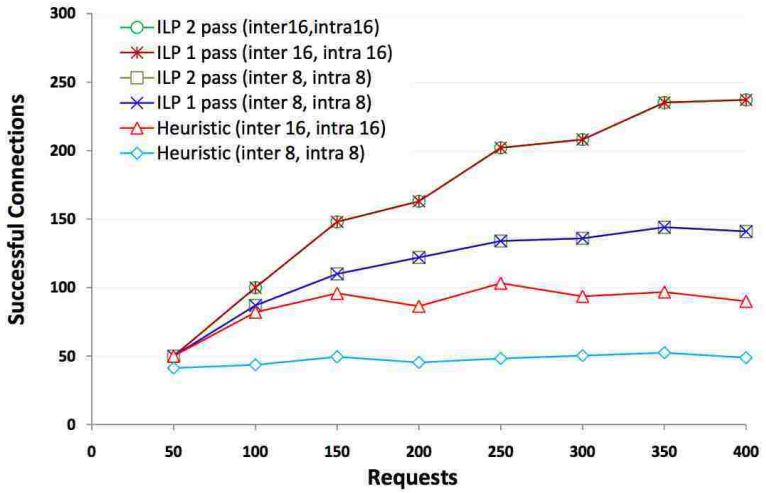

Figure 3.10: Successful requests, 6-domain (C_1,C_2=8 and C_1,C_2=16)

Initial tests are done for the 6-domain network using two different intra/inter-domain link sizes, i.e., C_1,C_2=8, 16 wavelengths, respectively. In particular both the single-pass optimization and double-pass re-optimization schemes are tested here and load-balancing path selection is enabled for the heuristic (shown to give better result, see [16]). The overall setup success rates for varying input request sizes (input loads) are then plotted in Figure 3.10. Overall, these results indicate relatively close performance between the ILP and (hierarchical routing) heuristic scheme at very low load regimes for $C_1, C_2 = 16$, i.e., below 50 requests. However, as input loads

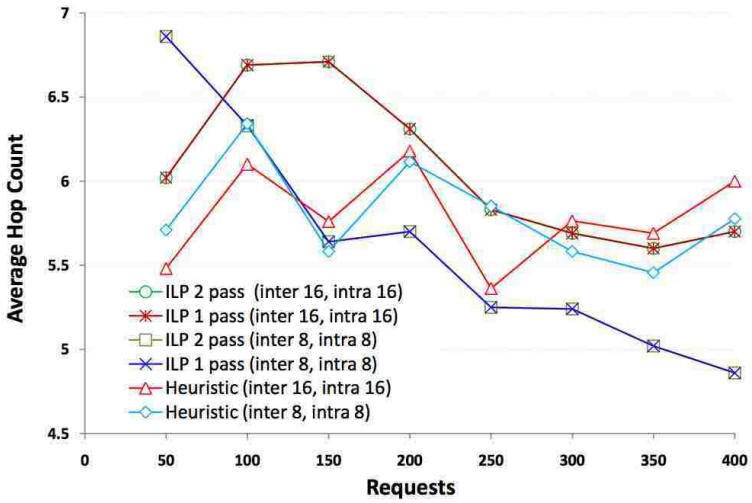

Figure 3.11: Average hop count, 6-domain (C_1,C_2=8 and C_1,C_2=16)

increase to moderate levels, the heuristic scheme quickly saturates around a certain number of successful connections, whereas the ILP schemes continue to improve. For example, the optimization solutions give over twice the number of successful setups at high input loads. In addition, the findings also show very little (no) improvement with the optional ILP re-optimization pass for this smaller topology. Note that additional simulation runs (not shown here) are also done with the multi-domain heuristic scheme [16] using simpler minimal hop selection. Overall, the findings here reveal almost identical (slightly higher) blocking and (slightly lower) hop count values versus load balancing. This is due to the infinite holding time assumption. As a result, minimal hop routing is not tested further.

Next, the corresponding average hop counts are measured and plotted in Figure 3.11 for successful lightpath setups. For the most part, these values indicate declining trends with increasing load, and this is expected as higher levels of contention generally result in increased failure rates for longer paths. The optional re-optimization pass also yields little change in the average hop count values. Also carefully note that these results exhibit higher variability at lower loads due to the reduced batch

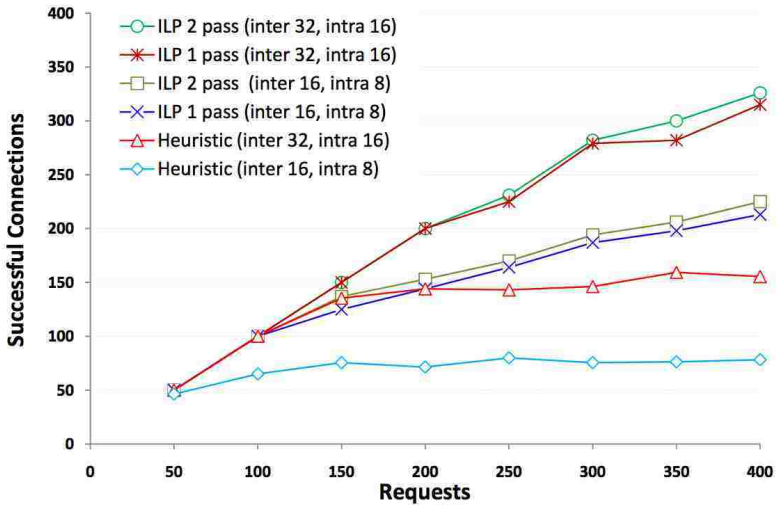

Figure 3.12: Successful requests, 6-domain ($C_2 = 2C_1$=16 and $C_2 = 2C_1$=32)

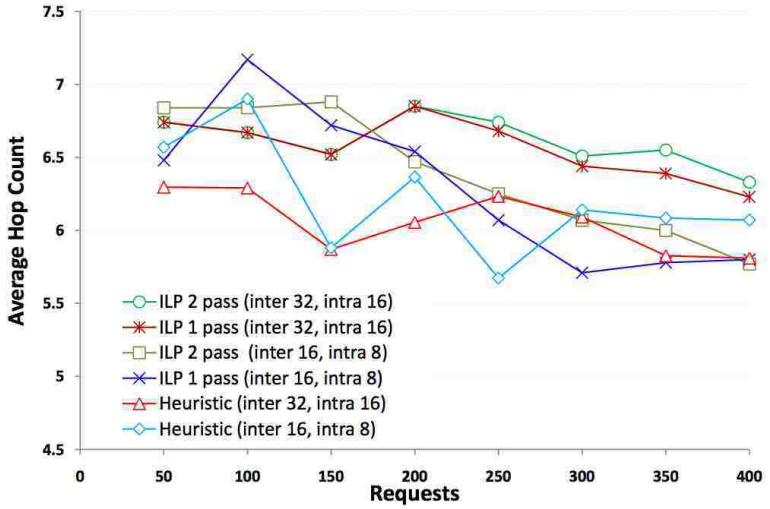

Figure 3.13: Average hop count, 6-domain ($C_2 = 2C_1$=16 and $C_2 = 2C_1$=32)

request (sample) sizes, but this tends to smooth out at heavier loads.

Next, tests are done for the 6-domain scenario with differing intra- and inter-domain link sizes. Namely, the number of inter-domain link wavelengths is set to twice the number of intra-domain wavelengths in order to model more realistic net-

works with increased inter-domain trunk sizes, i.e., $C_2 = 2C_1$. The associated setup success rates are then plotted in Figure 3.12 for varying request sizes and show similar behaviours to the equivalent-link scenario for varying input loads. For example, the ILP-based schemes closely track the hierarchical routing heuristic to about 150 requests (low-medium load regime) but then diverge and give much better performance at higher loads, i.e., over 150% increase in setup success rates (for $C_2 = 2C_1 = 32$ wavelengths). Furthermore, unlike earlier results for equivalent intra-/inter-domain link sizes, here the second re-optimization step does yield a slight improvement in setup success rates, i.e., about 3-5% lower blocking at higher loads. Meanwhile, the average hop count values are also shown in Figure 3.13 and generally indicate declining trends. However, the ILP-based schemes yield slightly higher resource consumption, as they are more successful at finding longer routes at higher load points.

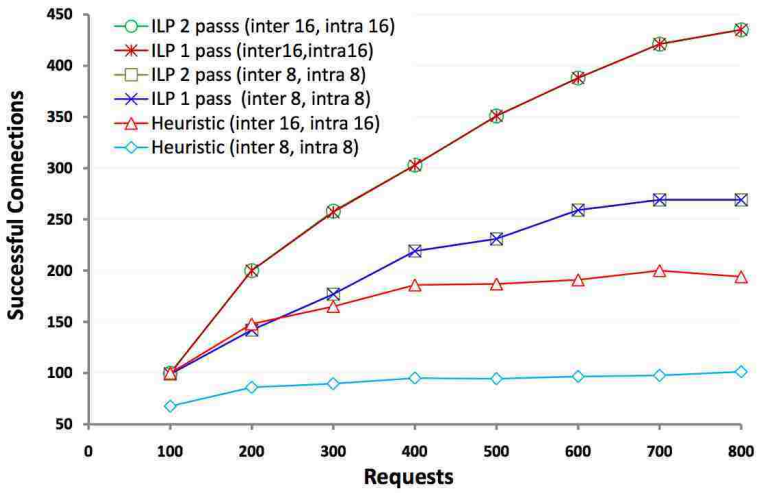

Figure 3.14: Successful requests, 16-domain (C_1,C_2=8 and C_1,C_2=16)

Next, the larger 16-domain modified NSFNET topology is tested for both the optimization and heuristic-based strategies. Namely Figure 3.14 first plots the number of successful setup requests for this network for two different intra- and inter-domain link sizes, i.e., $C_1, C_2 = 8$ and $C_1, C_2 = 16$ wavelengths. Again, these results

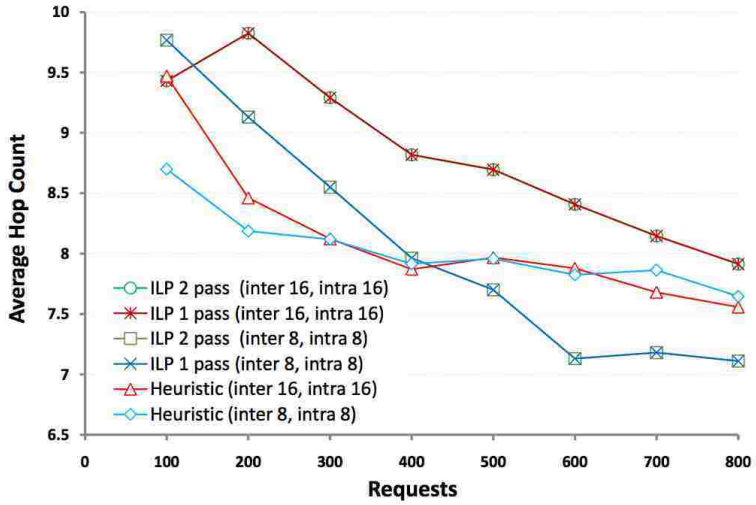

Figure 3.15: Average hop count, 16-domain (C_1,C_2=8 and C_1,C_2=16)

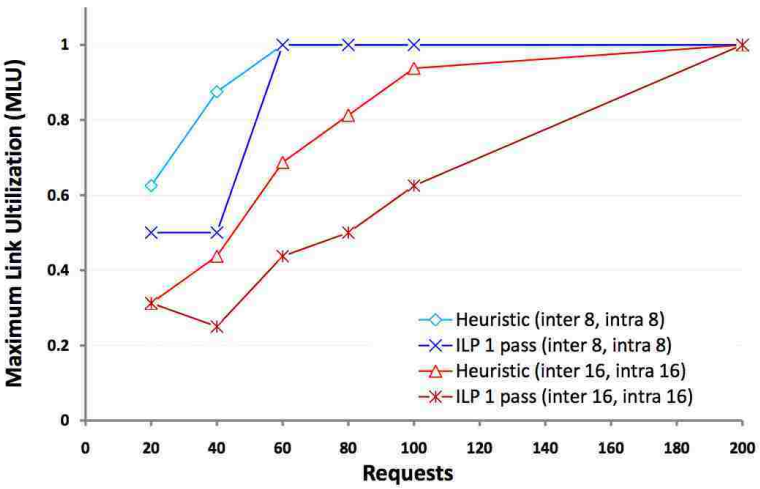

Figure 3.16: MLU values, 16-domain (C_1,C_2=8, C_1,C_2=16)

re-confirm earlier findings with the 6-domain network, and shown that the optimization schemes give well over twice the number of successful lightpath setups at higher loads. Again, the optional "re-optimization" yields very little reduction in blocking. The corresponding average hop counts are also plotted in Figure 3.15 and show a

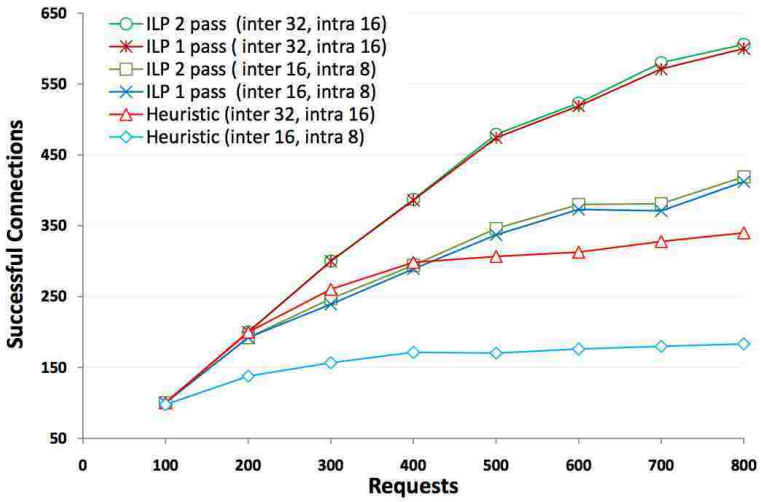

Figure 3.17: Successful requests ($C_2 = 2C_1 = 16$ and $C_2 = 2C_1 = 32$)

more definitive decline with increasing input loads (a-priori batch sizes) for both the ILP and heuristic schemes, i.e., about 30% lower at extremely high loads. Overall, this behavior is expected as increased loads drive up link utilization and lower the probability of finding longer domain-traversing sub-path routes with free (and continuous) wavelengths. Moreover the ILP objective function in Eq. (3.1) also takes into account resource usage, and hence this tries to minimize the number of links used. Now in order to further gauge resource utilization, the maximum link utilization on the physical inter-domain links is also plotted in Figure 3.16 for all the schemes. These results show that the heuristic strategy drives up resource utilization at a much faster rate and therefore achieves link saturation at lower input loads. By contrast, the ILP methods are much more efficient, i.e., less resource intensive, especially for larger link sizes, i.e., $C_1, C_2 = 16$ wavelengths.

Now careful analysis of the optimization results for the 16-domain topology indicates very few setup failures in the second intra-domain ILP stage. In other words, for equivalent intra-/inter-domain link sizes, blocking tends to occur mainly on *inter-*

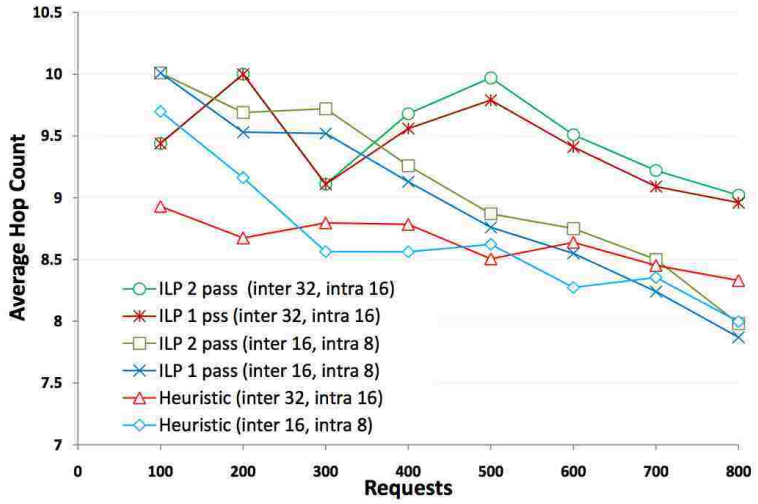

Figure 3.18: Average hop count ($C_2 = 2C_1 = 16$ and $C_2 = 2C_1 = 32$)

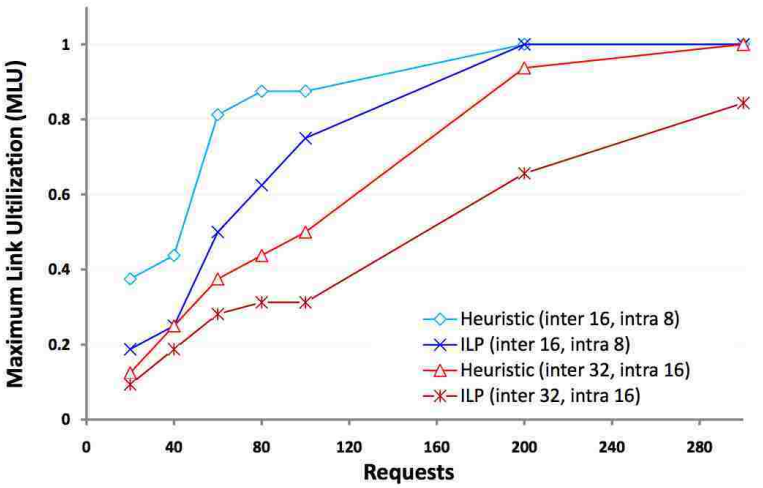

Figure 3.19: MLU values, 16-domain ($C_2 = 2C_1 = 16$, $C_2 = 2C_1 = 32$)

domain links. As a result, further tests are repeated for larger inter-domain link sizes in order to model more realistic link sizing strategies, i.e., $C_2 = 2C_1 = 16$ and $C_2 = 2C_1 = 32$ wavelengths. Namely, Figure 3.17 plots the number of successful setups and Figure 3.18 plots the corresponding average hop count values for this

modified scenario. Again, these findings indicate much lower blocking rates with the ILP-based strategies, averaging between 2-3 times more successful setups at medium-to-high loads. As per the similar runs with the 6-domain network, slight blocking reduction is also observed with the optional re-optimization pass, i.e., in the range of about 3%. The corresponding MLU values for inter-domain links are also shown in Figure 3.19. As per the earlier results in Figure 3.16, the ILP-based schemes give much slower usage growth with increasing load, indicating improved resource allocation versus the heuristic scheme.

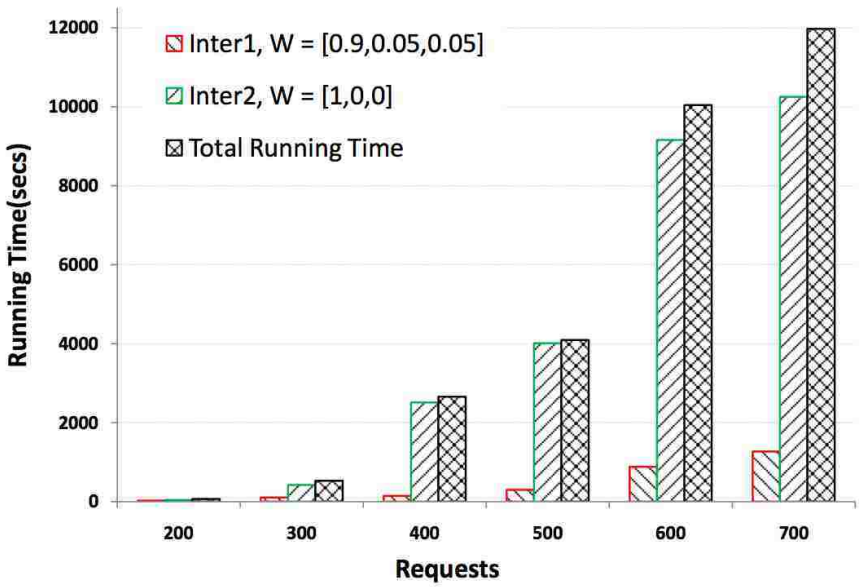

Figure 3.20: Sample ILP run times for 16-domain network (C_1,C_2=8)

Finally the overall run-times for the different optimization steps are also shown in Figure 3.20 for the smaller C_1,$C_2 = 8$ network scenario. These values show super-linear growth and indicate that the individual intra-domain optimizations over the $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$ graphs have the lowest computation times. This is expected since local domain topologies are much smaller than the larger (abstract) inter-domain topology, $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$, for this scenario.

Note that the optimization model is also tested extensively for differing weight values of the objective function, i.e., w_1, w_2, w_3 in Eqs. (3.1) and (3.2). Although these results are not presented here, the overall findings indicate that larger w_1 settings (favoring throughput maximization) give the best results in terms of reduced request blocking rates. Alternatively, pursuing resource minimization objectives, i.e., with larger w_2 values, yields a sharp decline in setup success rates (carried load), e.g., $w_2 = 0.05$ gives lower setup success rates. Hence all the results presented here use $w_1 = 0.90$, and equivalent assignments are also chosen for related weighting choices for the survivability-based optimizations in Chapters 4 and 5.

Chapter 4

Multi-Domain Path Protection Optimization

In general, most multi-domain network recovery schemes implement dedicated pre-provisioned protection to recover from single link failures, i.e., see survey in Section 2.3. Most notably, researchers have proposed a range of per-domain and end-to-end (hierarchical) protection strategies here. However, the former types are not very resource efficient as they route primary/backup path pairs along the same end-to-end domain sequence. Conversely, the latter schemes are much more flexible as they can leverage dynamic inter-domain routing state information to increase the level of domain "separation" along backup routes. As such, hierarchical protection schemes can give much more balanced resource engineering on inter-domain links.

Nevertheless, most hierarchical multi-domain protection schemes also use graph-based heuristic methodologies. As such, it is very difficult for network operators to gauge the true achievable performance of protection in multi-domain settings. In light of this, there is a further need to develop more formal optimization-based models here. Now a few efforts have studied p-cycle [88] protection design for multi-

domain recovery, as noted in Section 2.3.4. For example, the work in [37] proposes an ILP formulation to compute two separate sets of protection cycles, i.e., at the intra- and inter-domain levels, respectively. However, p-cycle techniques are designed to handle underlying fiber link failures and therefore essentially treat all demands as protected, i.e., no selectivity. As such, these solutions may be less applicable in practical settings where carriers will want more selective recovery based upon user needs and active inter-domain (inter-carrier) routing and policy state.

In light of the above, this chapter presents a novel optimization formulation for multi-domain lightpath protection against single link failures. The proposed scheme models realistic hierarchical routing/provisioning environments and extends the non-survivable optimization framework developed in Chapter 3. In particular, the focus here is on dedicated protection only, i.e., since shared protection methodologies may be less applicable in generalized multi-carrier settings owing to revenue and policy constraints [8]. Although shared protection can be considered in more specialized single-carrier multi-domain settings, this is left for future study (as mentioned in Chapter 6). The solution is now presented and its performance analyzed.

4.1 Survivability Optimization Formulation

An optimization-based solution is now presented to compute end-to-end link-disjoint path-pairs for an a-priori set of user lightpath demands. As per Chapter 3, the formulation assumes a hierarchical routing setup in which domains are represented using full-mesh topology abstraction. Furthermore, it is also assumed that domains are internally-transparent (all-optical) but support full wavelength conversion at their border OXC nodes. This overall solution is shown in Figure 4.1 and also uses a two-stage approach to mimic practical multi-domain heuristic designs. Namely, disjoint path-pairs are first computed at the inter-domain level subject to specific TE

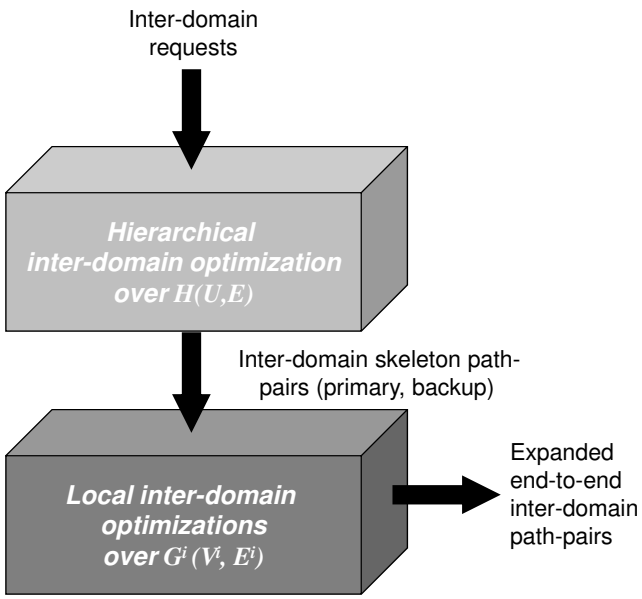

Figure 4.1: Two-step multi-domain ILP protection solution

constraints, followed by sub-path routes at the local intra-domain levels. Consider the details.

4.1.1 Constraints and Objectives

As per the case of non-survivable optimization in Chapter 3, a single generic formulation is proposed here for intra- and inter-domain path-pair route selection. Again, this is possible due to the assumption of full wavelength conversion at border gateway nodes and the decoupling of lightpath route and wavelength selection at the intra-domain level (both of which obviate the need to formulate wavelength assignment in the optimization). Therefore, the optimization problem is only presented for skeleton path-pair route provisioning over the abstract topology, $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$, and this

can readily be extended to the individual domain graphs as well, i.e., $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$.

Again, consider an a-priori set of protected user lightpath demands, denoted by the set of 2-tuples $\{(s_n, d_n)\}$, where n represents the request index, s_n is the source OXC node, d_n is the destination OXC node. As per Section 3.2, it is also assumed that users only request a single wavelength. Furthermore, let the variable f_n represent the number of wavelengths allocated for the n^{th} request, x_{km}^{nij} denote the number of wavelengths routed over link e_{km}^{ij} for the primary path for request n, y_{km}^{nij} denote the number of wavelengths routed over link e_{km}^{ij} for the backup path for request n, and α denotes a fixed MLU value (link utilization). Furthermore, it is also assumed that all abstract links in $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ have capacity C_2.

Given the above, a multi-objective ILP function F is now defined to compute link-disjoint primary/backup path-pairs over the abstract topology contingent to several TE objectives as follows:

$$\text{Max } F = w_1 \sum_{n \in \boldsymbol{N}} f_n - w_2 \sum_{n \in \boldsymbol{N}} \sum_{e_{km}^{ij} \in \boldsymbol{E}} (x_{km}^{nij} + y_{km}^{nij}) - w_3 \alpha = w_1 F_1 + w_2 F_2 + w_3 F_3 \quad (4.1)$$

where

$$F_1 = \sum_{n \in \boldsymbol{N}} f_n \quad (4.2a)$$

$$F_2 = - \sum_{n \in \boldsymbol{N}} \sum_{e_{km}^{ij} \in \boldsymbol{E}} (x_{km}^{nij} + y_{km}^{nij}) \quad (4.2b)$$

$$F_3 = -\alpha \quad (4.2c)$$

and w_1, w_2, and w_3 are fractional weights summing to unity. Carefully note that this function is very similar to that defined for non-survivable provisioning in Eq. (3.1), i.e., with additional terms for backup wavelength usage. Namely, it incorporates throughput maximization, F_1, resource minimization, F_2, as well as load-balancing, F_3. Now in order to generate valid solutions, the following constraints are defined

here:

$$\sum_{(j,m):e_{km}^{ij}\in \boldsymbol{E}} x_{km}^{nij} - \sum_{(j,m):e_{mk}^{ji}\in \boldsymbol{E}} x_{mk}^{nji} = \begin{cases} f_n; & if\ v_k^i = s_n \\ -f_n; & if\ v_k^i = d_n \\ 0; & otherwise \end{cases} ; n \in \boldsymbol{N} \tag{4.3}$$

$$\sum_{(j,m):e_{km}^{ij}\in \boldsymbol{E}} y_{km}^{nij} - \sum_{(j,m):e_{mk}^{ji}\in \boldsymbol{E}} y_{mk}^{nji} = \begin{cases} f_n; & if\ v_k^i = s_n \\ -f_n; & if\ v_k^i = d_n \\ 0; & otherwise \end{cases} ; n \in \boldsymbol{N} \tag{4.4}$$

$$x_{km}^{nij} + y_{km}^{nij} \leq f_n; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{4.5}$$

$$\sum_{n\in \boldsymbol{N}} (x_{km}^{nij} + y_{km}^{nij}) \leq \alpha C_2; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{4.6}$$

$$x_{km}^{nij} \in \{0,1\}; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{4.7}$$

$$y_{km}^{nij} \in \{0,1\}; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{4.8}$$

$$f_n \in \{0,1\}; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{4.9}$$

$$0 \leq \alpha \leq 1 \tag{4.10}$$

In particular, Eqs. (4.3) and (4.4) represent flow conservation constraints for the primary and backup paths between the incoming and outgoing flows at each (border) node. Meanwhile, Eq. (4.5) ensures that the primary and backup paths for the same request do not traverse the same links in $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$, i.e., link-disjointness constraint. Also, Eq. (4.6) restricts the total relative traffic load carried on an inter-domain link to under the pre-defined MLU value, i.e., below αC_2. Finally, Eqs. (4.7)-(4.9) represent binary constraints, and Eq. (4.9) bounds the MLU value to a positive fraction.

Overall, the above formulation has a total of $(N+2N|\boldsymbol{E}|)$ variables and $O(2N|\boldsymbol{U}|+N|\boldsymbol{E}|)$ equations, see Table 4.1 . Namely, a flow variable, f_n, is defined for each request, and further flow-specific variables are also defined for each link, i.e., x_{km}^{nij} and y_{km}^{nij} variables for the primary and backup paths, respectively. Now since this formula-

Equation	$\boldsymbol{H}(\boldsymbol{U},\boldsymbol{E})$	$\boldsymbol{G}^i(\boldsymbol{V}^i,\boldsymbol{L}^i)$
Eq.(4.3)	$N\lvert\boldsymbol{U}\rvert$	$N_i(\lvert\boldsymbol{V}^i\rvert)$
Eq.(4.4)	$N\lvert\boldsymbol{U}\rvert$	$N_i(\lvert\boldsymbol{V}^i\rvert)$
Eq.(4.5)	$N\lvert\boldsymbol{E}\rvert$	$N_i\lvert\boldsymbol{L}^i\rvert$
Eq.(4.6)	$\lvert\boldsymbol{E}\rvert$	$\lvert\boldsymbol{L}^i\rvert$
Eq.(4.7)	$N\lvert\boldsymbol{E}\rvert$	$N_i\lvert\boldsymbol{L}^i\rvert$
Eq.(4.8)	$N\lvert\boldsymbol{E}\rvert$	$N_i\lvert\boldsymbol{L}^i\rvert$
Eq.(4.9)	N	N_i
Eq.(4.10)	1	1
total	$O(2N(\lvert\boldsymbol{U}\rvert+\lvert\boldsymbol{E}\rvert)$	$O(2N_i(\lvert\boldsymbol{V}^i\rvert+\lvert\boldsymbol{L}^i\rvert))$

Table 4.1: Number of equations

tion has to be solved for the inter-domain topology as well as all intra-domain topologies, the maximum computational complexity will be dominated by the topology with the maximum total equation count, i.e., $\max_i\{2N(\lvert\boldsymbol{U}\rvert+\lvert\boldsymbol{E}\rvert), 2N_i(\lvert\boldsymbol{V}^i\rvert+\lvert\boldsymbol{L}^i\rvert)\}$, where N_i represents the number of local requests for domain i.

4.1.2 Approximation Solution

Expectedly, the above ILP formulation will pose very high computational complexity as the link-disjoint path-pair optimization problem is NP-hard [9]. Hence in order to address this concern, a further efficient approximation solution is proposed here based upon a *linear programming* (LP) approach. In particular, LP formulations can be used to approximate ILP solutions by relaxing binary constraints. Hence for the proposed ILP formulation in Section 4.1.1, the constraints in Eqs. (4.7)-(4.9) can be mapped to equivalent linear constraints as follows:

```
1: Solve the LP relaxation of ILP formulation, Eqs. (4.1)-(4.10)
2: for each request i do
3:     Identify all primary and backup paths, and check for link-disjoint paths
4:     if link-disjoint path pairs exists then
5:         Round f_i using randomized rounding
6:         if f_i is rounded to 1 then
7:             Test if the path-pair associated with largest allocated bandwidth for
               request i has a free wavelength. If not, test other path-pairs.
8:         end if
9:     end if
10: end for
```

Figure 4.2: Pseudocode for LP solution processing (rounding algorithm)

$$x_{km}^{nij} \in \{0,1\}; \;\; \Rightarrow \;\; 0 \leq x_{km}^{nij} \leq 1; \tag{4.11}$$

$$y_{km}^{nij} \in \{0,1\}; \;\; \Rightarrow \;\; 0 \leq y_{km}^{nij} \leq 1; \tag{4.12}$$

$$f_n \in \{0,1\}; \;\; \Rightarrow \;\; 0 \leq f_n \leq 1; \tag{4.13}$$

yielding real fractional values for the variables x_{km}^{nij}, y_{km}^{nij}, and f_n in the range of $[0, 1]$. This relaxed version can then be solved in polynomial time using any standard LP solver. However, the non-integral nature of the resulting LP solution will likely yield request splitting over multiple paths (with the respective allocated capacities summing to f_n). In addition, it may also yield several primary/backup paths, which may not necessarily be mutually link-disjoint. As a result, a final rounding algorithm is also proposed here, as shown in Figure 4.2. Overall, some sample runs with a basic 16-node NSFNET topology and moderate batch request sizes show very close blocking performance between the full ILP and its LP approximation solution, i.e., within 7-8% success rates.

4.2 Solution Approach

As shown in Figure 4.1, the proposed scheme also performs optimization at the inter- and intra-domain levels. Namely, the first stage performs optimization over the abstract graph $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ to obtain a set of skeleton primary/backup path-pairs. The second stage then optimizes domain-traversing local routes over the individual domain sub-graphs, $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$. These local segments are then inserted into their corresponding skeleton paths to generate the end-to-end link-disjoint primary/backup lightpaths, akin to the case of non-survivable lightpath routing in Chapter 3. Again, this "two-stage" approach mimics distributed multi-domain protection heuristics which compute and expand path-pair sequences. Consider some further details.

4.2.1 Hierarchical Inter-Domain Path-Pair Optimization

Carefully note that the optimization model presented in Section 4.1.1 does not differentiate between physical and abstract links, i.e., when applied over the graph $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$. As a result, this can result in overly-restrictive routing by preventing primary/backup routes from traversing the same *abstract* link, i.e., domain. Hence in order to remove this limitation, the ILP constraints are slightly modified to only enforce primary/backup link-disjointness for *physical* inter-domain links, i.e., Eq. (4.5). This modification will allow both the primary and backup routes (for a given request) to traverse common domains (abstract links) and will help lower blocking performance, i.e., equivalent to *link disjoint* (LD) protection strategy [63] (Section 2.3.1). Now once the ILP is formulated over $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$, it is solved using the proposed approximation algorithm in Figure 4.2 to generate an integer-based solution. Subsequently, the abstract links used by all successfully-computed skeleton path-pairs are identified to help generate the individual "domain-level" lightpath (sub-path) requests. Namely, these abstract links correspond to the local domain-traversing

lightpath segments that must be expanded (in the second optimization stage, Section 4.2.2) in order to generate complete explicit end-to-end path sequences. For tracking purposes, these extracted intra-domain requests are also assigned the same index as the original inter-domain request (from which they are extracted).

4.2.2 Local Intra-Domain Optimization

The second optimization stage takes the sub-path requests (abstract links) from the above skeleton path-pair optimization step (Section 4.2.1) and "expands" them over the individual domain topologies, $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$. Again, the relaxation solution in Figure 4.2 is re-used here to generate acceptable solutions in reasonable time. Now carefully note that link-disjointness requirements will only arise for those request indices whose primary and backup skeleton paths traverse the same domains. In all other cases, the link-disjointness constraints in Eq. (4.5) can be omitted when optimizing over the individual graphs, $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$, i.e., resulting in reduced computational complexity.

Now once the local sub-path routes have been computed, wavelength selection is done in order to ensure channel continuity across "all-optical" domains. Again, the MU wavelength assignment strategy is chosen here, as per the case of non-survivable optimization (Section 3.2.2). However, if a free wavelength cannot be found for a particular request along all its traversed (primary, backup) domains, then this request is dropped/failed. Otherwise if all intra-domain wavelengths are assigned, the explicit "end-to-end" multi-domain primary/backup lightpaths routes are generated by concatenating all the traversing domain-level sub-paths into their skeleton route paths (with same request index). As there is no requirement for wavelength selection on inter-domain links, any free wavelength can be assigned.

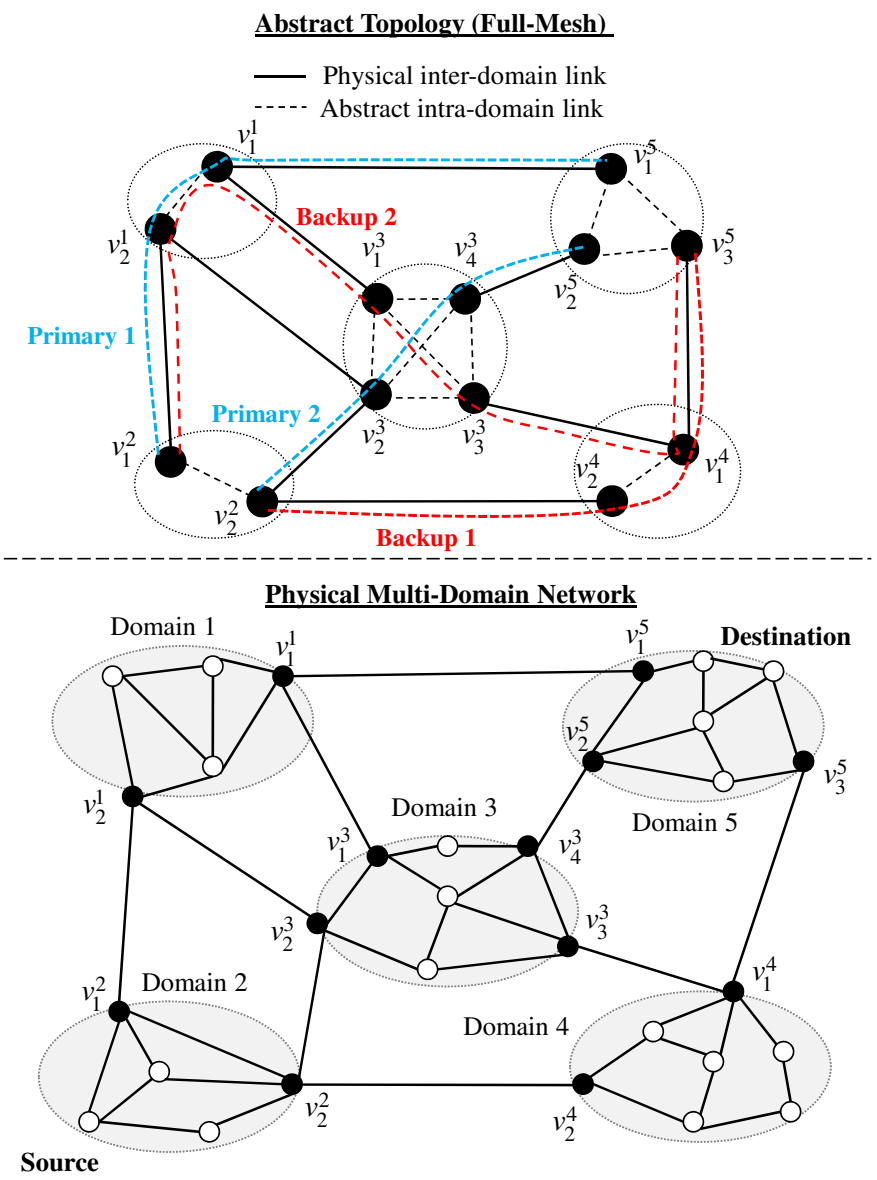

Figure 4.3: Skeleton loose route (LR) path-pair computation

An example of inter-domain skeleton path-pair computation and expansion is also shown in Figure 4.3 for a sample lightpath request between internal nodes in domains 2 and 5. In particular two different cases are shown here, with the primary and backup skeleton routes shown in blue and red, respectively. In the first case the primary/backup skeleton paths do not overlap except at the source and destination domains, i.e., primary 1 $\{v_1^2, v_2^1, v_1^1, v_1^5\}$, backup 1 $\{v_2^2, v_2^4, v_1^4, v_3^5\}$. Meanhwile in the second case the primary/backup skeleton routes intersect at domain 3, i.e., primary 2 $\{v_2^2, v_2^3, v_4^3, v_2^5\}$, backup 2 $\{v_1^2, v_2^1, v_1^1, v_1^3, v_3^3, v_1^4, v_3^5\}$. Hence intra-domain optimization at domain 3 will require link-disjointness constraints when expanding abstract links $v_2^3 - v_4^3$ (primary 2) and $v_1^3 - v_3^3$ (backup 2).

4.3 Performance Analysis

The proposed multi-domain protection optimization scheme is now tested and compared with the advanced protection heuristic scheme in [63]. In particular, this latter solution uses hierarchical routing and implements both LD and *domain disjoint* (DD) path pair computation (see Appendix A.2 for more details). Furthermore, tests are done using the same 6-domain and 16-domain network topologies used in Chapter 3, i.e., Figure 3.8 and 3.9, respectively. Again, inter-domain lightpath requests are generated in a random manner with varying a-priori batch sizes. Meanwhile, the objective function weights in Eq. (4.1) are set to w_1=0.99, w_2=0.005, and w_3=0.005 in order to emphasize throughput maximization.

Now the hierarchical ILP model in Section 4.1 is also coded using the PuLP linear programming modeler and then solved using the GLPK toolkit (as per Chapter 3). Meanwhile the multi-domain protection heuristic in [63] is simulated using *OPNET Modeler*$^{\text{TM}}$for the same input batch request sizes (with similar routing update thresholds and hold-down timer values as those used for the non-survivable simulations in

Chapter 3). In particular, load-balancing path selection is done here and only LD protection is tested, i.e., as DD protection yields much higher blocking, see [63] and Appendix A.2.

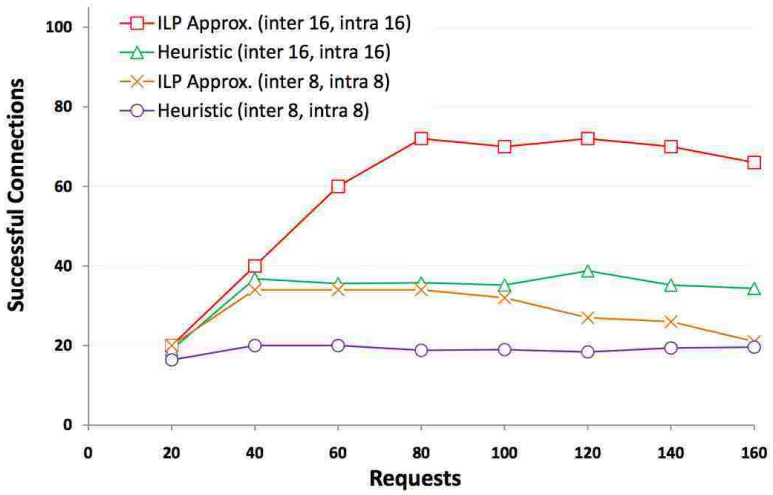

Figure 4.4: Successful requests, 6-domain ($C_1, C_2 = 8$, $C_1, C_2 = 16$)

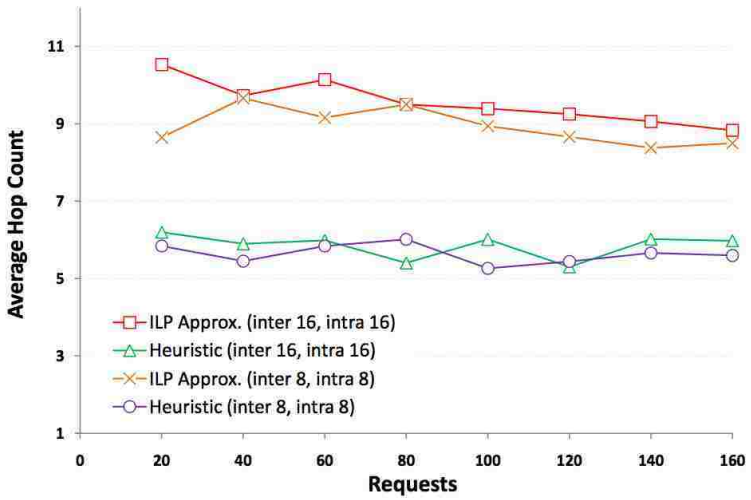

Figure 4.5: Average hop count, 6-domain ($C_1, C_2 = 8$, $C_1, C_2 = 16$)

Initial tests are done for the 6-domain network (Figure 3.8) using equivalent

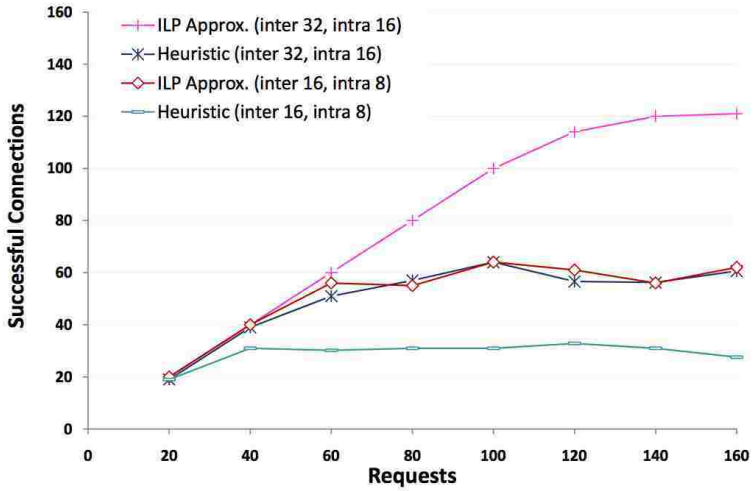

Figure 4.6: Successful requests, 6-domain ($C_2 = 2C_1 = 16$, $C_2 = 2C_1 = 32$)

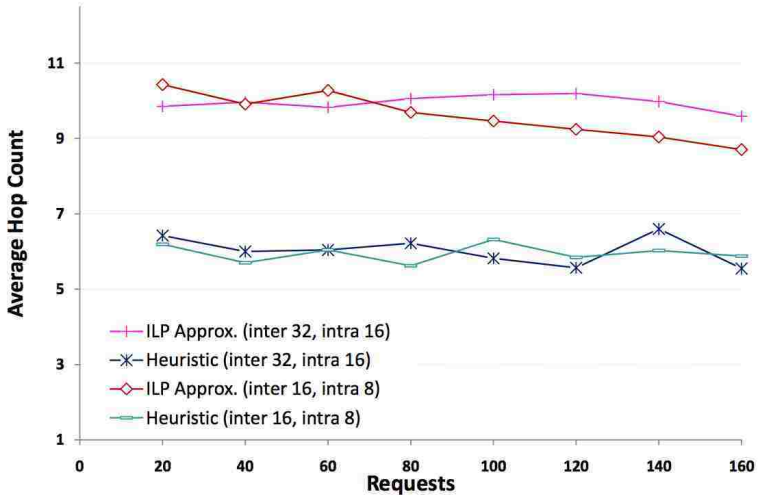

Figure 4.7: Average hop count, 6-domain ($C_2 = 2C_1 = 16$, $C_2 = 2C_1 = 32$)

intra-/inter-domain link sizes, i.e., $C_1, C_2 = 8$ and $C_1, C_2 = 16$ wavelengths. The resultant setup success rates are then plotted in Figure 4.4 with a request being counted as successful if and only if both its primary and backup lightpath routes are established. Overall these results show very close performance between the ILP and heuristic solutions at very low load regimes for both link sizes, i.e., below 40 requests.

However, with increasing request batch sizes, the performance gap between these two strategies quickly increases, with the ILP scheme giving almost twice the number of successful setups as the heuristic scheme. Carefully note that very large batch sizes (over 100 requests) tend to drive down the ILP success rates for the smaller link size scenario, i.e., $C_1, C_2 = 8$ wavelengths. This is due to the relaxation method described in Section 4.1.2. Namely, since link-disjointness is not *strictly* enforced here, the chance of link overlaps between primary and backup routes increases with larger batch requests, making the ILP scheme less effective. The average hop count values are also measured and plotted in Figure 4.5. Generally, the ILP scheme gives larger values here as compared to the heuristic strategy. For the most part, these values also indicate slight declines with increasing load, as longer paths have a higher chance of setup failure due to increased resource contention.

Additional tests are also done for the 6-domain network with more realistic link settings, i.e., with differing inter-/intra- domain links. Namely, inter-domain link sizes are set to twice those of the intra-domain links, i.e., $C_2 = 2C_1 = 16$ and $C_2 = 2C_1 = 32$ wavelengths. The associated success setup rates are then plotted in Figure 4.6 for varying request sizes and confirm the findings with the equivalent-link scenario. For example, for the case of $C_2 = 2C_1 = 32$ wavelengths, the hierarchical routing heuristic closely tracks the optimization-based scheme until about 60 requests. Subsequently, the heuristic solution enters saturation whereas the ILP scheme continues to improve, i.e., over twice the setup success rate. Commensurate average hop count values are also shown in Figure 4.7, and akin to Figure 4.5, show larger values for the ILP-based solution.

Next tests are done using the larger 16-domain backbone topology. In particular, the setup success rates for equivalent intra-/inter-domain link sizes are first plotted in Figure 4.8. Two different link sizes are also evaluated here, i.e., $C_1, C_2 = 8$ and $C_1, C_2 = 16$ wavelengths, and the findings show notably-higher success rates (lower

blocking) with the optimization approach for larger 16-wavelength links, i.e., almost 80% more setups at medium-to-high loads. In fact the heuristic scheme is only comparable at very low connection loads and its performance levels off much faster, i.e., saturation effect. However the relative improvement when using optimization for smaller 8 wavelength link sizes is somewhat lower, but still very significant, i.e., anywhere from 20-75% more accepted requests at medium-to-high loads.

Next, the average path lengths are also measured for the various schemes and plotted in Figure 4.9. Overall, these results show declining values with increasing load for the optimization approach. This is expected since shorter primary/backup routes will result in lower resource consumption, and thereby drive up throughput, i.e., Eq. (4.1). By contrast the heuristic scheme shows more of a flat trend for the average hop count values. This is due to the fact that this strategy provisions paths in a sequential manner, i.e., setup success is more dependent upon the order of the requests. Also, to further gauge resource utilization, the maximum link utilization on physical inter-domain links is also measured and plotted in Figure 4.10. These results show that the ILP scheme gives much slower utilization growth, indicating more balanced traffic distribution. This is particularly evident for larger link sizes, i.e., $C_1, C_2 = 16$ wavelengths. By contrast, the heuristic strategy consumes link resources in a much more uneven manner, leading to faster saturation.

Finally, tests are done for varying intra-/inter-domain link sizes. In particular, the wavelength capacities of inter-domain links are doubled versus the intra-domain links in order to model more realistic settings, i.e., $C_2 = 2C_1 = 16$ and $C_2 = 2C_1 = 32$ wavelengths (akin to Section 3.3). The number of successful lightpath setups are then plotted in Figure 4.11, and these findings show much better gains with the proposed optimization solution i.e., close to twice the setup success rates at medium-to-high loads. Moreover, the separation between the heuristic and optimization-based strategies for the smaller link sizes ($C_2 = 2C_1 = 16$ wavelengths) is also much more

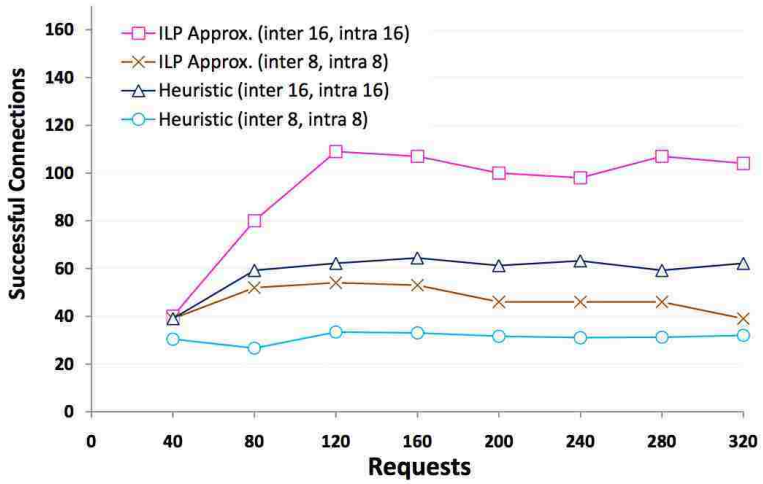

Figure 4.8: Successful requests, 16-domain ($C_1, C_2 = 8$, $C_1, C_2 = 16$)

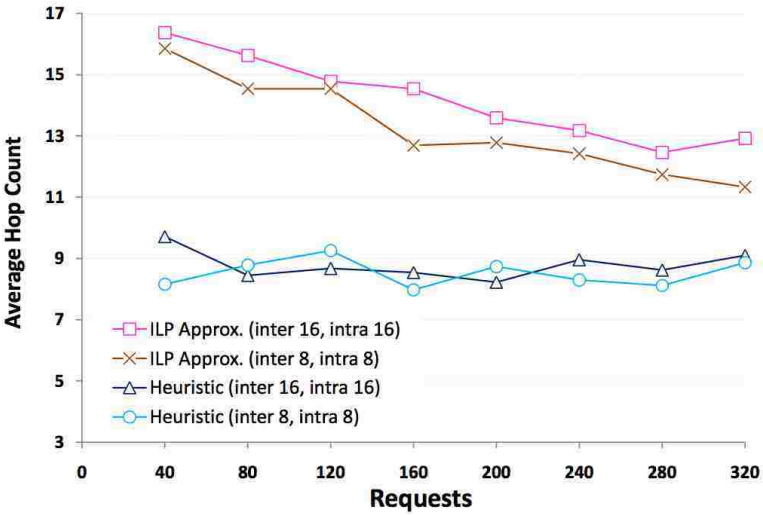

Figure 4.9: Average hop count, 16-domain ($C_1, C_2 = 8$, $C_1, C_2 = 16$)

noticeable here, i.e., close to 100% more successful setups as compared to Figure 4.8. The average path lengths for these scenarios are also shown in Figure 4.12 and follow the same trends as those in Figure 4.9 for equivalent intra- and inter-domain link sizes, i.e., slight hop count declines with the optimization scheme at higher loads.

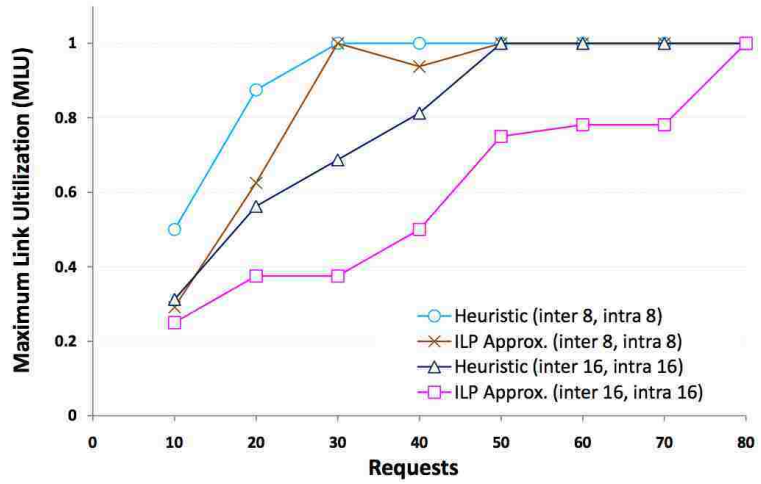

Figure 4.10: MLU values, 16-domain ($C_1, C_2 = 8$, $C_1, C_2 = 16$)

Finally, the corresponding maximum link utilization results are plotted in Figure 4.13. Similar to the findings with equivalent intra-/inter-domain link sizes in Figure 4.10, here the ILP-based scheme gives much slower resource saturation, i.e., more balanced traffic distribution as compared to the heuristic scheme.

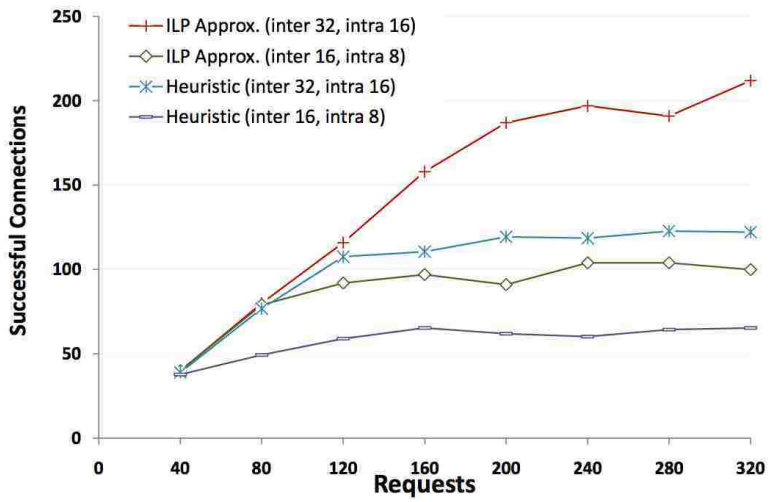

Figure 4.11: Successful requests, 16-domain ($C_2 = 2C_1 = 16$, $C_2 = 2C_1 = 32$)

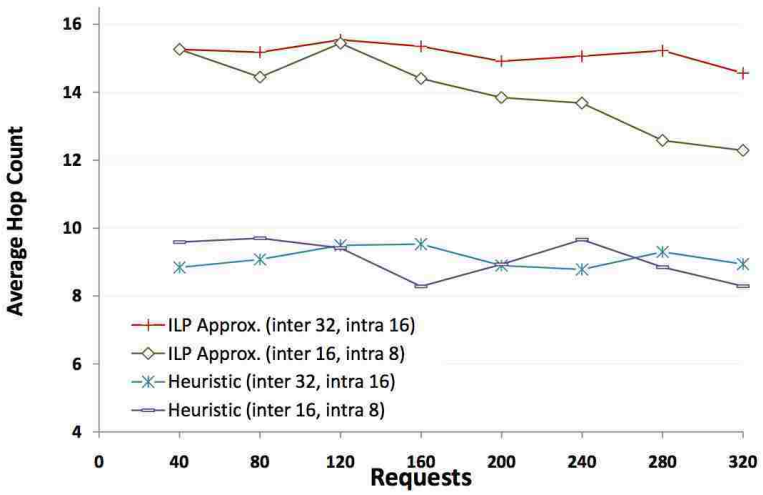

Figure 4.12: Average hop count, 16-domain ($C_2 = 2C_1 = 16$, $C_2 = 2C_1 = 32$)

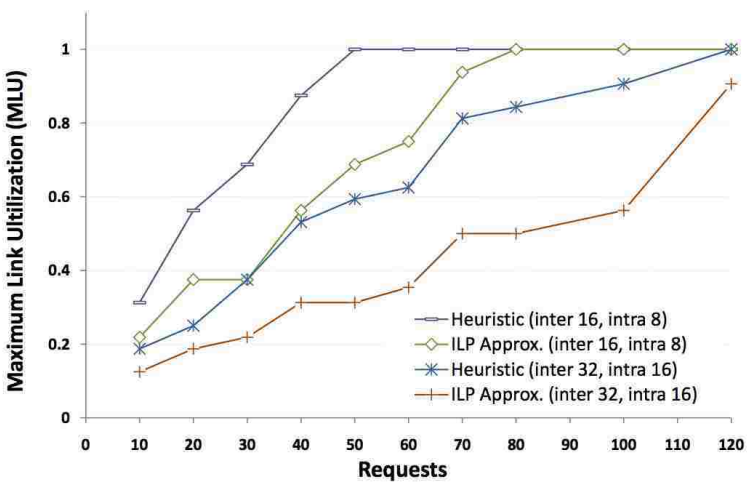

Figure 4.13: MLU values, 16-domain ($C_2 = 2C_1 = 16$, $C_2 = 2C_1 = 32$)

Chapter 5

Multi-Domain Path Protection Optimization Against Multiple Failures

Multi-failure network recovery is a relatively new topic that is gaining much attention today. In particular, large-scale national *information technology* (IT) infrastructures are increasingly vulnerable to catastrophic events, and multi-domain networks are particularly susceptible given their large geographic spread. Now in general, it is very difficult to pre-provision backup resources to handle multiple link/node failure combinations. Hence researchers have started to study advanced *probabilistic* protection schemes to minimize the risk exposure of primary/backup path-pairs in single-domain settings [68, 69, 89]. Extending upon this, new work has also proposed probabilistic protection recovery for distributed multi-domain networks experiencing multiple correlated (disaster type) failures [29] (Section 2.3.1). Namely, this approach introduces risk-based state information into the hierarchical routing process, i.e., topology abstraction, and then leverages it to compute primary/backup skeleton path-pairs to jointly minimize failure probabilities as well as lower TE costs.

Nevertheless, the multi-domain probabilistic protection scheme in [29] is still a heuristic-based solution. As a result, there is a further need to develop advanced optimization-based schemes for probabilistic multi-failure recovery in multi-domain optical networks. However, there are no known studies in this particular area. Along these lines, this chapter presents a novel solution which leverages the p-SRLG multi-failure stressor model in [68] and extends the survivability formulation presented in Chapter 4. In particular, critical link failure probability information is introduced into the topology abstraction model and then applied for multi-domain path protection optimization. Consider the full details now.

5.1 Optimization Formulation

The optimization-based solution is now presented. As per Chapters 3 and 4, the proposed formulation assumes a hierarchical routing setup with full-mesh topology abstraction. Furthermore, it is also assumed that domains are internally-transparent (all-optical) but support full wavelength conversion at their border OXC nodes. This overall solution also uses a two-stage approach (as shown in Figure 4.1) to mimic practical multi-domain designs. Namely, disjoint path-pairs are first computed at the inter-domain level subject to specific risk and TE objectives, followed by sub-path route expansion at the local intra-domain level.

5.1.1 Notation Overview

Before presenting the optimization model, the correlated multi-failure p-SRLG model in [68] is introduced along with its requisite notation. Namely, consider a multi-domain network defined using the same notation as in Section 3.1.1. Here a pre-defined set of correlated p-SRLG events (stressors) is defined for this network via

the set $\boldsymbol{R}$, where each event $r \in \boldsymbol{R}$ has an occurrence probability of $\pi_r \in [0,1]$, and $\sum_{r \in \boldsymbol{R}} \pi_r = 1$. Furthermore, it is also assumed that these events are sufficiently rare such that they are mutually exclusive, i.e., $P(r_1 \cap r_2) = 0$ for all $r_1 \neq r_2$. Next, appropriate failure regions are defined for each stressor, $r \in \boldsymbol{R}$, to model its localized impact regime. Leveraging this, conditional failure probabilities are then defined for each link e^{ij}_{km} (or l^{ij}_{km}) in the region of event r, i.e., $p^r_{i_k j_m}$. Namely, a link has a none-zero $p^r_{i_k j_m}$ value if it falls within the failure region of stressor event $r \in \boldsymbol{R}$. Furthermore, as per [68], it is also assumed that all link failures within a region (for a given stressor r) are independent.

5.1.2 Constraints and Objectives

An optimization formulation is now presented to jointly achieve TE and risk minimization objectives. Again this model is only detailed for inter-domain (skeleton) path-pair routing over the abstract topology, $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$, i.e., since this formulation can readily be extended to the individual domain graphs as well, $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$. Now consider an a-priori set of protected user lightpath requests, denoted by the set of 2-tuples $\{(s_n, d_n)\}$, where n represents the request index, s_n is the source OXC node, and d_n is the destination OXC node. As per Sections 3.2 and 4.1.1, it is also assumed that all requests are for a single wavelength. Furthermore, f_n is used to represent the number of allocated wavelengths for request n, x^{nij}_{km} is used to denote the number of wavelengths routed over link e^{ij}_{km} for the primary path for request n, and y^{nij}_{km} is used to denote the number of wavelengths routed over link e^{ij}_{km} for the backup path for request n. Now given single wavelength requests, necessarily x^{nij}_{km} and y^{nij}_{km} are binary variables in $\{0,1\}$. Furthermore, without loss of generality, it is also assumed that all *abstract* links in $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ have capacity C_2, and their conditional failure probabilities are given by the failure probability of the minimal risk path between the respective border nodes, i.e., by $\sum_{n \in \boldsymbol{N}_i} \sum_{r \in \boldsymbol{R}} \pi_r (1 - \prod_{l^{ii}_{km} \in \boldsymbol{L}^i} (1 - p^r_{i_k i m} x^{nii}_{km}))$.

Using the above notation and topology abstraction framework, the optimization formulation defines a multi-objective function to compute link-disjoint primary/backup path-pairs subject to several constraints. Specifically, the objectives here include throughput maximization, resource minimization, *and* risk minimization. Now the terms for first two objectives have already been defined in Chapter 4, i.e., see Eqs. (4.2a) and (4.2b), respectively. Hence the focus here is more on defining the latter term for risk minimization, as detailed next.

First, consider the conditional *joint* failure probability for a *single* path-pair contingent to event $r \in \boldsymbol{R}$. Based upon the notation introduced in Section 5.1.1, the probability that a link $e^{ij}_{km} \in \boldsymbol{E}$ will *not* be affected by event r is given by $1 - p^r_{i_k j_m}$. Now let $\boldsymbol{x} = \{x^{nij}_{km}\}$ denote the primary path links for request n, and let $\boldsymbol{y} = \{y^{nij}_{km}\}$ denote the backup path links for request n. Hence the failure probability for the primary path in event of r, $F^n_r(p^r, x)$, is given by:

$$F^n_r(p^r, x) = 1 - \prod_{e^{ij}_{km} \in \boldsymbol{E}} (1 - p^r_{i_k j_m} x^{nij}_{km}) \tag{5.1}$$

Similarly, a failure probability can also be computed for the backup path, i.e., $F^n_r(p^r, y)$. Now since the two paths $\boldsymbol{x}$ and $\boldsymbol{y}$ are necessarily link-disjoint, their failure probabilities are also independent due to the assumption that all link failures are independent contingent to event r. Hence the overall conditional failure probability (risk) of the link-disjoint path-pair for request n is given by:

$$F^n_r(p^r, x) F^n_r(p^r, y) = (1 - \prod_{e^{ij}_{km} \in \boldsymbol{E}} (1 - p^r_{i_k j_m} x^{nij}_{km}))(1 - \prod_{e^{ij}_{km} \in \boldsymbol{E}} (1 - p^r_{i_k j_m} y^{nij}_{km})) \tag{5.2}$$

Using the above, a multi-objective *integer nonlinear programming* (INLP) objective function F, is defined to incorporate risk minimization concerns into the link-disjoint

path-pair routing process as follows:

$$\begin{aligned}\text{Min } F &= w_1 \sum_{n \in \boldsymbol{N}} (1 - f_n) + w_2 \sum_{n \in \boldsymbol{N}} \sum_{e^{ij}_{km} \in \boldsymbol{E}} (x^{nij}_{km} + y^{nij}_{km}) + w_3 \sum_{n \in \boldsymbol{N}} \sum_{r \in \boldsymbol{R}} \pi_r F^n_r(p^r, x) F^n_r(p^r, y) \\ &= w_1 F_1 + w_2 F_2 + w_3 F_3 \end{aligned} \tag{5.3}$$

where

$$F_1 = \sum_{n \in \boldsymbol{N}} (1 - f_n) \tag{5.4a}$$

$$F_2 = \sum_{n \in \boldsymbol{N}} \sum_{e^{ij}_{km} \in \boldsymbol{E}} (x^{nij}_{km} + y^{nij}_{km}) \tag{5.4b}$$

$$F_3 = \sum_{n \in \boldsymbol{N}} \sum_{r \in \boldsymbol{R}} \pi_r F^n_r(p^r, x) F^n_r(p^r, y) \tag{5.4c}$$

$$\tag{5.4d}$$

and w_1, w_2, and w_3 are arbitrary weights, i.e., since F_1, F_2, and F_3 are not fractional values. Specifically, the first term, F_1, represents the penalty for provisioning failures, i.e., number of requests experiencing setup failures. Note that minimizing this penalty is the same as maximizing the overall throughput. Meanwhile the second term, F_2, represents the total resource consumption for all primary and backup paths (akin to Eq. (4.2b)). Finally the last term, F_3, sums the joint failure probability (risk) for all path-pairs. Overall, Eq. (5.3) tries to achieve a weighted balance between blocking reduction, resource minimization, and overall risk minimization. Note that an INLP formulation is necessary here due to the product terms in F_3. The requisite INLP constraints are now defined in order to generate a valid solution:

$$\sum_{(j,m):e_{km}^{ij}\in \boldsymbol{E}} x_{km}^{nij} - \sum_{(j,m):e_{mk}^{ji}\in \boldsymbol{E}} x_{mk}^{nji} = \begin{cases} f_n; & if\ v_k^i = s_n \\ -f_n; & if\ v_k^i = d_n \\ 0; & otherwise \end{cases} ; n \in \boldsymbol{N} \tag{5.5}$$

$$\sum_{(j,m):e_{km}^{ij}\in \boldsymbol{E}} y_{km}^{nij} - \sum_{(j,m):e_{mk}^{ji}\in \boldsymbol{E}} y_{mk}^{nji} = \begin{cases} f_n; & if\ v_k^i = s_n \\ -f_n; & if\ v_k^i = d_n \\ 0; & otherwise \end{cases} ; n \in \boldsymbol{N} \tag{5.6}$$

$$x_{km}^{nij} + y_{km}^{nij} \leq f_n; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{5.7}$$

$$\sum_{n\in \boldsymbol{N}} (x_{km}^{nij} + y_{km}^{nij}) \leq C_2; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{5.8}$$

$$x_{km}^{nij} \in \{0,1\}; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{5.9}$$

$$y_{km}^{nij} \in \{0,1\}; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{5.10}$$

$$f_n \in \{0,1\}; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E} \tag{5.11}$$

In particular, the above constraints are the same as those defined in Section 4.1. Namely, Eqs. (5.5) and (5.6) represent flow continuity constraints, Eq. (5.7) represents the link-disjointness constraint, Eq. (5.8) restricts the total traffic carried on each link to under its capacity, and Eqs. (5.9)-(5.11) are binary constraints on their corresponding variables.

5.1.3 Approximation Solution

In general, INLP problems have high computational complexity and are very difficult to solve. Moreover, this field of optimization research in itself is not very mature. In light of this, a further ILP-based approximation is developed to provide a more scalable solution. Foremost, consider the expansion of the joint (conditional) risk

failure probability for a link-disjoint path-pair $(\boldsymbol{x}, \boldsymbol{y})$ for request n in Eq. (5.2) as follows:

$$\begin{aligned} F_r^n(p^r, x)F_r^n(p^r, y) &= (1 - \prod_{e_{km}^{ij} \in \boldsymbol{E}} (1 - p_{i_k j_m}^r x_{km}^{ij}))(1 - \prod_{e_{km}^{ij} \in \boldsymbol{E}} (1 - p_{i_k j_m}^r y_{km}^{ij})) \\ &= 1 - \prod_{e_{km}^{ij} \in \boldsymbol{E}} (1 - p_{i_k j_m}^r x_{km}^{ij}) - \prod_{e_{km}^{ij} \in \boldsymbol{E}} (1 - p_{i_k j_m}^r y_{km}^{ij}) \\ &\quad + \prod_{e_{km}^{ij} \in \boldsymbol{E}} (1 - p_{i_k j_m}^r x_{km}^{ij}) \prod_{e_{km}^{ij} \in \boldsymbol{E}} (1 - p_{i_k j_m}^r y_{km}^{ij}) \\ &= \sum_{e_{km}^{ij} \in \boldsymbol{E}} \sum_{e_{pq}^{uv} \in \boldsymbol{E}} p_{i_k j_m}^r p_{u_p v_q}^r x_{km}^{nij} y_{pq}^{nuv} + HOT \end{aligned} \tag{5.12}$$

where the *high-order terms* (HOT) represent the sum of products of three or more risk probabilities. Now assuming low conditional failure probabilities, i.e., $p_{i_k j_m}^r \ll 1, \forall e_{km}^{ij} \in \boldsymbol{E}$, the HOT can be ignored and Eq. (5.4c) can be reduced to:

$$\begin{aligned} F_3 &= \sum_{n \in \boldsymbol{N}} \sum_{r \in \boldsymbol{R}} \pi_r F_r^n(p^r, x) F_r^n(p^r, y) \\ &\approx \sum_{n \in \boldsymbol{N}} \sum_{r \in \boldsymbol{R}} \sum_{e_{km}^{ij} \in \boldsymbol{E}} \sum_{e_{pq}^{uv} \in \boldsymbol{E}} \pi_r p_{i_k j_m}^r p_{u_p v_q}^r x_{km}^{nij} y_{pq}^{nuv} \\ &= \sum_{n \in \boldsymbol{N}} \sum_{r \in \boldsymbol{R}} \sum_{e_{km}^{ij} \in \boldsymbol{E}} \sum_{e_{pq}^{uv} \in \boldsymbol{E}} \pi_r p_{i_k j_m}^r p_{u_p v_q}^r z_{u_p v_q}^{ni_k j_m} \end{aligned} \tag{5.13}$$

where $z_{u_p v_q}^{ni_k j_m}$ are new binary variables that are introduced to replace the product terms $x_{lm}^{nij} y_{pq}^{nuv}$, i.e.,

$$z_{u_p v_q}^{ni_k j_m} \geq x_{lm}^{nij} + y_{pq}^{nuv} - 1; e_{km}^{in} \in \boldsymbol{E}, e_{pq}^{uv} \in \boldsymbol{E} \tag{5.14}$$

$$z_{u_p v_q}^{ni_k j_m} \in \{0, 1\}; n \in \boldsymbol{N}, e_{km}^{ij} \in \boldsymbol{E}, e_{pq}^{uv} \in \boldsymbol{E} \tag{5.15}$$

and superscript n denotes the request index, superscripts $i_k j_m$ come from the variable x_{km}^{nij}, and subscripts $u_p v_q$ come from the variable y_{pq}^{nuv}. In addition, further ILP

Equation	$\boldsymbol{H}(\boldsymbol{U},\boldsymbol{E})$	$\boldsymbol{G}^i(\boldsymbol{V}^i,\boldsymbol{L}^i)$
Eq.(5.5)	$N\lvert\boldsymbol{U}\rvert$	$N^i(\lvert\boldsymbol{V}^i\rvert)$
Eq.(5.6)	$N\lvert\boldsymbol{U}\rvert$	$N^i(\lvert\boldsymbol{V}^i\rvert)$
Eq.(5.7)	$N\lvert\boldsymbol{E}\rvert$	$N^i\lvert\boldsymbol{L}^i\rvert$
Eq.(5.8)	$\lvert\boldsymbol{E}\rvert$	$\lvert\boldsymbol{L}^i\rvert$
Eq.(5.9)	$N\lvert\boldsymbol{E}\rvert$	$N^i\lvert\boldsymbol{L}^i\rvert$
Eq.(5.10)	$N\lvert\boldsymbol{E}\rvert$	$N^i\lvert\boldsymbol{L}^i\rvert$
Eq.(5.11)	N	N_i
Eq.(5.14)	$N\lvert\boldsymbol{E}\rvert^2$	$N^i\lvert\boldsymbol{L}^i\rvert^2$
Eq.(5.15)	$N\lvert\boldsymbol{E}\rvert^2$	$N^i\lvert\boldsymbol{L}^i\rvert^2$
total	$O(2N\lvert\boldsymbol{E}\rvert^2)$	$O(2N^i\lvert\boldsymbol{L}^i\rvert^2)$

Table 5.1: Number of equations

constraints are also added to ensure that $z_{u_pv_q}^{ni_kj_m} = 1$ if and only if both x_{km}^{ij} and y_{pq}^{uv} are 1. In other words, the joint failure risk of the pair of links e_{km}^{ij} and e_{pq}^{uv} is counted only when both links carry traffic for request n, i.e., e_{km}^{ij} is used for the primary path and e_{pq}^{uv} is used for backup path. Finally, using the revised formulation to replace the product term of two variables by a single variable, $z_{u_pv_q}^{ni_kj_m}$, the original INLP problem can be approximated by an ILP formulation as follows:

$$\begin{aligned}\text{Min } F = {} & w_1 \sum_{n\in\boldsymbol{N}}(1-f_n) + w_2 \sum_{n\in\boldsymbol{N}} \sum_{e_{km}^{ij}\in\boldsymbol{E}} (x_{km}^{nij} + y_{km}^{nij}) \\ & + w_3 \sum_{n\in\boldsymbol{N}} \sum_{r\in\boldsymbol{R}} \sum_{e_{km}^{ij}\in\boldsymbol{E}} \sum_{e_{pq}^{uv}\in\boldsymbol{E}} \pi_r p_{i_kj_m}^r p_{u_pv_q}^r z_{u_pv_q}^{ni_kj_m} \qquad (5.16) \\ & \text{subject to: Eq.(5.5)} - \text{(5.11)}, \text{(5.14)}, \text{(5.15)}\end{aligned}$$

Overall, the above ILP formulation has a total of $(N + 2N|\boldsymbol{E}| + N|\boldsymbol{E}|^2)$ variables and $O(N|\boldsymbol{E}|^2)$ equations, see Table 5.1. Namely, akin to Section 4.1, a flow variable, f_n, is defined for each request. Flow-specific variables are also defined for each link, i.e., x_{km}^{nij} and y_{km}^{nij} for working and backup paths, respectively. Finally, the variables $z_{u_pv_q}^{ni_kj_m}$ are defined for all possible link-pair combinations, i.e., $N|\boldsymbol{E}|^2$ in total. Now since the above formulation must be solved for both the inter-domain

skeleton topology as well as all intra-domain topologies, the maximum computational complexity is dominated by the topology with the maximum equation count, i.e., $\max_i\{O(2N|\boldsymbol{E}|^2), O(2N^i|\boldsymbol{L}^i|^2)\}$, where N_i represents the local request for domain i.

5.2 Solution Approach

The proposed solution also pursues a two-stage optimization approach. Namely, the first part performs optimization over the abstract graph $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ to generate a set of skeleton path-pairs, whereas the second part performs local optimization over the individual domain sub-graphs, $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$, to resolve the domain traversing routes. Complete end-to-end paths are then obtained by inserting these local segments into their corresponding skeleton routes. Now most ILP formulations are solved in a batch manner to process requests simultaneously, e.g., as per Chapters 3 and 4. However, the ILP approximation in Section 5.1.3 poses much higher complexity since the number of $z_{u_p v_q}^{n i_k j_m}$ variables is $N|\boldsymbol{E}|^2$. Hence this ILP is only solved in a *sequential* manner instead to improve its scalability. Namely, each input request is optimized in a standalone fashion, and then its assigned capacity is removed along all its traversed intra-/inter-domain links, i.e., before optimizing the next request. Now clearly, request ordering can have a notable impact here. Hence further request batch shuffling is also done and the results averaged over multiple sequential runs.

5.2.1 Hierarchical Inter-Domain Minimal-Risk Path-Pair Optimization

As noted in Section 4.2.1, enforcing link-disjointness constraints, Eq. (5.7), for all abstract links in graph $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ can lead to overly-restrictive route selection. Hence the proposed solution only applies this constraint to *physical* inter-domain links and

solves the resultant formulation using the ILP approximation (Section 5.1.3). Subsequently, the domain-traversing sub-requests are identified by extracting all abstract links on the skeleton path-pairs for intra-domain optimization, as detailed next in Section 5.2.2.

5.2.2 Local Intra-Domain Optimization

As per the regular protection case in Section 4.2.2, the second optimization step takes the sub-path requests (abstract links) from the skeleton path-pairs and performs local optimization to find their explicit route sequences. Again, consider two cases here. First, if a domain is traversed by both the primary and backup skeleton routes, then the local optimization problem is a minimal-risk path-pair problem which can be solved using the same ILP approximation as in Section 5.1.3, i.e., over the local graph $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$. However, if the local domain is only traversed by one of the skeleton paths (either primary or backup), then the problem degenerates into a *single path* optimization. Hence a revised ILP formulation is used here for this simpler case. Namely, consider the end-to-end risk for a single path calculated as follows:

$$\begin{aligned} F_3' &= \sum_{n\in\boldsymbol{N}_i}\sum_{r\in\boldsymbol{R}} \pi_r(1 - \prod_{l^{ii}_{km}\in\boldsymbol{L}^i}(1 - p^r_{i_k i m} x^{nii}_{km})) \\ &= \sum_{n\in\boldsymbol{N}_i}\sum_{r\in\boldsymbol{R}} \pi_r(1 - \prod_{l^{ii}_{km}\in\boldsymbol{L}^i}(1 - p^r_{km} x^{nii}_{km})) \\ &= \sum_{n\in\boldsymbol{N}_i}\sum_{r\in\boldsymbol{R}} \pi_r(\sum_{l^{ii}_{km}\in\boldsymbol{L}^i} p^r_{i_k i_m} x^{nii}_{km}) + HOT \end{aligned} \tag{5.17}$$

where the HOT includes all products of two or more failure probabilities. Again, assuming low conditional link failure probabilities, the above expression is dominated by the $\sum_{l^{ii}_{km}\in\boldsymbol{L}} p^r_{i_k i_m} x^{nii}_{km}$ term, resulting in a simplified ILP objective function:

$$\text{Min } F = w_1 F_1 + w_2 \sum_{n \in \boldsymbol{N}} \sum_{l^{ii}_{km} \in \boldsymbol{L}^i} x^{nii}_{km} + w_3 \sum_{n \in \boldsymbol{N}_i} \sum_{r \in \boldsymbol{R}} \pi_r (\sum_{l^{ii}_{km} \in \boldsymbol{L}^i} p^r_{i_k i_m} x^{nii}_{km}) \tag{5.18}$$

subject to Eqs.(5.5), (5.9), (5.11) as well as

$$\sum_{n \in \boldsymbol{N}} x^{nij}_{km} \leq C_1; \; n \in \boldsymbol{N}, l^{ii}_{km} \in \boldsymbol{L}^i \tag{5.19}$$

In particular, Eq. (5.19) bounds the total traffic carried on each link to under its maximum link capacity C_1. Carefully note that this formulation can also be used to compute "risk-aware" full-mesh abstractions, i.e., where the risk of each *abstract* link is represented by the minimum failure probability of all its physical path links.

Finally, once the local sub-paths (single path or path-pairs) have been computed, a MU wavelength assignment strategy is used to assign the wavelength channels. Again, a request is only successful if all of its local domain-traversing sub-paths (lightpaths) are resolved. Carefully note that link resource levels in both the hierarchical graph $\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{E})$ and the local graphs $\boldsymbol{G}^i(\boldsymbol{V}^i, \boldsymbol{L}^i)$ must be updated after each sequential ILP request optimization, i.e., available wavelengths for *all* intra- and inter-domain links. Overall, the single-path optimization is still simpler than the path-pair optimization in Section 5.1.3. Hence the worst case complexity bounds derived in Section 5.1.3 will still apply, i.e., Table 5.1.

5.3 Performance Analysis

The multi-domain protection optimization scheme is tested and compared against the advanced protection heuristic in [29]. In particular, this latter solution uses hierarchical routing to build "risk-aware" topology abstractions and then computes link-disjoint path-pairs based upon joint failure probabilities and TE cost considerations, see Appendix A.3. Now all tests are done using the 6-domain network topology

in Figure 3.8 with four multi-failure stressor regions added, as shown in Figure 5.1. Note that additional tests with larger topology sizes are not done due to excessive ILP computational complexities. In addition, two different link sizes are tested, i.e., $C_2 = 2C_1 = 16$ wavelengths and $C_2 = 2C_1 = 32$ wavelengths. Furthermore, all light-path requests are generated in a random manner with varying a-priori batch sizes. Unlike the solutions in Chapters 3 and 4, however, these requests are optimized in a *sequential* manner, i.e., in the order in which they are generated. Meanwhile, the objective function weights in Eq. (5.3) are set to w_1=6, w_2=0.0001, w_3=1 for both the intra- and inter-domain optimization levels, i.e., to emphasize throughput maximization. As per the earlier chapters, the associated ILP formulations are also solved using a combination of the PuLP and GLPK toolkits.

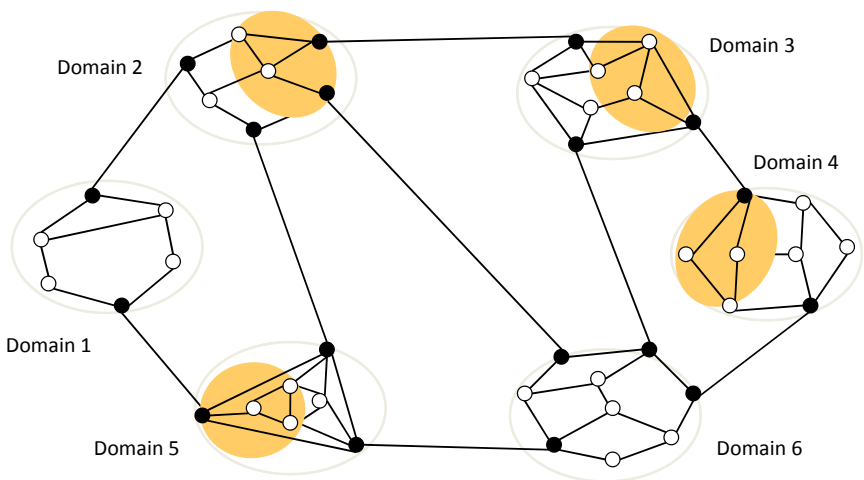

Figure 5.1: 6-domain network with 4 stressor regions (p-SRLG)

Initial tests are done for relatively high (independent) link failure probabilities

within the 4 p-SRLG regions, i.e., $p^r_{i_k j_m} = 0.8$, and all values are averaged over 5 independent runs. The associated setup success rates are first plotted in Figure 5.2 for varying batch request sizes, and results show very close performance between the optimization and heuristic strategies at lower load regimes for both link sizes. However as loads increase, the ILP (approximation) solution does much better, yielding almost 50% more successful setups. Meanwhile, Figure 5.3 also plots the *protection failure rate*, defined as the percentage of lightpaths experiencing both primary and backup failures over all successful setups. This is an important measure of resilience as it gauges service outage. These results indicate slightly lower protection failure rates for the ILP scheme for smaller request sizes, i.e., when both schemes have about the same setup success rate. However, with larger request sizes, the protection failure rate for the ILP scheme sometimes exceeds that of the heuristic strategy. This is likely due to the higher setup rate of the ILP scheme, which places more emphasis on throughput maximization. Hence more connections are established and some of them–especially those setup at higher loads–will experience higher joint-failure rates. Carefully note that the performance gap between the ILP and heuristic schemes for this metric is also more notable for lower link sizes, i.e., $C_2 = 2C_1 = 16$ wavelengths.

In addition, the number of *non-failed* lightpaths is also plotted in Figure 5.4. The results here show that the ILP scheme again gives more survivable lightpaths, i.e., almost 50% higher at increased loads. Finally, the average hop counts for the primary/backup paths are also measured and plotted in Figure 5.5. As observed earlier for the case of regular provisioning (Chapter 3) and simple link-disjoint protection provisioning (Chapter 4), the ILP scheme gives notably higher values here as compared to the heuristic strategy, i.e., ranging anywhere from 10-25%.

Next, some of the above tests are repeated for medium and low link failure probabilities, i.e., $p^r_{i_j k_m} = p = 0.4$ and 0.2, respectively. Namely, Figures 5.6 and 5.7 plot the number of successful setups for these two values, and the overall results are very

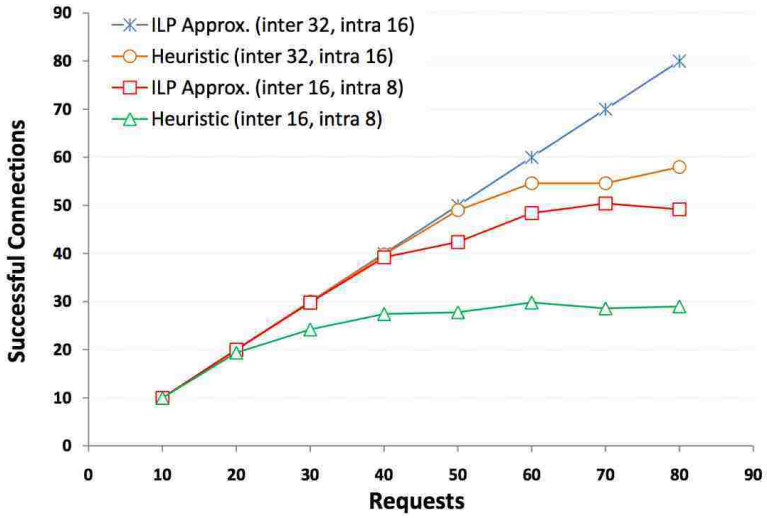

Figure 5.2: Successful requests ($p^r_{i_kj_m} = 0.8$/high)

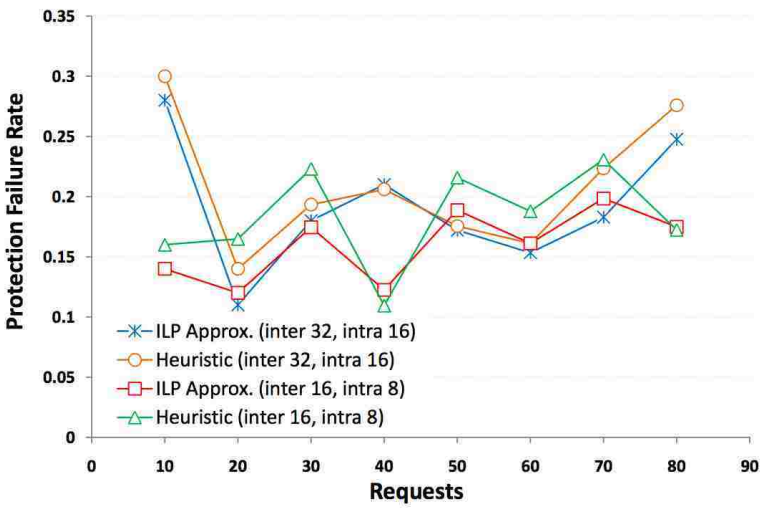

Figure 5.3: Protection failure rate ($p^r_{i_kj_m} = 0.8$/high)

similar to those for the high failure probability case, i.e., since link failure probabilities really have no effect on lightpath setup for both ILP and heuristic strategies. The respective protection failure rates for the medium and low link failure probabilities are also plotted in Figures 5.8 and 5.9, respectively. Again the performance gap

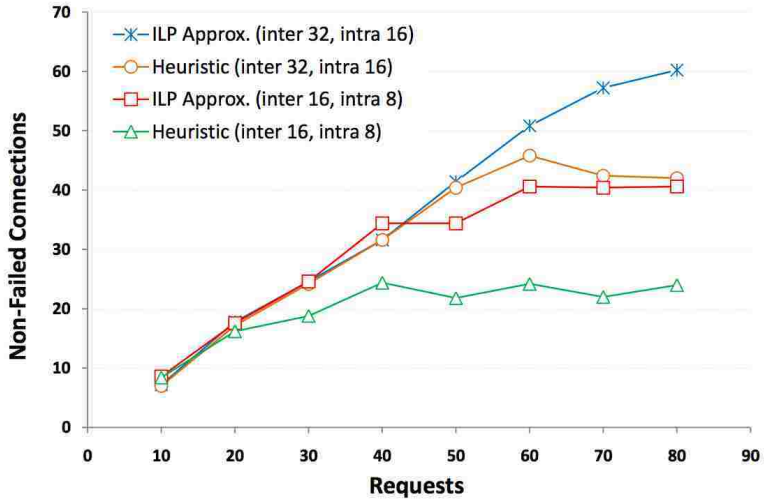

Figure 5.4: Non-failed requests ($p^r_{i_k j_m} = 0.8$/high)

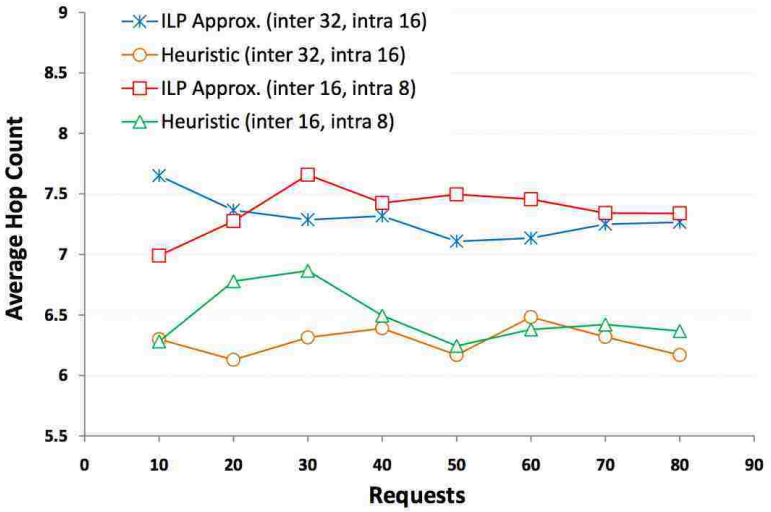

Figure 5.5: Average hop counts ($p^r_{i_k j_m} = 0.8$/high)

is more notable at lower request sizes, but becomes marginal at some higher load points. Finally, Figures 5.10 and 5.11 plot the numbers of non-failed lightpaths for the two link failure probabilities. Again, the ILP scheme gives much better lightpath survivability, particularly at higher load points, i.e., almost twice as many non-failed

lightpaths at higher loads for $C_2 = 2C_1 = 16$ wavelengths.

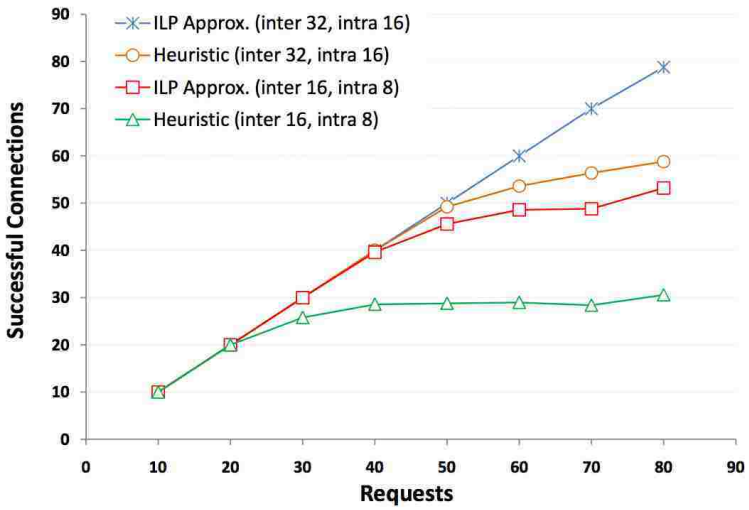

Figure 5.6: Successful requests ($p^r_{i_{kjm}} = 0.4$/medium)

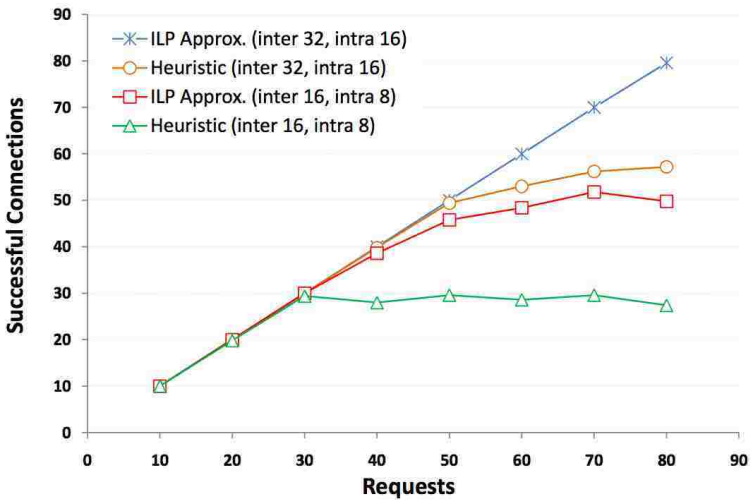

Figure 5.7: Successful requests ($p^r_{i_{kjm}} = 0.2$/low)

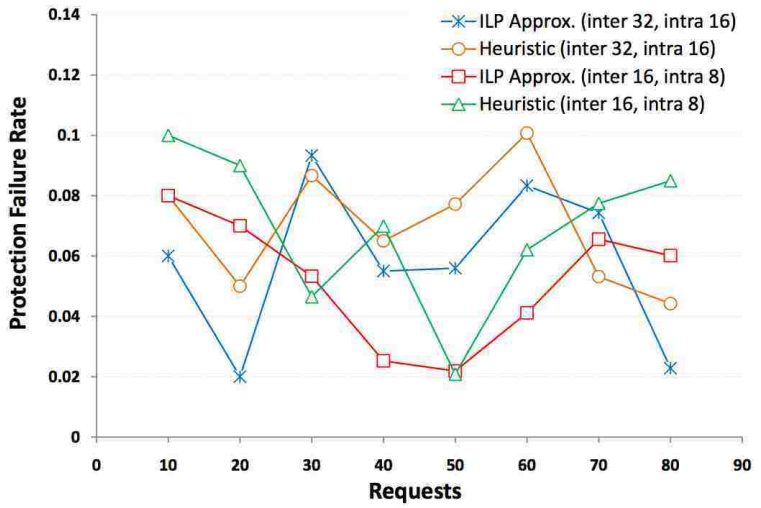

Figure 5.8: Protection failure rate ($p^r_{i_k j_m} = 0.4$/medium)

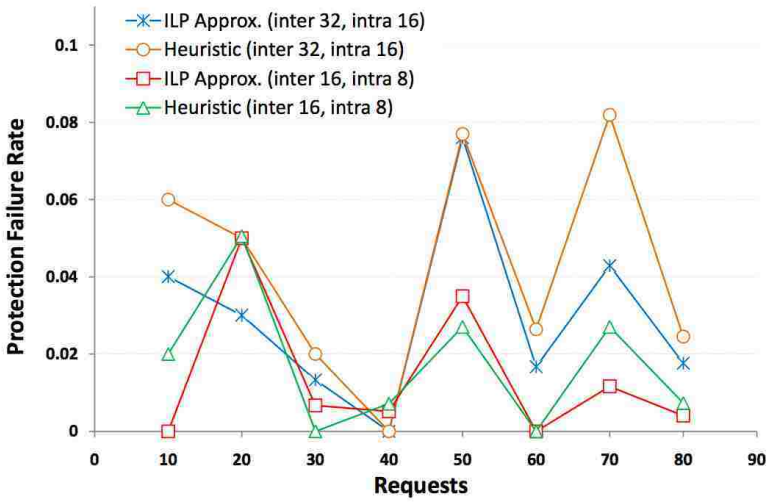

Figure 5.9: Protection failure rate ($p^r_{i_k j_m} = 0.2$/low)

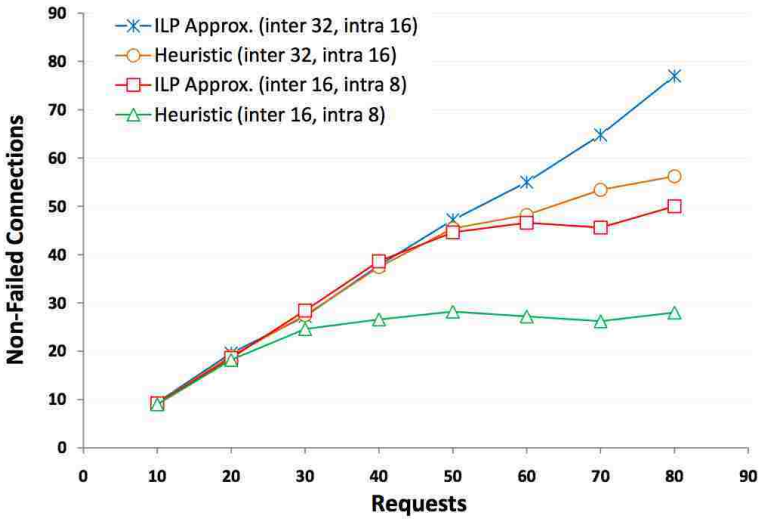

Figure 5.10: Non-failed requests ($p^r_{i_k j_m} = 0.4$/medium)

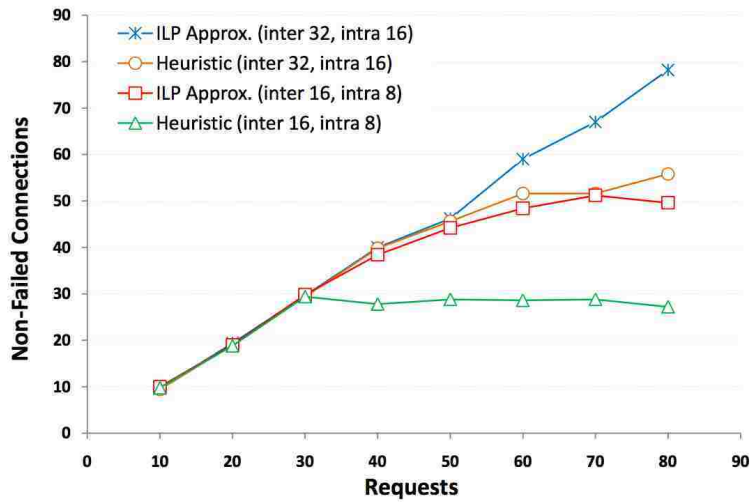

Figure 5.11: Non-failed requests ($p^r_{i_k j_m} = 0.2.$/low)

Chapter 6

Conclusions and Future Work

This dissertation presents one of the first optimization-based studies of multi-domain optical network provisioning and survivability. In particular, Chapter 2 reviews the existing background in the field, and then Chapter 3 proposes a formal ILP model for multi-domain lightpath provisioning pursuant to several TE objectives. Building upon this, Chapter 4 extends the formulation for single-link failure protection by optimizing diverse path-pairs at the intra- and inter- domain levels. Finally, Chapter 5 considers the case of multi-domain disaster recovery and presents a novel optimization model for probabilistic protection for multiple correlated failures. The conclusions from this work are now presented along with discussions on potential future research directions.

6.1 Conclusions

The research starts by addressing regular working-mode, i.e., non-survivable, lightpath provisioning optimization in Chapter 3. The solution framework models realistic settings in which all-optical domains are bounded by full opto-electronic border OXC

nodes, i.e., wavelength conversion/regeneration between domain boundaries for SLA monitoring. Furthermore, a hierarchical routing setup is also assumed, in which individual domains have limited global state knowledge in the form of full-mesh abstract topologies. Based upon these assumptions, the multi-domain lightpath provisioning problem is modeled as a two-stage ILP optimization for the ideal case of a-priori (batch) demands, i.e., pursuant to several TE objectives including throughput maximization, resource minimization, and load-balancing. Overall, this approach provides an invaluable reference against which to gauge other schemes, and the key findings from this effort reveal that:

- The ILP-based solution vastly outperforms an advanced multi-domain heuristic (using hierarchical link-state routing) [16] in terms of setup success rates. These gains range anywhere from 80-180% more accepted demands for two sample topologies and are most pronounced at higher loads (larger batch sizes). These findings also show the extent to which existing heuristics are sub-optimal and indicate much room for improvement.

- Multi-domain optimization yields slightly higher average hop counts than the heuristic scheme (by about 11%). These values also tend to decline with increasing batch request sizes.

- The ILP scheme also gives much slower utilization growth on inter-domain links for increasing input batch sizes. This indicates more efficient resource allocation versus the heuristic strategy which tends to achieve link saturation at much smaller batch sizes, i.e., 25% smaller batch request sizes for 80% maximum link utilization.

Chapter 4 extends the above optimization framework to implement multi-domain lightpath protection to support single link-failure recovery. Namely, a two-stage ILP model is proposed to optimize link-disjoint path-pairs at the skeleton inter-domain level, and these are then expanded at the local intra-domain level. Similar TE

objectives (to the working-mode only case) are also applied here and the overall results indicate that:

- The ILP-based solution outperforms an advanced multi-domain protection heuristic [63] in terms of setup success rates. In particular, findings for two sample multi-domain topologies up to 100% more accepted requests at higher loads. This also indicates much room to develop improved heuristic strategies.

- Multi-domain protection optimization also yields slightly higher average hop count values (by about 14%). These values also tend to decline with increasing input load, i.e., batch request sizes.

- The ILP model also gives much slower utilization growth on inter-domain links for increasing input batch sizes. This indicates more efficient resource allocation versus the heuristic strategy which gives much faster link saturation, i.e., 37.5% smaller batch request sizes for 100% maximum link utilization.

In general, single failure protection schemes are not very effective against large-scale disasters, i.e., since multiple network failures can easily impact both primary and backup routes. Hence improved solutions are required to directly incorporate the geographic proximity of link failures in such cases. To address this concern, Chapter 5 studies the highly-challenging case of multi-domain disaster recovery. In particular, stressors are modeled as random events, with each having a failure risk region comprised of a set of independently-failing links with close geographic proximity. Using this model, the two-stage optimization in Chapter 4 is modified to further incorporate risk minimization concerns. In particular, the problem is formulated as a more complex INLP scenario and then solved using an ILP approximation. The overall results from this study show:

- The optimization-based scheme outperforms an advanced multi-domain/multi-failure protection heuristic [29] in terms of setup success rates. In particular, the findings for a sample multi-domain topology reveal anywhere from 60-100% more accepted requests at higher loads. This indicates that there is much room to develop improved multi-failure recovery heuristics.

- The optimization solution also yields lower protection failure rates at low load regimes. However at higher loads, this approach gives very close (even slightly higher) protection failure rates due to its increased setup success rates.

- Multi-failure protection optimization gives higher number of survivable connections, ranging anywhere from 40-100% versus the heuristic strategy.

6.2 Future Work

Overall, this dissertation presents a strong basis from which to investigate further related problems in the area of multi-domain networking. Some of these avenues are now highlighted. Foremost, as noted in Chapter 1, the focus here is strictly on optimization design for optical WDM networks with wavelength (capacity, color) link-level constraints. However, multi-domain provisioning for more general (IP/MPLS and Ethernet) networks with dual bandwidth and delay link constraints is also very important and needs to be investigated from an optimization perspective. Since delay constraints are usually specified as floating point values, this effort will likely require more specialized *mixed ILP* (MILP) formulations.

Detailed research findings from this effort consistently show very significant performance gaps between the optimization and heuristic strategies, i.e., for both the survivable and non-survivable provisioning cases. As a result, future efforts can also look at developing improved heuristic methodologies to better track the optimiza-

tion results, with a particular focus on designing batch provisioning schemes. This is an area that has seen little focus to date, even within the context of single-domain networks.

Furthermore, the survivability-based optimization schemes in this dissertation only consider dedicated protection. Indeed, the more complex case of shared protection is not treated as it is generally less applicable in open multi-carrier settings, i.e., due to complex resource sharing concerns and likely policy restrictions between organizations. Nevertheless, the concept of shared protection is still viable in specialized single-carrier multi-domain networks. Along these lines, improved optimization models can be developed for resource sharing at both the inter- and intra-domain networking levels, i.e., between skeleton backup routes on the global abstract topology as well as domain-traversing path-pairs on the local network topologies.

Finally, optimization strategies can also be adapted for use in practical real-world settings. Namely, consider the fact that most requests usually arrive in a dynamic "on-demand" manner and are provisioned using heuristic algorithms. Nevertheless, many of these demands–particularly higher-bandwidth lightpath requests–tend to have relatively long holding times, on the order of weeks or even months. Hence operators can use optimization-based schemes to periodically re-optimize existing lightpath routes (that have already been provisioned in a sub-optimal heuristic manner). Overall, this approach can improve network resource efficiency and load balancing, thereby allowing operators to support increased user demands and improve their revenues. The design of appropriate batch re-optimization triggering policies will also be of particular interest here, i.e., periodic or threshold-based mechanisms.

Appendices

Appendix A

Multi-Domain Heuristic Strategies

As noted in the main dissertation, the proposed optimization schemes presented in Chapters 3-5 are tested against several advanced multi-domain heuristic strategies. These solutions use hierarchical routing to achieve an acceptable level of global topology/resource visibility, and then leverage this information to perform distributed path computation and protection. Along these lines, this Appendix presents an overview of some of these comparative heuristics, although interested readers are referred to the respective publications for complete details. In particular, a regular (non-survivable) multi-domain lightpath solution is first presented, followed by single and multi-link failure protection schemes, respectively.

A.1 Non-Survivable Provisioning

The work in Chapter 3 compares the performance of the non-survivable optimization strategy against a well-known distributed heuristic scheme in [15]. This solution assumes a hierarchical routing/provisioning architecture in which border gateway OXC nodes exchange complete information on physical inter-domain links as well as par-

tial (condensed) domain-internal information, see Figure 1.3. Now with regards to the latter, two topology abstraction schemes are proposed, i.e., simple node and full-mesh. The former restricts any domain-internal visibility exchange, whereas the latter summarizes intra-domain wavelength resource state between border nodes by computing abstract links. Namely, the available wavelength vector on a (full-mesh) abstract link is computed by searching the k-shortest paths between the respective border nodes and selecting the path with the maximum number of free wavelengths. Finally, a link-state routing protocol (hierarchical routing instance) is used to propagate updates for all physical inter-domain links and abstract links between the border gateway OXC nodes (abstract links for full-mesh abstraction only). In particular, a relative change triggering policy is used to generate updates if the change in a link's available wavelengths exceeds a *significance change factor* (SCF) and the time since its last update exceeds a *hold-down timer* (HT).

Overall, inter-domain link-state dissemination allows border OXC nodes to maintain detailed inter-domain *traffic engineering databases* (TEDB) with active "aggregated" network state, i.e., essentially a global abstract topology. This information is then used by domain-level path computing entities, e.g., such as PCE systems, to field incoming lightpath requests by first computing skeleton "end-to-end" loose routes over the abstract topology using shortest-path graph algorithms. Furthermore, two different link-weighting schemes are used here, including static (fixed) minimum hop count and dynamic load-balancing. The latter approach, in particular, is designed to avoid heavily-loaded links by setting their weights as inversely proportional to the number of free link wavelengths. Finally, these skeleton path routes are expanded using signaling to resolve the full intra-domain node sequences, i.e., via RSVP-TE protocol [5]. Namely, each intermediate domain computes a local domain-traversing lightpath segment between its ingress/egress border nodes (in the skeleton route) using lead-loaded path selection and MU wavelength selection. Overall findings show the best performance when coupling the heavier full-mesh topology

abstraction approach with load-balancing path provisioning, see [15] and also [16] for complete details.

A.2 Single-Failure Protection

Chapter 4 evaluates the proposed survivable optimization scheme against the recent multi-domain protection heuristic presented in [63]. This latter solution also uses a hierarchical routing/provisioning framework in which border gateway OXC nodes exchange inter-domain and abstract link-state information. However, in order to limit routing overheads/complexity, this heuristic approach re-uses the same (simple node, full-mesh) topology abstractions from [15]. Overall, this contrasts with other multi-domain protection strategies [24], [25], [26] which develop more complex abstractions to extract domain-internal path diversity state, i.e., to assist with subsequent link-disjoint backup path computation. Instead, the strategy in [63] relies upon signaling expansion procedures to achieve protection path diversity at the inter- and intra-domain levels. Consider some more details here.

Using the above routing information, domain-level path computation entities field incoming lightpath protection requests by computing diverse primary/backup skeleton path-pairs. In particular, two different strategies are proposed here, i.e., *domain-disjoint* (DD) and *link-disjoint* (LD). The former computes primary/backup skeleton routes without any common physical inter-domain links and intermediate domains (as minimal overlap is necessary at the source and destination domains). Namely, this is done by computing k-shortest link-disjoint path-pairs (with no intermediate domain overlaps) and selecting the pair with the minimum aggregate cost. Meanwhile the latter scheme is less restrictive and allows primary/backup skeleton paths to traverse common intermediate domains. However, this scheme cannot guarantee intra-domain link disjointness at overlapping domains due to limited domain-internal

visibility from the simpler topology abstraction. Hence sequential signaling is used to fully resolve all local domain-traversing segments and ensure primary/backup path diversity. Furthermore, Suurballe's path-pair algorithm is also used during the local expansion phase to overcome potential trap topology concerns. As per [15], both of these schemes can be further coupled with static (minimum hop) or dynamic (load-balancing) link weight selection. Overall findings show the best blocking performance when the link-disjoint scheme is coupled with load-balancing operation, see [63] for complete details. Hence this particular variant is used for comparison purposes in Chapter 4. However, note that the original scheme in [63] is presented for regular bandwidth-provisioning networks and not optical wavelength routing networks. Hence slight modifications are made to perform wavelength assignment using the MU policy after local domain path (path-pair) expansion, i.e., all-optical domains.

A.3 Multi-Failure Protection

The study in Chapter 5 focuses on probabilistic protection optimization across multiple domains and leverages a newly-proposed heuristic scheme in [29] for comparison purposes. Namely, this scheme uses the same probabilistic multi-failure p-SRLG model proposed in [68] and re-used in Section 5.1.1. Here large disaster events are represented by an a-priori set of static stressor regions, with each event having a fixed occurrence probability and vulnerable link set (independent failure probabilities). Using this failure model, a novel "risk-aware" full-mesh topology abstraction scheme is first derived to extract and propagate p-SRLG state information at the inter-domain level, i.e., by computing failure probabilities for the abstract links, see [29] for details. Note that this contrasts with the single-failure multi-domain protection heuristic in [15] which does not introduce any specialized abstraction state for recovery purposes (as detailed above in Appendix A.2).

Leveraging this new inter-domain routing state, two different "risk-aware" skeleton path-pair algorithms are then proposed in [29]. Namely, the first algorithm implements a greedy solution which computes a minimum-risk primary skeleton path and then prunes it to re-compute a minimum-risk backup skeleton path. However, studies have shown that pure risk-minimization strategies can yield very lengthy routes, thereby causing high resource inefficiency and blocking, e.g., see single-domain study in [89]. Hence a second heuristic solution is also proposed to jointly incorporate both risk minimization and TE efficiency concerns into the skeleton path-pair computation. Namely, k-shortest (skeleton) paths are first computed over the abstract topology along with their respective link-disjoint backup paths. Subsequently, the skeleton primary/backup path-pair with the lowest failure probability dot product is selected for distributed signaling expansion, see [16] for complete details. Overall results show much improved resource efficiency (and even multi-failure recovery rates) with the latter joint risk/TE scheme. As a result, this specific variant is used for comparison purposes in Chapter 5. Again, the original scheme in [29] is presented for regular bandwidth-provisioning networks and not optical wavelength routing networks. Hence slight modifications are also made to perform wavelength assignment using the MU policy after local domain path (path-pair) expansion.

References

[1] N. Rao, W. Wing, S. Carter, and Q. Wu, "Ultrascience Net: Network Testbed for Large-Scale Science Applications," *IEEE Communications Magazine*, vol. 43, no. 11, pp. S12–S17, 2005.

[2] "Terabit Networks for Extreme Scale Science," United States Department of Energy workshop report, February 2011. [Online]. Available: http://science.energy.gov/ascr/$\sim$/media/ascr/pdf/programdocuments/docs/Terabit_Networks_Workshop_Report.pdf

[3] G. Bell, "Network as Instrument–Faster Data for California," 2013. [Online]. Available: http://www.cenic.org/publications/cenictoday/1Q2013/1Q2013CT_3.html

[4] A. Kasim, P. Adhikari, N. Chen, N. Finn, N. Ghani, M. Hajduczenia, P. Havala, G. Heron, M. Howard, and L. Martini, *Delivering Carrier Ethernet: Extending Ethernet Beyond the LAN*, 1st ed. New York, NY: McGraw-Hill, Inc., 2008.

[5] G. Bernstein, B. Rajagopalan, and D. Saha, *Optical Network Control: Architecture, Protocols, and Standards*. Boston, MA: Addison-Wesley, 2003.

[6] B. Mukherjee, *Optical WDM Networks (Optical Networks)*. New York, NY: Springer, 2006.

[7] A. Farrel *et al.*, "A Path Computation Element (PCE)-Based Architecture," *IETF RFC 4655*, August 2006.

[8] N. Ghani, M. Peng, and A. Rayes, "Provisioning and Survivability in Multi-Domain Optical Networks," in *WDM Systems and Networks*, N. N. Antoniades, G. Ellinas, and I. Roudas, Eds. Springer New York, 2012, pp. 481–519.

[9] M. Chamania and A. Jukan, "A Survey of Inter-Domain Peering and Provisioning Solutions for the Next Generation Optical Networks," *IEEE Communications Survey and Tutorials*, vol. 11, no. 1, March 2009.

[10] "Moving Towards Terabit/Sec Scientific Dataset Transfers." California Institute of Technology (CalTech) Press Release, SuperComputing (SC) 2012, October 2012.

[11] T. Korkmaz and M. Krunz, "Source-Oriented Topology Aggregation with Multiple QoS Parameters in Hierarchical Networks," *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, vol. 10, no. 4, pp. 295–325, October 2000.

[12] B. St Arnaud *et al.*, "BGP Optical Switches and Lightpath Route Arbiter," *Optical Networks*, vol. 2, no. 2, pp. 73–81, March/April 2001.

[13] M. Blanchet, F. Parent, and B. St Arnaud, "Optical BGP (OBGP): Inter AS Lightpath Provisioning," *IETF Draft draft-parent-obgp-01.txt*, March 2001.

[14] S. Sanchez-Lopez *et al.*, "A Hierarchical Routing Approach for GMPLS-Based Control Plane for ASON," in *IEEE International Conference on Communications (ICC)*, Seoul, Korea, June 2005.

[15] Q. Liu, M. Kok, N. Ghani, and A. Gumaste, "Hierarchical Routing in Multi-Domain Optical Networks," *Computer Communications*, vol. 30, no. 1, pp. 122–131, December 2006.

[16] Q. Liu, N. Ghani, N. Rao, A. Gumaste, and M. Garcia, "Distributed Inter-Domain Lightpath Provisioning in the Presence of Wavelength Conversion," *Computer Communications*, vol. 30, no. 18, pp. 3362–3375, December 2007.

[17] O. Yu, "Intercarrier Interdomain Control Plane for Global Optical Networks," in *IEEE Internatioanl Conference on Communications (ICC)*, New York City, NY, June 2004.

[18] M. Yannuzzi *et al.*, "Interdomain RWA Based on Stochastic Estimation Methods and Adaptive Filtering for Optical Networks," in *IEEE Global Communications Conference (GLOBECOM)*, San Francisco, CA, November 2006.

[19] M. Yannuzzi *et al.*, "Toward A New Route Control Model for Multi-Domain Optical Networks," *IEEE Communications Magazine*, vol. 46, no. 6, pp. 104–111, June 2008.

[20] T. Takeda, E. Oki, and K. Shiomoto, "Diverse Path Setup Schemes in Multi-Domain Optical Networks," in *IEEE International Conference on Broadband Communications, Networks and Systems (BROADNETS)*, London, UK, September 2008.

[21] J. Vasseur *et al.*, "A Backward Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineered Label Switched Paths," *IETF RFC 5441*, April 2009.

[22] M. Rahnamay-Naeini, J. Pezoa, G. Azar, N. Ghani, and M. Hayat, "Modeling Stochastic Correlated Failures and their Effects on Network Reliability," in *IEEE International Conference on Computer Communications and Networks (ICCCN)*, Maui, HI, August 2011.

[23] N. Ghani, Q. Liu, D. Benhaddou, N. Rao, and T. Lehman, "Control Plane Design in Multidomain/Multilayer Optical Networks," *IEEE Communications Magazine*, vol. 46, no. 6, pp. 78–87, 2008.

[24] A. Sprintson, M. Yannuzzi, A. Orda, and X. Masip-Bruin, "Reliable Routing with QoS Guarantees for Multi-Domain IP/MPLS Networks," in *IEEE International Conference on Computer Communications (INFOCOM)*, Anchorage, AK, May 2007.

[25] D. Truong and B. Thiongane, "Dynamic Routing for Shared Path Protection in Multidomain Optical Mesh Networks," *OSA Journal of Optical Networking*, vol. 5, no. 1, pp. 58–74, January 2006.

[26] D. Truong and B. Jaumard, "Using Topology Aggregation for Efficient Shared Segment Protection Solutions in Multi-Domain Networks," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 9, pp. 96–107, 2007.

[27] C. Huang and D. Messier, "A Fast and Scalable Inter-Domain MPLS Protection Mechanism," *Journal of Communications and Networks*, vol. 6, no. 1, pp. 1–8, March 2004.

[28] F. Ricciato, U. Monaco, and D. Ali, "Distributed Schemes for Diverse Path Computation in Multidomain MPLS Networks," *IEEE Communications Magazine*, vol. 43, no. 6, pp. 138–146, 2005.

[29] F. Xu, N. Min-Allah, S. Khan, and N. Ghani, "Diverse Routing in Multi-Domain Optical Networks with Correlated and Probabilistic Multi-Failures," in *IEEE International Conference on Communications (ICC) Workshop on New Trends in Optical Network Survivability*, Ottawa, Canada, June 2012.

[30] F. Xu, M. Peng, M. Esmaeili, N. Ghani, and A. Rayes, "Multi-Domain Restoration with Crankback in IP/MPLS Networks," *Optical Switching and Networking*, vol. 8, no. 1, pp. 68 – 78, 2011.

[31] F. Xu, T. Das, M. Peng, and N. Ghani, "Evaluation of Post-Fault Restoration Strategies in Multi-Domain Networks," in *IEEE International Symposium on Advanced Networks and Telecommunication Systems (ANTS)*, Bombay, India, December 2010.

[32] H. Zang, J. Jue, and B. Mukherjee, "A Review of Routing and Wavelength Assignment Approaches for Wavelength-Routed Optical WDM Networks," *Optical Networks Magazine*, vol. 1, no. 1, pp. 47–60, January 2000.

[33] R. Ramaswami and K. Sivarajan, "Routing and Wavelength Assignment in All-Optical Networks," *IEEE/ACM Transactions on Networking*, vol. 3, no. 5, pp. 489–500, October 1999.

[34] K. Christodoulopoulos, K. Manousakis, and E. Varvarigos, "Comparison of Routing and Wavelength Assignment Algorithms in WDM Networks," in *IEEE Global Communications Conference (GLOBECOM)*, New Orleans, LA, December 2008.

[35] A. Farkas, J. Szigeti, and T. Cinkler, "P-Cycle Based Protection Schemes for Multi-Domain Networks," in *IEEE International Workshop on Design of Reliable Communication Networks (DRCN)*, Island of Ischia, Naples, Italy, October 2005.

[36] J. Szigeti, R. Romeral, T. Cinkler, and D. Larrabeiti, "P-Cycle Protection in Multi-Domain Optical Networks," *Photonic Network Communications*, vol. 17, no. 1, pp. 35–47, 2009.

[37] B. Jaumard, D. Kien, M. Toulouse, and A. Hoang, "P-Cycle Based Protection Mechanisms in Multi-Domain Optical Networks," in *IEEE International Conference on Communications and Electronics (ICCE)*, Hue Royal City, Vietnam, August 2012.

[38] J. Moy, *OSPF: Anatomy of an Internet Routing Protocol*. Boston, MA: Addison-Wesley, 1998.

[39] A. Farrel and I. Bryskin, *GMPLS: Architecture and Applications*, 1st ed. San Francisco, CA: Morgan Kaufmann, 2006.

[40] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels," *IETF RFC 3209*, December 2001.

[41] A. Farrel *et al.*, "Crankback Signaling Extensions for MPLS and GMPLS RSVP-TE," *IETF RFC 4920*, July 2007.

[42] J. Vasseur and J. L. Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)," *IETF RFC 5440*, March 2007.

[43] G. Bernstein, D. Cheng, and D. Pendarakis, "Domain to Domain Routing using GMPLS, OSPF Extension V1.1," *Optical Internetworking Forum (OIF)*, July 2002.

[44] "Architecture for the Automatic Switched Optical Networks (ASON)," *ITU-T Rec. G. 8080/Y.1304*, 2001.

[45] T. Zhang, Y. Cui, Y. Zhao, L. Fu, and T. Korkmaz, "Scalable BGP QoS Extension with Multiple Metrics," in *IEEE International Conference on Networking and Services (ICNS)*, Washington, DC, 2006.

[46] W. Xu and J. Rexford, "MIRO: Multi-Path Interdomain Routing," in *ACM SIGCOMM*, Pisa, Italy, September 2006.

[47] X. Yang and B. Ramamurthy, "Inter-Domain Dynamic Routing in Multi-Layer Optical Transport Networks," in *IEEE Global Communications Conference (GLOBECOM)*, San Francisco, CA, December 2003.

[48] F. Hao and E. Zegura, "On Scalable QoS Routing: Performance Evaluation of Topology Aggregation," in *IEEE Conference on Computer Communications (INFOCOM)*, Tel Aviv, Israel, March 2003.

[49] B. Awerbuch and Y. Shavitt, "Topology Aggregation for Directed Graphs," *IEEE/ACM Transactions on Networking*, vol. 9, no. 2, pp. 82–90, February 2001.

[50] K. Liu, K. Nahrstedt, and S. Chen, "Routing with Topology Abstraction in Delay-Bandwidth Sensitive Networks," *IEEE/ACM Transactions on Networking*, vol. 12, no. 1, pp. 17–29, February 2004.

[51] G. Maier, C. Busca, and A. Pattavina, "Topology-Information Periodic Updates in Multi-Domain ASON Networks with Topology Aggregation," *Fiber and Integrated Optics*, vol. 27, no. 4, pp. 265–277, 2008.

[52] Q. Liu *et al.*, "Routing Scalability in Multi-Domain DWDM Networks," *Photonic Network Communications*, vol. 17, no. 1, pp. 63–74, 2009.

[53] Y. Yu *et al.*, "On the Efficiency of Inter-Domain State Advertising in Multi-Domain Networks," in *IEEE Global Communications Conference (GLOBECOM)*, Honolulu, HI, November 2009.

[54] J. Vasseur, A. Ayyangar, and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)," *IETF RFC 5152*, February 2008.

[55] S. Dasgupta, J. De Oliveira, and J. Vasseur, "Path-Computation-Element-Based Architecture for Interdomain MPLS/GMPLS Traffic Engineering: Overview and Performance," *IEEE Network*, vol. 21, no. 4, pp. 38–45, July/August 2007.

[56] Y. Zhao *et al.*, "Routing and Wavelength Assignment Problem in PCE-Based Wavelength-Switched Optical Networks," *Journal of Optical Communication and Networks*, vol. 2, no. 2, pp. 196–205, April 2010.

[57] R. Casellas, *et al.*, "Enhanced Backwards Recursive Path Computation for Multi-area Wavelength Switched Optical Networks Under Wavelength Continuity Constraint," in *IEEE/OSA Optical Fiber Communication Conference (OFC)*, San Diego, CA, March 2009.

[58] F. Xu, M. Esmaeili, C. Xie, N. Ghani, M. Peng, and Q. Liu, "Enhanced Crankback Signaling for Multi-Domain Traffic Engineering," in *IEEE International Conference on Communications (ICC)*, Cape Town, South Africa, May 2010.

[59] M. Esmaeili, F. Xu, N. Ghani, C. Xie, M. Peng, and Q. Liu, "Enhanced Crankback for Lightpath Setup in Multi-Domain Optical Networks," *IEEE Communications Letters*, vol. 14, no. 5, pp. 480–482, 2010.

[60] P. Cholda, A. Mykkeltveit, B. Helvik, O. Wittner, and A. Jajszczyk, "A Survey of Resilience Differentiation Frameworks in Communication Networks," *IEEE Communications Surveys Tutorials*, vol. 9, no. 4, pp. 32–55, 2007.

[61] S. Seetharaman, A. Jukan, and M. Ammar, "End-to-End Dedicated Protection in Multi-Segment Optical Networks." [Online]. Available: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.131.3016\&rep=rep1\&type=pdf

[62] J. W. Suurballe, "Disjoint Paths in a Network," *Networks*, vol. 4, no. 2, pp. 125–145, 1974.

[63] F. Xu, F. Gu, H. Alazemi, M. Peng, and N. Ghani, "Novel Path Protection Scheme for Multi-Domain Networks," in *IEEE Global Communications Conference (GLOBECOM) Joint Workshop on Complex Networks and Pervasive Group Communication*, Houston, TX, December 2011.

[64] B. Jaumard and D. Truong, "Backup Path Re-optimizations for Shared Path Protection in Multi-Domain Networks," in *IEEE Global Communications Conference (GLOBECOM)*, San Francisco, CA, December 2006.

[65] D. Truong and B. Jaumard, "Recent Progress in Dynamic Routing for Shared Protection in Multidomain Networks," *IEEE Communications Magazine*, vol. 46, no. 6, pp. 112–119, 2008.

[66] M. Motiwala, M. Elmore, N. Feamster, and S. Vempala, "Path Splicing," in *ACM SIGCOMM*, Seattle, WA, 2008.

[67] N. Wang, Y. Guo, K. Ho, M. Howarth, and G. Pavlou, "Fast Network Failure Recovery Using Multiple BGP Routing Planes," in *IEEE Global Communications Conference (GLOBECOM)*, Honolulu, HI, December 2009.

[68] H. Lee, E. Modiano, and K. Lee, "Diverse Routing in Networks With Probabilistic Failures," *IEEE/ACM Transactions on Networking*, vol. 18, no. 6, pp. 1895–1907, 2010.

[69] Q. She, X. Huang, and J. Jue, "Maximum Survivability Under Multiple Failures," in *IEEE/OSA Optical Fiber Communication Conference (OFC)*, Los Angeles, CA, March 2006.

[70] D. Papadimitriou, S. Dharanikota, R. Hartani, D. Griffith, and Y. Xue, "Inference of Shared Risk Link Groups," *IETF Draft draft-many-inference-srlg-02.txt*, May 2002.

[71] F. Paolucci, F. Cugini, L. Valcarenghi, and P. Castoldi, "Enhancing Backward Recursive PCE-Based Computation (BRPC) for Inter-Domain Protected LSP Provisioning," in *Optical Fiber communication/National Fiber Optic Engineers Conference (OFC/NFOEC)*, San Diego, CA, February 2008.

[72] D. Staessens, D. Colle, U. Lievens, M. Pickavet, P. Demeester, W. Colitti, A. Nowe, K. Steenhaut, and R. Romeral, "Enabling High Availability Over Multiple Optical Networks," *IEEE Communications Magazine*, vol. 46, no. 6, pp. 120–126, 2008.

[73] L. Guo, "LSSP: A Novel Local Segment-Shared Protection for Multi-Domain Optical Mesh Networks," *Computer Communications*, vol. 30, no. 8, pp. 1794 – 1801, 2007.

[74] F. Xu, M. Peng, M. Esmaeili, M. Rahnamay-Naeini, S. Khan, N. Ghani, and M. Hayat, "Post-Fault Restoration in Multi-Domain Networks With Multiple Failures," in *IEEE Military Communications Conference (MILCOM)*, San Jose, November 2010.

[75] W. Grover and D. Stamatelakis, "Cycle-Oriented Distributed Preconfiguration: Ring-like Speed With Mesh-like Capacity for Self-planning Network Restoration," in *IEEE International Conference on Communications (ICC)*, Atlanta, GA, June 1998.

[76] W. Grover and et al, "Design of Survivable Networks Based on p-Cycles," *Handbook of Optimization in Telecommunications*, pp. 391–434, 2006.

[77] J. Doucette, D. He, W. Grover, and O. Yang, "Algorithmic Approaches for Efficient Enumeration of Candidate p-Cycles and Capacitated p-Cycle Network Design," in *IEEE International Workshop on Design of Reliable Communication Networks (DRCN)*, Banff, Alberta, Canada, October 2003.

[78] A. Kodian and W. Grover, "Failure-Independent Path-Protecting p-Cycles: Efficient and Simple Fully Preconnected Optical-Path Protection," *IEEE Journal of Lightwave Technology*, vol. 23, no. 10, pp. 3241–3259, 2005.

[79] H. Drid, B. Cousin, M. Molnar, and N. Ghani, "Graph Partitioning for Survivability in Multi-Domain Optical Networks," in *IEEE International Conference on Communications (ICC)*, Cape Town, South Africa, May 2010.

[80] A. Chiu, G. Choudhury, R. Doverspike, and G. Li, "Restoration Design in IP Over Reconfigurable All-Optical Networks," in *IFIP International Conference on Network and Parallel Computing (NPC)*, Dalian, Liaoning, China, September 2007.

[81] R. Krishnaswamy and K. Sivarajan, "Algorithms for Routing and Wavelength Assignment Based on Solutions of LP-Relaxations," *IEEE Communications Letters*, vol. 5, no. 10, pp. 435–437, 2001.

[82] J. Crichigno, W. Shu, and M. Wu, "Throughput Optimization and Traffic Engineering in WDM Networks Considering Multiple Metrics," in *IEEE International Conference on Communications (ICC)*, Cape Town, South Africa, May 2010.

[83] S. Yuan and J. Jue, "Dynamic Lightpath Protection in WDM Mesh Networks Under Wavelength-Continuity and Risk-Disjoint Constraints," in *IEEE Global Communications Conference (GLOBECOM)*, Dallas, TX, December 2004.

[84] D. Xu, Y. Xiong, C. Qiao, and G. Li, "Trap Avoidance and Protection Schemes in Networks With Shared Risk Link Groups," *IEEE Journal of Lightwave Technology*, vol. 21, no. 11, pp. 2683–2693, 2003.

[85] Y. Zhu, A. Jukan, and M. Ammar, "Multi-Segment Wavelength Routing in Large-Scale Optical Networks," in *IEEE International Conference on Communications (ICC)*, Anchorage, AL, May 2003.

[86] D. Banerjee and B. Mukherjee, "Wavelength-Routed Optical Networks: Linear Formulation Resource Budgeting Tradeoffs, and a Reconfiguration Study," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 598–607, May 2000.

[87] J. Crichigno, J. Khoury, W. Shu, M. Wu, and N. Ghani, "Dynamic Routing Optimization in WDM Networks," in *IEEE Global Communications Conference (GLOBECOM) 2010*, Miami, FL, December 2010.

[88] R. Asthana, Y. Singh, and W. Grover, "p-Cycles: An Overview," *IEEE Communications Surveys Tutorials*, vol. 12, no. 1, pp. 97–111, 2010.

[89] O. Diaz, F. Xu, N. Min-Allah, M. Khodeir, M. Peng, S. Khan, and N. Ghani, "Network Survivability for Multiple Probabilistic Failures," *IEEE Communications Letters*, vol. 16, no. 8, pp. 1320–1323, 2012.