Modelling and Design of Resilient Networks under Challenges

By

Egemen K. Çetinkaya

Submitted to the graduate degree program in Electrical Engineering & Computer Science and the Graduate Faculty of the University of Kansas in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Chairperson: Prof. James P.G. Sterbenz

Prof. Dr. Georg Carle

Prof. Tyrone E. Duncan

Prof. Victor S. Frost

Prof. Bo Luo

Prof. Deep Medhi

Prof. Gary J. Minden

Date Defended: 2-December-2013

The Dissertation Committee for Egemen K. Çetinkaya
certifies that this is the approved version of the following dissertation:

Modelling and Design of Resilient Networks under Challenges

Chairperson: Prof. James P.G. Sterbenz

Date approved: 2-December-2013

Abstract

Communication networks, in particular the Internet, face a variety of challenges that can disrupt our daily lives resulting in the loss of human lives and significant financial costs in the worst cases. We define challenges as external events that trigger faults that eventually result in service failures. Understanding these challenges accordingly is essential for improvement of the current networks and for designing Future Internet architectures. This dissertation presents a taxonomy of challenges that can help evaluate design choices for the current and Future Internet. Graph models to analyse critical infrastructures are examined and a multilevel graph model is developed to study interdependencies between different networks. Furthermore, graph-theoretic heuristic optimisation algorithms are developed. These heuristic algorithms add links to increase the resilience of networks in the least costly manner and they are computationally less expensive than an exhaustive search algorithm. The performance of networks under random failures, targeted attacks, and correlated area-based challenges are evaluated by the challenge simulation module that we developed. The GpENI Future Internet testbed is used to conduct experiments to evaluate the performance of the heuristic algorithms developed.

Page left intentionally blank.

Acknowledgments

I would like to sincerely thank my advisor, Professor Dr. James P.G. Sterbenz, for his support during my studies. I owe my deepest gratitude to Dr. Sterbenz as I learned substantial knowledge from him in the field of *communication networks* and his teachings in the non-technical areas also shaped my thinking and became inspirational since we started working together. I thank the committee members: Prof. Dr. Georg Carle, Dr. Tyrone E. Duncan, Dr. Victor S. Frost, Dr. Bo Luo, Dr. Deep Medhi, and Dr. Gary J. Minden for their valuable feedback for improvement of this dissertation.

I feel very fortunate to have worked with Mohammed J.F. Alenazi during the implementation of the topology design and optimisation algorithms. Mohammed implemented the algorithms presented in Chapter 5 using the Python programming language, while we worked together developing and testing the algorithms and coauthoring papers. It was a pleasure to work with Andrew M. Peck during the development of the multilevel graph model presented in Section 4.3. Andrew implemented the multilevel graph model in Python, and we developed and tested this framework and coauthored papers together. The KU-TopView (KU Topology Viewer) graphical user interface developed by Justin P. Rohrer and Parker Riley was helpful in analysing the graphs visually. A preliminary version of the KU-CSM framework was developed by Rabat Anam Mahmood. It was a pleasure to work with Amit Dandekar and Justin P. Rohrer to significantly improve this framework. Amit specifically improved the area-based challenge models by incorporating Computational Geometry Algorithms Library (CGAL) into the KU-CSM framework. Dongsheng Zhang's help during the experimentation setup on GpENI testbed was outstanding. I would like to thank Yufei Cheng for helping generating the tedious adjacency matrices. Tools generated by Abdul Jabbar, Justin P. Rohrer, Mohammed J.F. Alenazi, and Yufei Cheng were extremely useful to complete some of the repetitious tasks. I also would like to thank Prof. David Hutchison, Dr. Paul Smith, Dr. Marcus Schöller, and the other members of the ResiliNets group for their work on the resilient network archi-

tecture, which is foundational to this work. I would like to thank the Information and Telecommunication Technology Center (ITTC) network system administrators and the ITTC administrative staff for their support during the development of this dissertation. Michael Hulet, Wesley Mason, Charles Henry, and Paul Calnon were always helpful in assisting whatever computing problems I faced.

I received support during my doctoral studies through of my advisor's grants by the NSF FIND (Future Internet Design) Program under grant CNS-0626918 (Postmodern Internet Architecture), NSF grant CNS-1219028 (Resilient Network Design for Massive Failures and Attacks), EU FP7 FIRE Programme ResumeNet project (grant agreement no. 224619), NSF GENI program (GPO contract no. 9500009441), NSF grant CNS-1050226 (Multilayer Network Resilience Analysis, and Experimentation on GENI), by the Battelle Institute under contract number NFP0069666: Interdomain Resilience, and T&E/S&T Program through the Army PEO STRI Contracting Office, contract number W900KK-09-C-0019 for AeroNP and AeroTP: Aeronautical Network and Transport Protocols for iNET (ANTP). I was partially supported by a International Foundation for Telemetering (IFT) fellowship and by teaching assistant positions at the Department of Electrical Engineering and Computer Science of the University of Kansas.

Early in my career at Sprint, my manager Patricia August's support to start the doctoral studies at the University of Kansas was encouraging. I am also indebted to my colleagues for their support in the early years of my career: Ernest Alvarez, Scott Wilder, Hui Wang, Mike Church, Dwight Doan, John Huff, Ben Watson, and Young Zhao at Sprint and Greg Hull and Cody Heinrich at Nortel Networks. I am also delighted to have the support of my friends especially: Koray Sarandal, Atacan Kadıoğlu, Pir A. Shah, Burak Gökoğlu, Buğra Çankaya, Nejat Sarandal, and Plarent Tirana.

Finally, I am very grateful to my dear family for their never-ending support. I feel blessed to have my brother Erdem A. Çetinkaya and my parents Bekir and Engin A. Çetinkaya.

Contents

List of Figures

List of Tables

Page left intentionally blank.

Chapter 1

Introduction and Motivation

Communication networks enable us to exchange information globally and are considered to be a critical infrastructure [8]. As society's dependence on communication networks in general and the Internet in particular increases, a disruption in the communication system has greater effects on the users. The impact of communication network disruptions has been observed as mostly financial losses [9, 10]. The financial impact of a malicious activity can be on the order of billions of dollars, as in the case of the Code Red worm attack [11]. Communication network disruptions also have the potential to result in human losses [12].

A *challenge* is a characteristic or condition that may manifest as an adverse event or condition that impacts normal operation for which the network is designed [1]. A challenge triggers *faults*, and a fault may manifest itself as an *error*. The error may propagate to cause delivered services to *fail* [13–16]. In the context of communication networks, some of these challenges include: human errors, malicious attacks, large-scale disasters, environmental challenges, unusual but legitimate traffic, infrastructure dependencies, and socio-political and economical factors [1, 2, 5, 17–27].

Networks in general, and the Internet in particular, are prone to perturbations. The Internet's susceptibility to disruptions was emphasised even in the early developmental

phases [28]. Understanding these challenges can help designing resilient network protocols and architectures. On the other hand, identification and categorisation of network challenges has been considered difficult [29, 30]. Categorising communication network challenges can help us understand the impact of disruptions, improve existing network resilience, as well as aid in designing the Future Internet architectures and protocols.

The Internet can be examined at the physical, IP, router, PoP (point of presence), and AS (autonomous system) level from a topological point of view [31]. At the bottom is the physical topology consisting of elements such as fibre and copper cables, point-to-point wireless links, ADMs (add drop multiplexers), cross-connects, and layer-2 switches. The logical level consists of devices operating at the IP-layer. A PoP is a collection of routers in a geographic location, and PoP-level topology is the interconnection of the PoPs. The AS-level topology represents how provider networks peer with each other at IXPs (Internet exchange points) and private peering points [32]. Understanding the evolution of the Internet from a multilevel point of view is more realistic than examining its properties at individual levels. On the other hand, the primary focus of previous studies has been on the logical aspects of the topology, since tools have been developed to collect, measure, and analyse IP-level properties of the Internet (e.g. Rocketfuel [33]). However, given that physical networks provide the means of connecting nodes in the higher levels, the study of physical connectivity is an important area of research [34–36]. Furthermore, it is essential to model the impact of large-scale disasters and attacks against the physical infrastructure using the physical-level graph [23]. There are only a few studies that analyse graphs holistically from a multilevel point of view [37–39], but in very specific contexts.

Another important aspect of modelling physical graphs is the *cost* of networks, which is particularly important to consider when designing physical level networks. Moreover, from a network design perspective, it is important to design networks that are *resilient*

yet *less costly*. Unfortunately, these two objectives fundamentally oppose one another. There are no other known studies that provide structural- and cost-based comparisons of geographic graph models applied to graphs with node locations that are constrained to those of actual physical graphs.

Algorithms and mechanisms are necessary to defend networks and to make them resilient against challenges [2]. The design and optimisation of cost-efficient networks that are resilient against challenges and attacks has been studied by many researchers over the past few decades [40–45], but the resilient network design problem is NP-hard [46, 47]. Moreover, networks cannot have unlimited resilience due to cost constraints. Therefore, topological design and optimisation requires developing intelligent algorithms so that a designer can select optimum parameters to achieve resilience in a cost efficient manner.

Communication networks have evolved tremendously over the past several decades, offering a multitude of services while becoming an essential critical infrastructure in our daily lives. While this evolution is still progressing, user expectations from these networks are increasing in terms of performance and dependability. Understanding network behaviour under perturbations can improve today's networks performance, as well as lead to a more resilient and survivable Future Internet. We cannot thoroughly study the effects of challenges in live networks without impacting users. Testbeds are useful, but do not provide the scope and scale necessary to understand the resilience of large, complex networks, although progress is being made in this direction [48, 49]. Simulations arguably provide the best compromise between tractability and realism to study challenges, however this is nontrivial [50].

1.1 Problem Statement

It is essential to understand what features of a challenge contributes to the worst level of degradation in the services delivered by the communication networks. The impact of challenges on network services can be alleviated by mechanisms that increase network redundancy, diversity, and connectivity. However, hardening the networks to withstand perturbations comes with added costs. We would like to find tradeoffs between increased resiliency and added cost for realistic network development. Moreover, it is nontrivial to evaluate network performance and rigorous methods are required to evaluate the performance of networks when faced by challenges. Therefore, our *thesis statement* is as follows:

> Modelling communication network challenges can be useful to understand the impact of such perturbations on networks and can be foundational to improve resilience and tolerance to challenges of existing networks as well as the design of Future Internet architectures.

1.2 Proposed Solution

We propose a graph-theoretical approach to model, design, and evaluate resilient networks. First, we systematically identify a wide spectrum of challenges and categorise them in order to gain a better understanding of challenge impact on service failures.

As a second crucial step, existing graph models are examined. Physical-level topologies provide service to higher levels; however, topological data for physical-level networks is not readily available to study networks holistically. We obtain several networks' physical-level topology data based on a third-party map and subsequently analyse structural characteristics of physical- and logical-level networks. We propose important properties

of graph generators for modelling physical-level topologies. Furthermore, we develop a realistic multilevel graph model to analyse critical infrastructures and their interdependencies.

While network resiliency can be improved by simply adding nodes and links to an existing graph, such additions come with an increased cost. Therefore, we develop heuristic algorithms that generate cost-constrained resilient graphs. We use two graph metrics – algebraic connectivity and path diversity – to generate cost-efficient resilient networks. The heuristic algorithms developed are computationally less expensive than an exhaustive algorithm.

Moreover, we develop a simulation-based framework to evaluate the network performance. We build challenge models that simulates random failures, targeted attacks, and correlated area-based failures using the ns-3 network simulator. Finally, we conduct experiments on the GpENI Future Internet testbed to evaluate the graph algorithms we developed.

1.3 Contributions

The main contributions of this dissertation are as follows:

1. Contribute to the identification of resilience disciplines and their interrelationships. This enables one to recognise resilience disciplines easily.

2. Comprehensive identification of known and potential challenges. This enables one to see the wide spectrum of what can go wrong with communication networks.

3. Categorisation of challenges. This enables one to better understand the challenge relationships and the spectrum of threat models.

4. Spectral analysis of different size and order networks for comparison of their structural properties. The normalised Laplacian spectra of critical infrastructure networks are studied to analyse their structural properties.

5. Development of a graph-theoretical framework to evaluate the performance of multilevel and multiprovider graphs. This framework analyses multilevel graphs instead of single level analysis that obscures realistic analysis.

6. Development of two heuristic algorithms that improve connectivity of graphs in a cost-efficient manner. The algorithms aim to improve two graph measurements: algebraic connectivity and path diversity. The optimisation algorithms we have developed are computationally less costly than an exhaustive optimisation.

7. A methodology to analyse network performability when faced by perturbations. We model challenges such as targeted attacks, random failures, and correlated area-based challenges using the ns-3 network simulator.

8. Evaluation of the graph optimisation algorithm in a Future Internet testbed. The GpENI testbed is used to evaluate the heuristic algorithm outcome.

9. Contributions to the ns-3 open source project by providing an example code. The code written in C++ has been available since the standard release of ns-3.10. The code builds topologies based on user-provided adjacency matrix and node coördinates.

10. Construction of physical-level topologies of four service provider networks (AT&T, Level 3, Sprint, TeliaSonera), a research network (Internet2), a hypothetical fibre network (CORONET), and the US interstate highway topology. These maps are publicly available on the KU-TopView [51] webpage for future use by the research community.

1.4 Relevant Publications

The research presented in this dissertation has resulted in a number of publications, including the following.

Journal Articles

7. **Egemen K. Çetinkaya**, Mohammed J.F. Alenazi, Andrew M. Peck, Justin P. Rohrer, and James P.G. Sterbenz, "Multilevel Resilience Analysis of Transportation and Communication Networks," *Telecommunication Systems.* (accepted in July 2013)

6. Jacek Rak, Mario Pickavet, Kishor S. Trivedi, Javier Alonso Lopez, Arie Koster, James P.G. Sterbenz, **Egemen K. Çetinkaya**, Teresa Gomes, Matthias Gunkel, Krzysztof Walkowiak, and Dimitri Staessens, "Future Research Directions in Design of Reliable Communication Systems," *Telecommunication Systems.* (accepted in July 2013)

5. James P.G. Sterbenz, Deep Medhi, Byrav Ramamurthy, Caterina Scoglio, Justin P. Rohrer, **Egemen K. Çetinkaya**, Ramkumar Cherukuri, Xuan Liu, Pragatheeswaran Angu, Andy Bavier, and Cort Buffington, "The GpENI Testbed: Network Infrastructure, Implementation Experience, and Experimentation," *Computer Networks.* (accepted in July 2013 with minor modifications)

4. James P.G. Sterbenz, David Hutchison, **Egemen K. Çetinkaya**, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller, and Paul Smith, "Redundancy, Diversity, and Connectivity to Achieve Multilevel Network Resilience, Survivability, and Disruption Tolerance **(Invited Paper)**," *Telecommunication Systems.* (accepted in April 2012)

3. **Egemen K. Çetinkaya**, Dan Broyles, Amit Dandekar, Sripriya Srinivasan, and James P.G. Sterbenz, "Modelling Communication Network Challenges for Future Internet Resilience, Survivability, and Disruption Tolerance: A Simulation-Based Approach," *Telecommunication Systems*, Volume 52, Issue 2, pp. 751 – 766, February 2013.

2. James P.G. Sterbenz, **Egemen K. Çetinkaya**, Mahmood A. Hameed, Abdul Jabbar, Shi Qian, and Justin P. Rohrer, "Evaluation of Network Resilience, Survivability, and Disruption Tolerance: Analysis, Topology Generation, Simulation, and Experimentation **(Invited Paper)**," *Telecommunication Systems*, Volume 52, Issue 2, pp. 705 – 736, February 2013.

1. James P.G. Sterbenz, David Hutchison, **Egemen K. Çetinkaya**, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller, and Paul Smith, "Resilience and Survivability in Communication Networks: Strategies, Principles, and Survey of Disciplines," *Computer Networks*, Volume 54, Issue 8, pp. 1245 – 1265, June 2010.

Peer-Reviewed Conference Proceedings

8. **Egemen K. Çetinkaya**, Mohammed J.F. Alenazi, Yufei Cheng, Andrew M. Peck, and James P.G. Sterbenz, "On the Fitness of Geographic Graph Generators for Modelling Physical Level Topologies," in *Proceedings of the 5th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, Almaty, September 2013.

7. Mohammed J.F. Alenazi, **Egemen K. Çetinkaya**, and James P.G. Sterbenz, "Network Design and Optimisation Based on Cost and Algebraic Connectivity," in *Proceedings of the 5th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, Almaty, September 2013.

6. **Egemen K. Çetinkaya** and James P.G. Sterbenz, "A Taxonomy of Network Challenges," in *Proceedings of the 9th IEEE/IFIP International Conference on the Design of Reliable Communication Networks (DRCN)*, Budapest, March 2013, pp. 322 – 330.

5. **Egemen K. Çetinkaya**, Andrew M. Peck, and James P.G. Sterbenz, "Flow Robustness of Multilevel Networks," in *Proceedings of the 9th IEEE/IFIP International Conference on the Design of Reliable Communication Networks (DRCN)*, Budapest, March 2013, pp. 274 – 281.

4. **Egemen K. Çetinkaya**, Mohammed J.F. Alenazi, Justin P. Rohrer, and James P.G. Sterbenz, "Topology Connectivity Analysis of Internet Infrastructure Using Graph Spectra," in *Proceedings of the 4th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, St. Petersburg, October 2012, pp. 752 – 758.

3. Justin P. Rohrer, **Egemen K. Çetinkaya**, and James P.G. Sterbenz, "Progress and Challenges in Large-Scale Future Internet Experimentation using the GpENI Programmable Testbed," in *Proceedings of the 6th ACM International Conference on Future Internet Technologies (CFI)*, Seoul, June 2011, pp. 46 – 49.

2. James P.G. Sterbenz, **Egemen K. Çetinkaya**, Mahmood A. Hameed, Abdul Jabbar, and Justin P. Rohrer, "Modelling and Analysis of Network Resilience **(Invited Paper)**," in *Proceedings of the 3rd IEEE International Conference on Communication Systems and Networks (COMSNETS)*, Bangalore, January 2011, pp. 1 – 10.

1. **Egemen K. Çetinkaya**, Dan Broyles, Amit Dandekar, Sripriya Srinivasan, and James P.G. Sterbenz, "A Comprehensive Framework to Simulate Network Attacks

and Challenges," in *Proceedings of the 2nd IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, Moscow, October 2010, pp. 538 – 544.

Technical Reports

2. James P.G. Sterbenz, Justin P. Rohrer, and **Egemen K. Çetinkaya**, "Multilayer Network Resilience Analysis and Experimentation on GENI," The University of Kansas, Lawrence, KS, USA, Technical Report ITTC-FY2011-TR-61349-01, July 2010.

1. Abdul Jabbar, Qian Shi, **Egemen K. Çetinkaya**, and James P.G. Sterbenz, "KU-LocGen: Location and Cost-Constrained Network Topology Generator," The University of Kansas, Lawrence, KS, USA, Technical Report ITTC-FY2009-TR-45030-01, December 2008.

Extended Abstracts

2. **Egemen K. Çetinkaya**, Justin P. Rohrer, and James P.G. Sterbenz, "Resilience of Backbone Provider Networks," in *IEEE INFOCOM Student Workshop*, Orlando, FL, March 2012.

1. Justin P. Rohrer, **Egemen K. Çetinkaya**, and James P.G. Sterbenz, "Resilience Experiments in the GpENI Programmable Future Internet Testbed," in *Proceedings of the 11th Würzburg Workshop on IP: Joint ITG and Euro-NF Workshop (EuroView)*, Würzburg, August 2011, pp. 29 – 30.

1.5 Summary

In this chapter we give the overview and motivation performed for this dissertation. The rest of this dissertation is organised as follows: Chapter 2 presents background, consisting in part of architectural work contributed by this author and his advisor, as well as a number of others in the ResiliNets research group. The second chapter also contains related work, primarily consisting of previous network design and optimisation algorithms as well as resilience evaluation methodologies. Chapter 3 presents a taxonomy of network challenges that is developed based on a comprehensive survey of existing and potential challenges. Modelling of complex networks is presented in Chapter 4. This consists of understanding structural differences between critical infrastructures using graph spectra, development of a multilevel graph model for realistic analysis of complex networks, and analysis of synthetic graph generators to model physical-level topologies. Network algorithms that optimise a given graph by increasing the resilience properties in the least costly manner are presented in Chapter 5. The resilience metrics that are considered are algebraic connectivity and path diversity, while the cost of networks are captured in terms of total link length. Chapter 6 presents simulation methodology to evaluate network resilience and experiments run on the Future Internet testbed to evaluate the algorithm outcome. Conclusions and future work are presented in Chapter 7. Finally, a complete set of plots used to analyse multilevel graphs is presented in Appendix A and a complete set of graph optimisation results is presented in Appendix B.

Page left intentionally blank.

Chapter 2

Background and Related Work

In this chapter we present background and related work required to understand the rest of this dissertation. First, we present resilience disciplines and the ResiliNets strategy in Section 2.1. A brief overview of network challenges is presented in Section 2.2. Existing graph optimisation algorithms are presented in Section 2.3. Finally, related work on network resilience evaluation methods and models are presented in Section 2.4, and we conclude in Section 2.5.

2.1 Resilience Disciplines

The ResiliNets architectural framework [1–3] provides a strategy and set of principles to alleviate the impact of challenges. *Resilience is the ability of the network to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation* [1, 2, 22, 24]. The resilience disciplines can be divided into two categories as shown in Figure 2.1: challenge tolerance and trustworthiness. The challenge tolerance category considers *adverse events or conditions* that result in operationally degraded networks. The trustworthiness category studies the *measurable characteristics* such as dependability (including reliability and availability), security, and performability. Trustworthiness and challenge tolerance are related by robustness and complexity in the

system. Robustness is the ability and measure of networks to remain trustworthy against challenges. Furthermore, mechanisms to improve resilience characteristics of the networks can add complexity to the system design, which must be managed such that it doesn't *decrease* resilience.

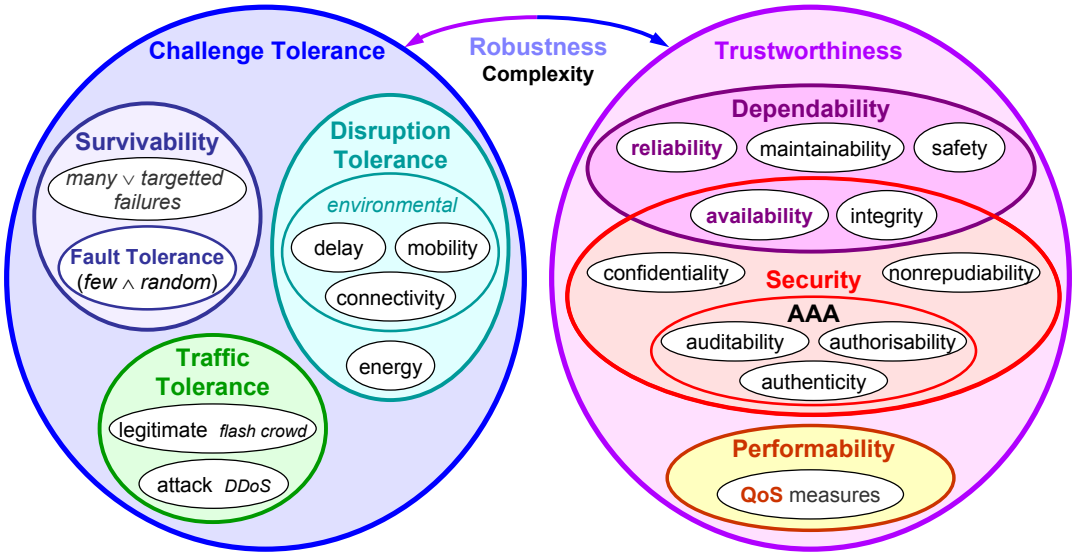

Figure 2.1: Resilience disciplines [1, 2]

Challenge tolerance can be further categorised into survivability, disruption tolerance, and traffic tolerance. Survivability tolerates many correlated or targeted failures [52, 53], whereas fault-tolerance tolerates only a few random failures. Therefore, we consider fault-tolerance to be a subset of the survivability discipline. Traffic tolerance resists challenges against *normal* traffic flows, as well as unusual but legitimate traffic, such as flash crowds in the Internet [17,54–56]. A DDoS attack with a malicious objective attempts to disrupt the services provided by the network. Lastly, disruption tolerance tolerates challenges that are inherent in mobile wireless communication environments. For example delay, mobility, and connectivity are the challenges that need to be dealt with in MANETs (mobile ad hoc networks), mobile cellular networks, DTN (delay tolerant networks), and highly mobile networks. Special purpose wireless networks such as wireless sensor

networks (WSNs) also have to tolerate limited energy constraints.

The challenge tolerance of networks can be increased via the ResiliNets strategy [1,2,24], formalised as D^2R^2+DR, shown in Figure 2.2. Real-time D^2R^2 mechanisms include defence, detection, remediation, and recovery. Long-term DR mechanisms include diagnosis and refinement. While the short-term steps in the ResiliNets strategy provide control aspects of the networks to bring the service levels to their original operating conditions in real-time, long-term steps primarily involve improving service levels as the network evolves. The first step for preserving the resilience of a network involves *defensive* measures. Defence mechanisms can be passive or active. Passive defence primarily involves structural improvement of the network. Two such mechanisms are placing redundant components within the network in order to achieve fault-tolerance and increasing the geographic and mechanism diversity of the network to mask multiple failures for survivability [3]. An example of an active defence includes firewalls that filter anomalous traffic. Next, *detection* is required to discover if the defensive measures have been penetrated [57–59]. After detection of abnormal conditions, the effects of the adverse event or condition should be *remediated* to provide the best possible level of service constrained by available resources. For example, after a power blackout, using a limited number of portable power generators only in selected base stations in a cellular telephony network can provide limited service. *Recovery* involves bringing the operations to the original and normal state [2,21,60] including redeploying destroyed infrastructure. The long-term DR outer loop involves diagnosis as a first step. Diagnosis involves localisation of faults and root-cause analysis [61]. Once the faults are identified by root-cause analysis, the network can be *refined* to improve defence, detection, remediation, and recovery (D^2R^2) in the future for a given and predicted challenges.

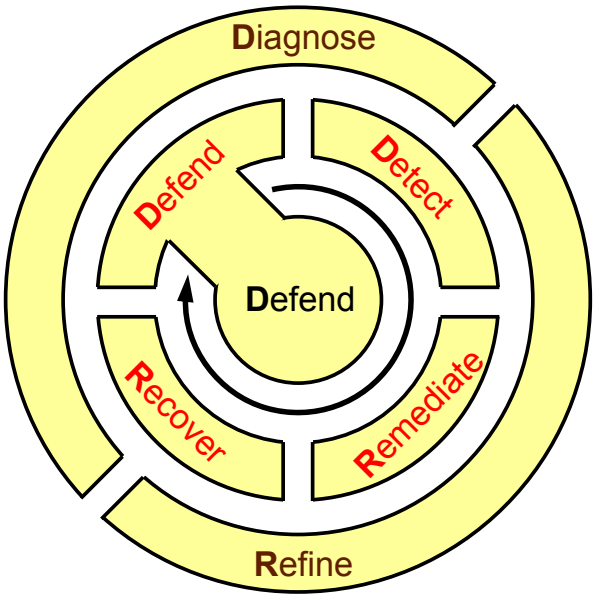

Figure 2.2: ResiliNets strategy [1, 2]

2.2 Communication Network Challenges

In this section, we present major challenge groupings and details will be explained in Chapter 3. The challenges can be broadly listed as follows:

1. **Large-scale disasters:** Critical network infrastructure can be impacted by *large-scale natural disasters* (e.g. hurricanes, earthquakes, tsunamis). The observed service failures are geographically correlated. Furthermore, large-scale disasters can be caused by humans, (e.g. EMP–electromagnetic pulse weapons and power blackouts).

2. **Socio-political and economic challenges:** Social, political, and economic challenges caused by deliberate human actions can threaten resilient communication. Canonical examples include collateral damage to communication networks due to terrorism, nationwide Internet outage due to political decisions, and peering disputes to increase market share.

3. **Dependent failures:** *Service failure* at a lower layer is a challenge to higher layers. For example, if the service provided by a routing layer (discovery of the end-to-end paths) fails to converge, the transport layer sees it as a challenge to its ability to transfer data end-to-end. Moreover, infrastructure dependencies can result in failure of services delivered.

4. **Human errors:** Non-malicious human action such as BGP misconfiguration is a challenge to networks. Moreover, incompetence of operational personnel or designers can result in catastrophic failures.

5. **Malicious attacks:** Deliberate attempts to disrupt service, such as targeted hardware and software attacks, are challenges to networks. Furthermore, damage may be worse if the attack targets Internet control protocols, since the impact can be global.

6. **Unusual traffic:** Unusual but legitimate traffic, such as a flash crowd on the Internet, is a challenge to communication. Moreover, this type of challenge may vary depending upon the specific network. For example, the telephone network is designed to handle the load on Mother's Day but not the load during a catastrophe such as the 9/11 terrorist attacks.

7. **Environmental challenges:** Communication is challenged by phenomena that is inherent in the communication environment. Examples include mobility of nodes in an ad-hoc network, weakly connected channels, and unpredictably long delays.

2.3 Graph Algorithms to Improve Network Resilience

Network design is a NP-hard problem [46, 47] that has been studied in the past decades by many network researchers [41–45,62–66]. The design process includes constructing the

network from the ground up including placement of nodes [44, 45] and providing connectivity among nodes to enable services. The optimisation process includes improvement of the network for one or multiple objectives. Network optimisation can be accomplished by means of rewiring while keeping the number of edges constant [67] or by means of adding new links to improve the connectivity of graphs [68]. Moreover, the design process is different for backbone and access networks, since the topological structure of these networks fundamentally differ [44, 45, 62].

Network design and optimisation objectives are cost, capacity, reliability, and performance [43–45]. Network cost is incurred by the number of nodes required, capacity of nodes required, and number of links. Topological connectivity is another objective that can be measured by means of many graph metrics such as average degree, betweenness, closeness, and graph diversity [41,42,67,69–71]. In this work, we measure the connectivity of a graph in terms of algebraic connectivity [72] and path diversity metrics [71, 73].

2.3.1 Optimisation Based on Algebraic Connectivity

Algebraic connectivity $a(G)$ is defined as the second smallest eigenvalue of the Laplacian matrix [72]. The Laplacian matrix of G is: $L(G) = D(G) - A(G)$ where $D(G)$ is the diagonal matrix of node degrees, $d_{ii} = \deg(v_i)$, and $A(G)$ is the symmetric adjacency matrix with no self-loops. The algebraic connectivity of a complete graph (i.e. full mesh) is n where n is the number of nodes, and it is 0 for a disconnected graph with more than one component.

Topology design using algebraic connectivity has been studied by several researchers [67, 68, 74]. It has been shown that algebraic connectivity is more informative and accurate than average node degree when characterising network resilience [74]. Moreover, we have shown algebraic connectivity [75, 76] and diversity [71] are predictive of flow

robustness of graphs. Three synthetically generated topologies (i.e. Watts-Strogatz *small-world*, Gilbert *random*, Barabási-Albert *scale-free*) have been optimised using edge rewiring in which the objective is to increase the algebraic connectivity [67]. It was shown that algebraic connectivity increases the most if edges are rewired between weakly connected nodes. Another study optimised synthetically generated Erdős-Rényi random and Barabási-Albert graphs in terms of adding links to the existing topology [68]. It was concluded that adding links between a low degree node and a random node is computationally less expensive than an exhaustive search.

2.3.2 Optimisation Based on Path Diversity

Algorithms and mechanisms are necessary to defend networks and to make them resilient against challenges [2]. One such mechanism is *diversity* and it has been the subject of many published works in the field of network resilience. Diversity is used to enhance bandwidth, delay, and loss rate of media streaming applications [77]. Path diversity is used in the optical domain to route around failed nodes or to split traffic for a better utilisation of network resources [78]. Diverse routing is necessary for multihoming to improve the service delivery of provider networks [79, 80]. While the path diversity of a graph is essential for survivable design, it can be improved by addition of links in a given graph using an optimisation algorithm.

A path between a source s and a destination d is the set of nodes and links that form a loop-free connection. Diverse paths between node pairs strengthen the ability of a network to withstand attacks and correlated failures. If the alternative paths have no common node or link they are disjoint, and if there are common network nodes or links, they are partially disjoint. Path diversity has been studied from a topological perspective [81–83], as well as in terms of multipath routing [78, 80, 84–90], and multipath transport [71, 91]. Further, multipath routing has been studied to improve the QoS (Quality

of Service) of networks [86, 87], the resilience of interdomain routing [84, 88], and the survivability in optical networks [78, 89, 90]. Moreover, finding disjoint paths between a node pair is considered to be a NP-complete problem [92, 93]. Next, we explain path diversity and path diversity of a graph [71, 73].

Path Diversity Definition

Given a shortest path and an alternative path between two nodes in a graph, the path diversity of the alternative is defined as the ratio of the number of disjoint elements (nodes *and* links) between the shortest path and alternative path to the number of elements in the shortest path. Given a (source s, destination d) node pair, a path P between them is a set containing all links L and all intermediate nodes N [71],

$$P = L \cup N \tag{2.1}$$

and the *length* of this path $|P|$ is the combined total number of elements in L and N. Let the shortest path between a given (s, d) pair be P_0. Then, for any other path P_k between the same source and destination, the definition of the diversity function [71, 84] $D(P_k)$ with respect to P_0 as:

$$D(P_k) = 1 - \frac{|P_k \cap P_0|}{|P_0|} \tag{2.2}$$

The path diversity has a value of 1 if P_k and P_0 are completely disjoint and a value of 0 if P_k and P_0 are identical. This measure captures the diversity with respect to both nodes and links on alternative paths [71]. As an example, let us find the path diversity of the paths between node 0 and 2 in the graph shown in Figure 2.3.

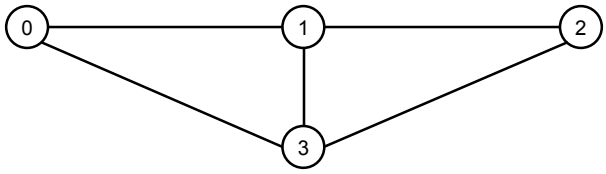

Figure 2.3: Path definition example

There are four possible paths namely: $P_0 = [0, 1, 2]$, $P_1 = [0, 3, 2]$, $P_2 = [0, 1, 3, 2]$, and $P_3 = [0, 3, 1, 2]$. The shortest path in this graph is P_0 so its path diversity is zero. However, to calculate the path diversity for P_1, we first convert the paths to the path elements sets. For P_0, the path elements set is $\{(0, 1), 1, (1, 2)\}$. The tuples $(0, 1)$ and $(1, 2)$ represent the links while the element 1 represents the node 1 in the path. We do not include the source and destination nodes in order to have a path diversity of 1 when the two paths are fully disjoint. Using the same method, the path elements set for P_1 is $\{(0, 3), 3, (3, 2)\}$. Using Equation 2.2, $D(P_1) = 1 - \frac{0}{3} = 1$. To find the path diversity for P_2, the path elements set for P_2 is $\{(0, 1), 1, (1, 3), 3, (3, 2)\}$. Again, using Equation 2.2, $D(P_2) = 1 - \frac{2}{3} = \frac{1}{3}$. Finally the path diversity of P_3, using the same procedure, is $D(P_3) = 1 - \frac{2}{3} = \frac{1}{3}$. We note that converting a path to its elements method assumes directed graphs. In this work, since we study the path diversity of undirected graphs, we modify the method to work with undirected graphs. Thus, to construct a path to element set, we start with the same method for the directed graph. Then, for each link (a, b) in the resulting set, we add (b, a) to the set.

Capturing Path Diversity of a Graph

TGD (total graph diversity) is the average of the EPD (effective path diversity) values of all node pairs in a given graph [71] and TGD measures the path diversity of a graph as a single value. EPD is the normalised sum of path diversities for a selected set of paths connecting a node pair (s, d). First, we find the k diverse paths using the algorithm

presented in Section 5.2.1. Then, we remove zero diversity paths from the list of returned paths because they do not add any additional diversity. The returned diverse path is denoted as $P_{s,d} = \{P_1, P_2, \ ... \ , P_m\}$, where $m \leq k$, since-zero diversity paths are removed from the set. To calculate EPD, we use the exponential function:

$$\text{EPD} = 1 - e^{-\lambda k_{sd}} \tag{2.3}$$

where k_{sd} is the sum of all non-zero diversity paths defined as:

$$k_{sd} = \sum_{i=1}^{m} D(P_i) \tag{2.4}$$

where $D(P_i)$ is the non-zero path diversity of the i-th path with respect to the P_0. In Equation 2.3, λ is an experimentally determined constant that scales the impact of k_{sd} based on the utility of this added diversity [71]. For a given pair of nodes, the range of EPD is between $[0, 1)$ where 0 means that there is no diversity in between the two nodes as there are no alternative paths connecting the pair. When the EPD approaches 1, it means that it has a high path diversity [71].

2.4 Models, Simulation, and Experimentation

Performance evaluation of networks is a vast field, involving many disciplines such as computer science, engineering, mathematics, and physics. We limit our discussion on this topic to performance evaluation of network resilience via analytical models, simulation models, and experimentation testbeds.

2.4.1 Analytical Models

In this section, we briefly provide the mathematical foundations, quantitative methods, and evaluation frameworks that analyse network resilience.

Mathematical Foundations

Reliability theory has a rich history [94] and is defined as the probability of being in the up state during a time interval under specified conditions [95]. It is important to mention the important reliability metrics first. MTTF (mean time to failure) is the average time that a component remains operational. MTTR (mean time to repair) is the average time to repair a component. MTBF (mean time between failures) is the average time between the down states of a component. These reliability metrics can be formalised as:

$$\text{MTBF} = \text{MTTF} + \text{MTTR} \tag{2.5}$$

Given the failure rate λ, MTTF can be denoted as [96]:

$$\text{MTTF} = \frac{1}{\lambda} \tag{2.6}$$

Reliability is the probability of a component being in an up state in a given time interval under specified operating conditions and is denoted as [95, 96]:

$$R(t) = e^{-\lambda t} \tag{2.7}$$

Then, unreliability is:

$$Q(t) = 1 - R(t) = 1 - e^{-\lambda t} \tag{2.8}$$

Average repair time is a function of repair rate μ:

$$\text{MTTR} = \frac{1}{\mu} \tag{2.9}$$

Availability is the probability of a component being in an up state in the future and is denoted:

$$A = \frac{\text{uptime}}{\text{uptime} + \text{downtime}} = \frac{\text{MTTF}}{\text{MTTF} + \text{MTTR}} \tag{2.10}$$

Thus, unavailability is:

$$U = \bar{A} = 1 - A = \frac{\text{MTTR}}{\text{MTTF} + \text{MTTR}} \tag{2.11}$$

Reliability and availability concepts can be easily misunderstood. While reliability refers to probability of failure-free operation during an *interval*, availability refers to probability of failure-free operation at an *instant of time* [97]. To further clarify the concepts, consider the following two network services: transactional service in which a small request is followed by a large response and connection-oriented service in which an end-to-end connection is established before information can flow. For transactional service availability is important since a low repair time and a relatively high operational time is required to complete the transaction; however, individual transactions can still be completed for a system with a high failure rate. On the other hand, for connection-oriented service, the systems needs to stay up, thus requires a low failure rate, but repair time can be

relatively high. Thus, it is the designer's choice to carefully consider either reliability or availability as more important for a given service.

Dependability disciplines (cf. Figure 2.1) such as availability and reliability provide service measurements of components and *subsystems* in terms of binary states (e.g. up and down states). However, a *degradable system* can provide an *acceptable* level of service in the presence of challenges and faults. This is particularly important for networks that are complex systems of systems, in which subsystems fail but the overall network continues to operate. The performability [98–101] of a system S over a period T is defined as:

$$\text{Perf}(B) = P[Y \in B] \tag{2.12}$$

in which Y is the performability variable and can take the values from the the measurable accomplishment set B.

Quantitative Methods

Reliability modelling has been long studied and a variety of models have been proposed to study reliability characteristics quantitatively [50, 96, 102, 103]. Modelling requires metrics to evaluate outcomes [102] and modify the model accordingly. The quantitative models can be categorised as combinatorial and non-combinatorial [50, 96, 104]. Combinatorial models include RBD (reliability block diagrams), FT (fault trees), and attack trees [50]. These methods can give first-cut results of the model outcome but they fall short of representing complex interactions. For example, dependent failure analysis using fault-trees can be intractable. Non-combinatorial reliability models include Markov processes [50,96] and petri nets [105,106]. Non-combinatorial methods can provide solutions to failure models that have dependent events; however, for complex systems such as the

Internet, this can fall short of a complete failure space analysis, in which case simulations can be useful abstraction.

Evaluation Frameworks

There have been several frameworks to evaluate resilience disciplines [2, 107–109]. The seminal IFIP 10.4 work about dependability is the basis for explaining the faults and their taxonomies [13]. Dependability is the discipline that encompasses reliability, availability, integrity, maintainability, and safety [13]. Dependability, security, and QoS together form the resilience trustworthiness disciplines [2].

Resilience of a network can be evaluated using service and operational metrics [110–112]. In this framework, resilience is formalised as transitions of the network state in a two-dimensional state space quantifying network operational state and network service parameters as shown in Figure 2.4. The resilience is measured using the area under the trajectory $1 - \mathbb{R}$ from the initial state to a challenged state $S_0 \rightarrow S_c$.

Another major discipline is survivability [52,53,113,114] and it is a multi-layer function [3, 4, 39, 115, 116]. A survivability specification is defined with six tuples {S, E, D, V, T, P} [114, 117]. Each tuple designation is as follows:

- S: set of acceptable level of service specification
- E: set of relative service values in varying operating conditions
- D: set of environmental conditions a service encounters
- V: set of user-perceived service values
- T: set of valid transitions between acceptable levels of service
- P: set of probabilistic requirements on the operation of a service

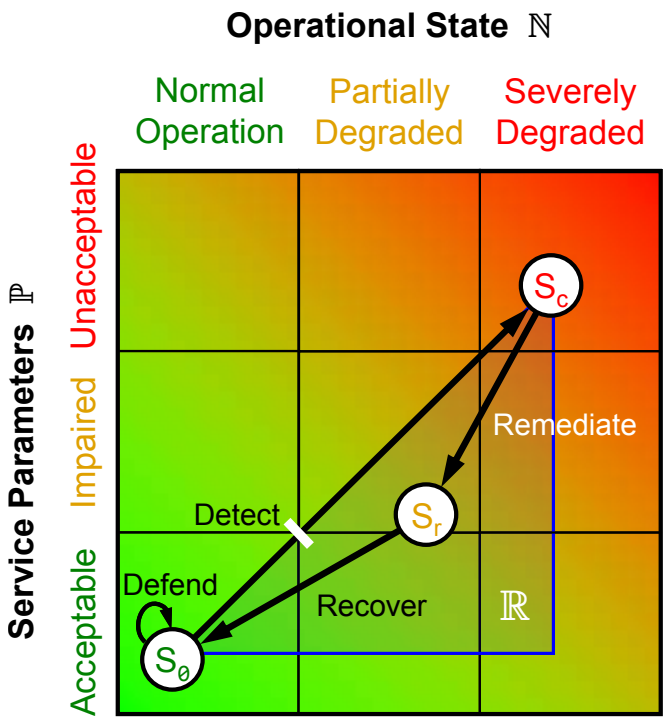

Figure 2.4: Resilience evaluation in two-dimensional state space [3]

The ATIS survivability model provides a survivability function composed of a network's Unservability (U), Duration (D), and Extent (E) during a service outage [4,116]. A given service outage's severity is categorised as catastrophic, major, or minor, depending on the UDE triplet combination as shown in Figure 2.5.

Temporal characteristics of a service failure are given in the ATIS survivability model [4] is shown in Figure 2.6 overlaid with the ResiliNets D^2R^2 strategy (cf. Figure 2.2). The performability of a specific network service, $P(t)$, without any adverse event or condition is at 1. After a challenge, the performability drops to P_a. Remediation takes place for a duration of t_r that is greater than challenge duration $t_r \geq t_c$. Finally the network services are recovered to the original normal state at the end of duration t_R. It should be noted that a challenge duration can be very small, such as duration of a lightning strike.

Another dependability evaluation model was presented by the ANSA framework [6]. In

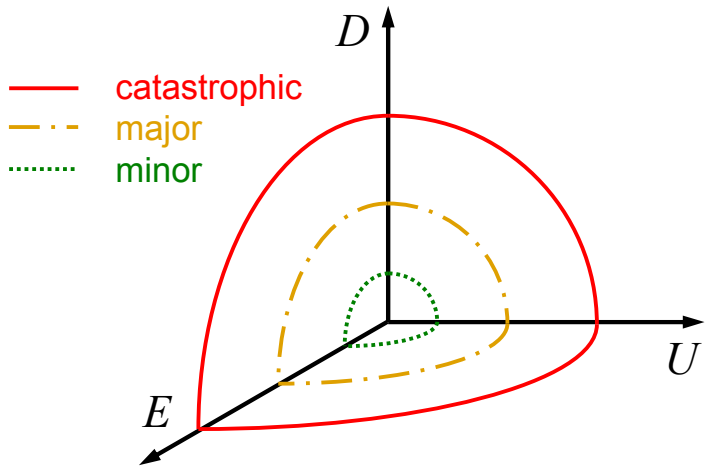

Figure 2.5: Severity of a service outage (adapted from [4])

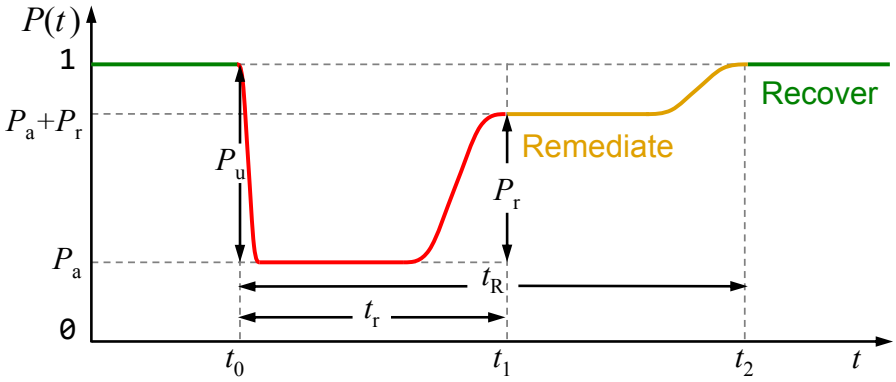

Figure 2.6: Temporal characteristics of a challenge [5] (based on [4])

this model, failures are modelled in the *value* and *time* domain. When the *expected* service from a component in the system deviates, a failure occurs. In other words, the unexpected occurrence of a service o_u results in failures. Three cases of expected values and occurrence relationship are shown in Figure 2.7. In the first case if the service occurrence o_1 deviates from a particular expected value, e_1, the service failure occurs. The expected value (e_2) can be within a range ($v_2 - v_2'$) within a time ($t_2 - t_2'$), and the expected value (e_3) can be time varying.

A multidimensional survivability model has been developed using the Continuous Time

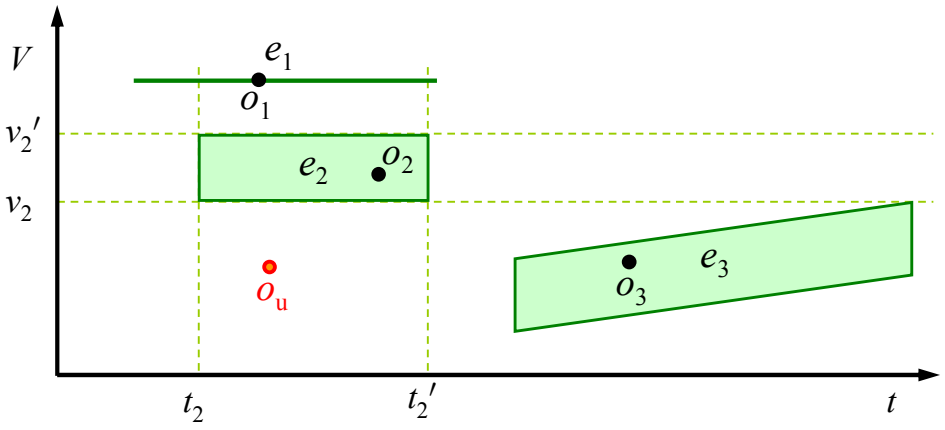

Figure 2.7: Time and value of an event and its occurrence (based on [6])

Markov Chain (CTMC) [7, 118, 119]. State transitions of a system with n components is shown in Figure 2.8. In this model, performability of the system is depicted horizontally with the arrival rate λ and service rate μ. State transitions to model system availability is depicted vertically with failure rate γ and repair rate τ.

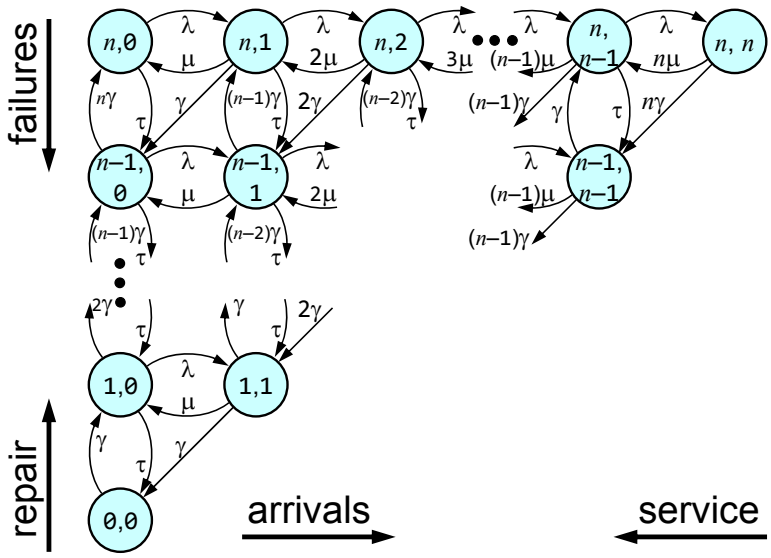

Figure 2.8: Two dimensional Markov model (based on [7])

2.4.2 Simulation Models

Simulations provide the tools to analyse networks in complex scenarios. On the other hand, modelling and simulating network performance under challenge conditions is non-trivial [50, 120, 121]. There have been several studies that analyse different aspects of networks under challenges. Next, we briefly present the past simulation models that study graph-based, large-scale correlated failures, attacks, and wireless medium challenges.

Topological Simulations

Random topologies faced by random node or link failures have been surveyed [69, 122–124]. However, given the *complexity* of the Internet, random topologies do not reflect realistic analysis. Network topologies faced by targeted attacks by the degree of connectivity are shown to have local effect [125], since higher degree nodes (i.e. access PoPs) reside on the edge of the network. Statistical properties of logical topologies under degree-based attacks for static and dynamic evolving cases have been investigated [126, 127]. It was shown that a massive attack on Internet connectivity as a whole is not feasible [127]. Vertex and edge attacks against the wired topologies have been studied [128–130]. The attacks are based on static and dynamic topologies, in which graph metrics, such as degree of connectivity and betweenness are recalculated dynamically after the attacks.

Large-scale Geographic Simulations

Recently, there has been an interest in characterising the behaviour of *physical topologies* under geographically correlated failures. Logical topologies can be useful to study random failures and their impacts, as well as to analyse the performability of the networks under attacks [20]. However, physical topologies are *necessary* to study the impact of geographically correlated failures of the physical infrastructure. While some studies model

geographic challenges dynamically (i.e. temporally and spatially evolving), others model these as static only.

There are several studies that utilised MATLAB to model static geographic area-based challenges. Algorithms that find a worst-case line segment and circular cut are presented and numerical results are shown using MATLAB [131–135]. 2-terminal and all-terminal methods of identifying worst large geographic challenges have been evaluated in MATLAB for wired and wireless topologies [136, 137]. Survivability of large-scale regional failures has been modelled, considering the performance of path restoration after a failure [138]. Both static and dynamic area-based challenges are modelled to study required restoration capacity and time [139]. Unfortunately, the simulation platform has not been explicitly stated in these studies [138,139]. Wired and wireless telecommunication system performance (i.e. call blocking) can be analysed using the Network-Simulation Modeling and Analysis Research Tool (N-SMART) [140, 141]. The failure scenarios consider the geographic locations of network resources and static challenge models.

Attack Simulations

Distributed denial-of-service (DDoS) attacks have been simulated using the ns-2 simulator and performance of legitimate user bandwidth is analysed for different queueing algorithms [142], but DDoS attacks generally are not targeted against infrastructure such as routers. An agent-based simulation model is presented for attack scenarios [143]. Another agent-based simulation framework based on OMNeT++ has been developed [144]. Worm propagation models that impact end-hosts at large-scale (50,000 hosts) have been developed using the Georgia Tech Network Simulator (GTNetS) [145]. A *policy-based* resilience evaluation simulator based on SSFNet [146] has been implemented [27].

Wireless Simulations

Simulation models to analyse network resilience in the wireless medium have not been extensively studied. A toolkit has been previously implemented in ns-2 for simulating obstacles; however, it lacks jammers and impairments [147,148]. Jammer and impairment models in ns-3 have been recently implemented [20, 23, 149–151].

2.4.3 Challenge Experimentation

Analytical and simulation models help our understanding of the network challenges to build better network architectures. However, to realistically evaluate communication network challenges at scale, testbeds are required. Programmable networks provide the necessary infrastructure to evaluate new protocols and architectures [48, 152–155]. Two recent initiatives of programmable networks are GENI (Global Environment for Network Innovations) [156] in the US and FIRE (Future Internet Research and Experimentation) [157] in the EU (European Union). DETER is a testbed that is medium-scale for network security experiments [158, 159]. A recent small-scale testbed effort is provided jointly by the SecSI and LHS labs to provide security experiments [160]. GpENI (Great Plains Environment for Network Innovation) testbed, which is a part of the GENI program, was deployed to provide programmable testbed infrastructure at a global scale [48], with 40 sites in 20 nations. Experiments to evaluate network resilience is proposed [49].

2.5 Summary

Based on this literature survey, we can conclude that network design and optimisation have been active areas of research. The objectives of these studies vary in terms of graph metric that is being optimised. Moreover, several techniques to evaluate the performance

of networks are presented. But it is also clear that much remains to be done, which is the focus of this dissertation.

Page left intentionally blank.

Chapter 3

Challenge Model and Taxonomy

Identifying, understanding, and analysing challenges to communication networks is essential to increase the defence, detection, and remediation capabilities of existing networks and the development of Future Internet architectures and protocols. Starting with the right threat model is essential for the cost-efficient and resilient system design [161]. Therefore, a better understanding of the challenges and possible impacts on networks and services is essential for improving existing networks and designing the Future Internet.

This work has resulted in a publication in which we developed the challenge taxonomy [161]. The remaining sections of this chapter are organised as follows: In Section 3.1 comprehensive examples of challenges are presented. Next, based on the identified challenges, we present a taxonomy of challenges that can help assess resilient designs and mechanisms in Section 3.2. Finally, we conclude the chapter in Section 3.3.

3.1 Challenge Examples and Impacts

In this section, we present a comprehensive survey of challenges that have affected communication networks, as well as potential scenarios that might impact networks in the future. We organise these challenges within each of the challenge tolerance disciplines:

survivability, traffic tolerance, and disruption tolerance. We do not claim that this is a full list of challenge examples; however, we strive to be as comprehensive as possible. In some well-studied sub-disciplines, such as DDoS (Distributed Denial of Service) attacks, we refer the reader to existing comprehensive surveys. The examples demonstrate *what can go wrong with communication networks*. The identification of these events primarily relies on scholarly articles, blogs from Internet monitoring companies such as Renesys [162], Arbor Networks [163], BGPMon [164], and reports from organisations such as RIPE [165], ICANN [166], NANOG [167], and ENISA [168]. In rare instances we use news websites to corroborate the examples.

3.1.1 Survivability

Survivability is a resilience discipline that is concerned with correlated failures and attacks [2, 52, 53, 113]. Survivability is defined as: *the capability of a system to fulfill its mission, in a timely manner, in the presence of threats such as attacks or large-scale disasters* [2, 52, 113]. On the other hand, the fault tolerance discipline that is a subset of the survivability discipline is concerned with *few and random* faults in a system and is not adequate to defend against attacks and large-scale correlated failures.

Several network attack mechanisms exist and a comprehensive coverage of security attacks are beyond the scope of this study. However, there are comprehensive surveys that present attack scenarios and examples for a variety of networks [169–179]. Instead, we focus on challenge examples of *large-scale correlated failures* such as: natural disasters, human-made disasters, challenges to the Internet waist, and socio-political and economic events that disrupt communication network services on a large scale. We also note that some of the examples in this section also cover the fault tolerance subset aspects (e.g. collision of satellites). We begin by discussing disasters, which can be broadly categorised as *natural* or *human-made*.

Natural Disasters

Canonical examples of natural disasters include earthquakes, tsunamis, hurricanes, ice storms, and wildfires. Depending on its origin, a natural disaster can be categorised as terrestrial, meteorological, or cosmological. Next, we survey some of the recent well-known *natural* disasters that had significant impacts on communication infrastructure.

An earthquake of Richter magnitude 7.1 hit Taiwan in 2006. The Hengchun earthquake caused disruption to seven out of nine submarine cables in the Luzon Strait that eventually caused the loss of primary BGP (border gateway protocol) peerings [180]. Even though backup AS (autonomous system) paths were utilised, human intervention was needed to optimise the traffic flow in the Asia-Pacific research and education networks. It is suggested that geographically diverse physical topology, QoS-aware BGP, integrated traffic management, and post disaster emergency communications are crucial for continuation of network services [181]. The 2008 Wenchuan earthquake in China caused a complete communication halt within the most-affected areas. The 8.0 magnitude earthquake caused immense infrastructure damage: 3,897 telecom offices, 28,714 mobile cellular towers, 28,765 km of fibre-optic cables, and 142,078 telecom poles were ruined [182]. Power disruptions also impacted operation of the communication infrastructure. The number of telephone calls initiated from the impacted areas was 10 times higher than normal, and the number of calls from the rest of China to the impacted areas was 6 times higher [182]. Recently, the 9.0 magnitude Fukushima earthquake and the resulting tsunami hit Japan in March 2011 [183]. Surprisingly, the overall network impact of the huge earthquake was limited: out of the 6,000 network prefixes only 100 of them were withdrawn from the global routing table [184–186]. On the other hand, 1.5 million telephone lines were disrupted during the peak of the disaster [187]. The limited global impact was reasoned due to the *rich* connectivity of Japan with the rest of the

world.

Tsunamis, which are caused by seismic activity such as earthquakes, landslides, and volcanic eruptions, can also damage network infrastructure. The 2004 Indian Ocean tsunami, which was caused by the 9.0 magnitude Sumatra-Andaman earthquake in the middle of the ocean, impacted all public telecommunication services in the Maldives Island of Sri Lanka [188]. Another major earthquake caused the tsunami that damaged network infrastructure occurred in 2010 on the southern coast of Chile [187].

Hurricanes are another type of natural disaster. The United States was hit hard in 2005 by category 5 hurricanes Katrina, Rita, and Wilma [141, 187, 189]. *Flooding* in the areas hit by hurricane Katrina also impacted communication infrastructure [189]. Three million telephone lines were cut, 2,000 cellular telephone towers were damaged, and major damage to radio broadcast stations was observed in the impacted areas [189]. Internet access in the Gulf Coast region was impaired; however, the impact of Katrina on the Global Internet routing was limited [190]. In September 2008 category 2 hurricane Ike caused large-scale disruptions in multiple US states [191–193]. Hurricane Irene affected tens of networks in each state that it passed through along the east coast of US [194]. The network failures during hurricane Irene were due to power outages and equipment damage, as in the case with the hurricane Gustav and Ike [191, 192, 195, 196]. More recently, hurricane Sandy caused major disruptions in the Northeast US. Datacenters flooded and more than 5% of network prefixes were withdrawn in New York and New Jersey region [197, 198]. Traceroute measurements indicate that international traffic was rerouted from Ashburn, VA PoP (point of presence) for almost 24 hours when the PoP in New York City was not operational [199, 200]. Traceroute [201] and ping [202] probes using ICMP echo request messages indicate that it took up to five days for some ASes to become operational again. Moreover, less than 1% of a provider's customers lack network services even after four months of the hurricane Sandy's initial impact [203].

Thunder and rain storms affect access links, particularly in wireless networks. A study identified that residential customers across several ISPs were *unreachable* during thunder and rain storms [204]. Disruptions to communication caused by rain storms in the Great Plains region of the US are common and the impact of these challenges depends on the rain rate and the geographic area that they cover [205–207].

Derechos are windstorms that are associated with showers and thunderstorms [208]. The Ohio Valley / Mid-Atlantic Derecho of June 2012, which caused deaths and financial losses, also impacted communication networks. In particular, 77 PSAPs (public safety answering points – also known as 9-1-1 call centers) in six states serving 3.6 million users lost some degree of connectivity [209]. The most notable cause of failures was lack of functional backup power systems.

Ice storms, a type of meteorological event, can result in disaster. In 1998 Canada and the northeastern United States were hit by an ice storm that the US government declared a disaster in the impacted areas. Power loss and damage to telecommunication infrastructures (e.g. antennas and towers) were observed; however, overall network performance outside of the directly-affected area was not significantly degraded [210].

Geomagnetic storms can impact communication systems [211–214]. The impact of the geomagnetic storms on telegraph systems has been reported as early as 1847 in England [213]. In 1958, *world-wide* radio fade-outs and disruptions to transatlantic submarine cables were observed in a 10 hour window of a storm's impact [212]. One of the most severe magnetic storms in history resulted in a power blackout for over 9 hours in Québec, Canada in 1989 [213, 214]. The impact of geomagnetic storms on the commercial communication and GPS satellites has also been documented [211].

Human-made Disasters

Humans can be the cause of disasters and the consequences of such events can be catastrophic. Human-made disasters can be the result of simply ignoring an early warning in a system's operation or can be the result of a malicious act such as terrorism. The *target*, *objective*, and *intent* of the human actions can also vary.

An electromagnetic pulse (EMP) attack with a *malicious objective* can severely disrupt the communications infrastructure [215]. In 1962, during the test of such weapons, the disruptions to telecommunication infrastructure were observed 1,000 miles away from the test field [216]. EMP bombs can severely disrupt cellular telephony networks and the power grid [217]. Such attacks can be costly and restoration of telecommunication equipment could take as long as 27 months [218].

Space debris is another challenge to communication networks and it has been observed in several occasions. On 10 February 2009, Iridium 33 satellite collided with the decommissioned Russian communications satellite Cosmos 2251 at an altitude of 790 km over Siberia in low earth orbit [219]. As the rerouting of the traffic was expected to take 3 days and moving a spare satellite as a replacement was planned for 30 days, possible disruptions to Iridium services were announced [220]. Such an example of satellite communication challenge was reported as a risk factor during the annual SEC (securities and exchange commission) filing [221]. Furthermore, the electronic hijacking of satellite communications has been repeatedly documented [222]. We also note that, in addition to the challenges that directly impact the satellite communications, the GPS (Global Positioning System) satellites are critical for the operation of communication networks. They supply location information for location-based routing and precise timing information for many network services including SONET (synchronous optical networking) and mobile telephony. Therefore, a failure in the GPS system can result in failure of our daily

communication systems. Alternative systems to GPS for location and timing information has been recommended to reduce the dependency among communication systems [223].

Fires rarely impact communication infrastructures, but a few significant events involving fires have caused considerable damage [224]. According to a survey that included 27 telecom companies in the US representing 93% of the customer base, 189 fires occurred during the 1988–1992 period [224]. Prior to the survey period, between 1982 and 1987 there were also reports of 6 fires affecting service. Out of the 189 fires during the 1988–1992 period, 35 of them impacted service. The causes of the fires were categorised according to initiation: telco or peripheral equipment, power equipment, personnel (either vendor or carrier staff), natural (e.g. lightning, wind, and flood), outside facility, and building systems. Out of the 189 fires, 63 were initiated by power equipment (e.g. AC/DC, batteries, diesel generators). Among the 41 fires impacting service between 1982 and 1992, 13 were caused by power equipment. 10 out of the 195 fires were caused by natural phenomena of which 7 affected service [224].

A major fire accident occurred in 1988 at the Hinsdale central office of the Illinois Bell company in Hinsdale, IL [12, 225–227]. Some reported that the cause of the fire was lightning [227], while others reported that the cause was a damaged power cable [226]. 35,000 residential and business customers were out of telephone service, 9-1-1 emergency services were disrupted in the impacted areas, 50% of the cellular service in Chicago area was affected, directory assistance service was down, and the impact was seen as far as 50 miles away from the Hinsdale central office [227]. The service was fully restored almost one month after the fire [225]. Two other fires in the New York City area caused long-term disruptions to the public switched telephony system in 1975 and 1987 [226].

On 18 July 2001, a fire erupted after CSX Transportation train derailed in the Howard Street Tunnel in Baltimore, MD [228]; the root cause of the derailment was never identi-

fied [229]. The fire caused damage to fibre-optic cables that ran through the tunnel. To make matters worse, these cables were used by several major ISPs and the single point of failure in this case resulted in disruptions and slowdowns across the country [230]. Recovery of the damaged cables took 36 hours [228, 230]. BGP prefix updates showed burstiness on that day [231]. Both of the Hinsdale central office and Baltimore Tunnel fires show that *diversity* is a crucial property for a resilient network topology [2, 226], that is fault-tolerant redundancy is not sufficient if the redundant components share the same fate.

Cable cuts can be caused by natural phenomena such as earthquakes [180], chafe, corrosion, and submarine landslides [232]. Sharkbites [233] and equipment failures [234] also cause disruption of submarine cables. Cable cuts can also result from man-made hazards to submarine cables, including fish trawling, ship anchors, ocean mining, and sabotage. According to data collected from 1879 to 1980 covering a 101-year period, a total of 1061 submarine cable failures was observed [232]. 75% of the submarine cable failures were attributed to fish trawling, chafe and corrosion [232]. Recently, submarine cables were damaged in three different occasions over a two-day period [235]. Two of the cable breaks occurred near Alexandria, Egypt on 30 January 2008 impacting 13 countries in Africa, the Middle East, and Asia [236–238]. The third submarine cable break happened near Dubai, UAE on 1 February 2008 [239]. Network performance metrics (latency, jitter, throughput, and packet loss) degraded until the cables were fixed between Europe, the US, and India [239]. The submarine cable breaks were restored in about 10 days [239]. The analysis showed that to cope with such incidents *geographic diversity* is needed [235]; however, laying additional cables on alternative routes can be very expensive. As with the case of the Baltimore tunnel, multiple providers' cable frequently share routes through constrained geography such as the Suez Canal [240].

The frequency of fibre-optic cable damages in sub-surface and aerial media is not any

better than the frequency of submarine cable disruptions. Between September 1992 and February 1993 there were a total of 160 fibre-optic cable failures [241]. Among the 160 failures, 33 of them were large outages affecting 30,000 customers in the telecommunications sector. The causes of fibre-optic cable failures include: dig-ups, vehicle, process, power line, rodent, sabotage, fire, firearm, flood, and excavation. Dig-ups (sometimes referred as *backhoe fades*) caused 60% of the total outages [241, 242]. A recent accidental fibre cut in Georgia caused the loss of Internet services in Armenia for five hours in 2011 [243]. Multiple independent fibre cuts were detected in the Telenor backbone network in Norway [244]. Events ranging from a crashing airplane damaging the aerial fibre to vandals shooting fibre-optic cables in an ISP can be the causes of cable cuts [245].

A pandemic is a spread of disease that can impact large populations across the globe. In the case of a *biological warfare*, it originates with a malicious objective. We did not find any examples of pandemics impacting communication networks; however, for continuity of operations during a pandemic, planning, preparedness, response, and recovery actions are often necessary [246]. The potential impact of an influenza pandemic on telecommunications and information technology could be significant [247]. In the worst case scenario, decreases in network maintenance and increases in telecommuter traffic can result in network disruptions. Furthermore, critical operations and maintenance staff may be afraid to report to work to keep the network operational.

Power blackouts can cause network failures due to *interdependencies* among critical infrastructures [29, 248–253]. During August 2003, 50 million people were affected in the Northeast US and parts of Canada [249, 250]. One month later, 55 million people in Italy, France, and Switzerland were impacted in another event [29]. 15 million people in Western Europe were impacted by the November 2006 blackouts [249]. Communication networks in the Latin American countries of Brasil, Paraguay, and Uruguay were impacted by the blackouts in November 2009 [254]. Despite the number of networks

and ASes in the various blackout areas, the characteristics of the network outages show similarities [255]. The average outage duration for large network service providers ranged from 12 to 33 hours in the Northeast US blackout of 2003 [255]. The impact of the blackouts are regional, rather than global in scale [255, 256]. 150 large-scale PSTN outages in the US over an 8 year period due to power related outages have been analysed and 20% of these have been identified as high impacting outages in terms of lost customer minutes [257]. More recently, in late July of 2012, northern India experienced power blackouts over two consecutive days that impacted more than 600 million people, and the impact was observed in other critical infrastructures including the communication networks [258–260].

Human errors result in large-scale network disruptions, which are the cause of 50% of the outages in PSTN [261]. For example, although the Hinsdale central office fire was not human initiated, that fact that the operator initially *ignored* the alarm resulted in the late arrival of the firefighters, which in turn resulted in *severe fire damage* [225, 226]. Locating so many potential points of failure in the same building can be attributed to *human incompetence*. Furthermore, during the event of the Baltimore Tunnel fire, the decision of the network designers to use redundant fibres along the same geographic route resulted in *severe* disruptions. From a security point of view, most threats come from humans [262]; however, designing systems to tolerate human errors is difficult [263], and requires redundancy, diversity, and heterogeneity [3] in addition to traditional security mechanisms.

Finally, software vulnerabilities can be exposed in a variety of scenarios. The root cause of AT&T's SS7 (signalling system 7) outage in 1991 [225] that was one of the many SS7 outages in the 1990s [264] impacting the US PSTN, was determined to be a misplaced `break` statement [265]. There are some other interesting examples that have challenged the Global Internet [266, 267]. A Juniper BGP edge router software bug caused global

Internet outages for two hours in November 2011 [268, 269]. More recently, the GitHub service was disrupted globally due to a software bug in which switches in the GitHub network did not populate the MAC address table properly resulting 18 minutes outage [270].

Challenges to the Internet Waist

The Internet is a collection of ASes (autonomous systems) from a topological point of view. Each of these networks or domains consists of end systems, routers, and links that connect them. The protocols enable end systems to communicate with each other. Openness of the Internet architecture enables applications to be developed independent of the network infrastructure [271–273]. The hourglass metaphor reflects the Internet architecture and it has a *narrow waist* in the middle containing IP (Internet Protocol). Thus, IP is the minimally required element [274]. However, in practice the narrow waist is not so narrow. The Domain Name System (DNS) [275,276] and Border Gateway Protocol (BGP) [277] are both required for practical use of the Internet and thus increase the size of the waist as shown in Figure 3.1. Evolving shapes of the hourglass model are presented in an entertaining article [278].

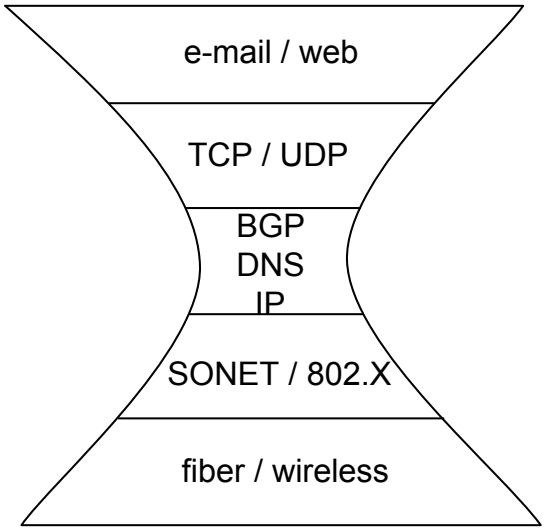

Figure 3.1: Internet hourglass waist model

BGP is the interdomain routing protocol that peers ASes. Each AS is a single administrative domain, in which the network is managed by a single entity. ASes announce IP prefix reachability information to other ASes, enabling the construction of an AS topology. BGP utilises policy-based path-vector routing in which each domain administrator decides how to advertise AS reachability information. The early deployment of the BGP was in 1989 and the current BGP-4 was deployed in 1993 [279]. There are a number of BGP threats, vulnerabilities, security attack objectives, attack mechanisms, and defenses against malicious activity [280–285]. In addition to intentional attacks targeted directly against BGP and accidental misconfiguration of BGP, its operation has deviated from normal due to the increased number of updates resulting from worm activity, power blackouts [250, 256], and power outages after disasters [193, 286].

Worms such as Morris, Code Red, Nimda, Blaster, and Slammer have been the source of the Internet disruptions [170]. BGP operation has been impacted as a result of those worms [287–290]. The analysis of BGP during the highest point of Code Red worm impact indicated that BGP updates increased 10-fold [287]. On the other hand, another analysis concluded that the enormous BGP activity during these worm attacks was an artifact of the monitoring process [288]. Increased BGP activity in one event was due to BGP session resets that occurred at the monitoring RIPE NCC (Réseaux IP Européens Network Coordination Centre) RIS (Routing Information Service) [291]. Each BGP session reset caused entire routing table transfers. Even though BGP was stable, the monitoring process was affected by the worms [288]. The Slammer worm caused an eight-fold increase in the number of BGP update messages and some parts of the Internet were not reachable on that day [289]. Further investigation revealed that the instabilities caused by the frequent update messages were due to a few edge ASes that might have been hit hard by the Slammer worm [289]. In August 2003, the Blaster worm impacted the hosts using the Microsoft Windows operating system [292]. The number of BGP

update messages also increased significantly during the Blaster worm [293].

Prefix hijacking occurs when an AS announces an IP address prefix (i.e. destination) that it does not belong to itself. When a prefix is hijacked, it can cause a blackhole effect in which packets do not reach the intended destination, resulting in a denial of service attack. The hijacker can also perform a man-in-the-middle attack by eavesdropping on the hijacked traffic and forwarding it to the original destination after interception [294–296]. Between 1997 and 2009, there were 15 high-profile prefix hijacking events according to a study compiled from NANOG (North American Network Operators' Group) mailing list archives [297]. On the other hand, between 1997 and 2001, in a span of 1279 days, 138,225 multiple origin AS conflicts were observed in which legitimate causes accounted for multiple origin AS conflicts such as multi-homing [298]. According to a 21-day study, all BGP misconfiguration errors are short-lived (less than a day) and 0.2–1.0% of BGP table entries are affected by the misconfigurations [299]. Next, we present some of the most well-known *hijacking events* that cover events from different continents, countries, service providers, and time zones showing that prefix hijacking has no borders.

Spammers can hijack IP prefixes to send unsolicited commercial e-mail [300]. Spammers can use unused IP address blocks (IP prefixes) to send spam. However, to make tracing difficult, they prefer to announce short IP prefixes (e.g. /8 instead of more specific /24) for short periods of time [300]. In 2003, an IP prefix hijacking case was not resolved for two months since the spam started [301]. However, we did not find any global impacting event to BGP routing caused by the spammers in the literature.

On April 1997, AS 7007 owned by the Florida Internet Exchange and used by MAI Network Services started announcing the entire global routing table to their upstream provider Sprintlink [298,302]. The more specific (i.e. de-aggregation) /24 announcements from an AS 7007 router started at 11:30 UTC and continued for 45 minutes until MAI

disabled the link to the upstream provider Sprintlink. Thus the router manufactured by Bay Networks was unplugged from the network [303, 304]. Despite taking the router offline, advertisements to the Global Internet continued. Later, the issue in which withdrawn messages were not updated with the rest of the global routing tables was found to be due to a Cisco router bug [305].

In December 2004, AS 9121 used by Türk Telekom (TTNet) announced over 106 K bad prefixes to its upstream provider AS 6762 (Telecom Italia) [306]. The peer router in Telecom Italia was configured with a relatively low value for the maximum number of prefixes it could accept from TTNet, however the router configuration was not saved. Eventually, during the 45 minutes of bad prefix announcements 70% of Internet routes were originated from Turkey, causing significant disruptions to the Global Internet [295, 307].

AS 27506 used by Con Edison Communications started announcing routes to its existing and *former* customers through itself in January 2006 [296,308]. If the traffic was delivered to its existing customers then there was no issue for them. However, AS 27506 also announced prefixes that did not belong to it. There were two occasions of prefix hijacking. One was around 05:05 UTC lasting for 17 minutes and the other started at around 08:30 UTC announcing Panix (AS 2033) routes as well [308,309]. Despite the upstream provider Verio's route filtering, the hijacked routes were announced since the filters were based on the outdated public RADb (Router Arbiter Database) registry [309].

In February 2008, we witnessed one of the most publicised BGP hijacking event. In this case, upon receiving a Pakistani government order to block three IP addresses on the YouTube service, Pakistan Telecom (AS 17557) started advertising /24 prefixes with the intention to block those sites [310, 311]. The prefix that was being announced was more specific than the /22 that YouTube actually announces. PCCW Global (AS 3491),

which is the upstream provider to Pakistan Telecom, started to announce the bad prefixes coming from Pakistan Telecom. The incorrect routes to YouTube via Pakistan Telecom poisoned the global routing tables in about two minutes [310,312]. Eventually, YouTube started announcing /24 and more specific /25 prefixes to re-route the traffic. The resulting impact was inaccessible YouTube videos up to 2 hours depending on the vantage point [310, 313]. There were three issues in this event: governmental censorship on the Global Internet [313], incorrect operation of Pakistan Telecom (instead of blocking three IP addresses started the blackholing effect [314]), and reliance of global routing on a transitive trust model [313].

In another hijacking event around the same time as Pakistan Telecom's YouTube hijacking, AboveNet (AS 6461) started announcing 194.9.82.0/24 on March 2008 [315]. The hijacked prefix was owned by Africa Online Kenya (AS 36915). The announcements by AboveNet located in the US were believed to diffuse into the global routing tables faster since they are more *centrally located* in the AS-topology [315].

On November 2008 Companhia de Telecomunicações do Brasil Central (CTBC) from AS 16735 announced the full BGP routing table several times in a 75-minute period [316]. Fortunately, there was no impact on the global routing tables in spite of the leak of 267,947 distinct prefixes [317,318]. The full routing table leak was identified by one of the RIPE RIS monitors that peered with the AS 16735. However, the incorrect prefixes were filtered by the upstream provider of CTBC (Companhia de Telecomunicações do Brasil Central) and the only impact was slow traffic for many hours for CTBC customers [318].

Another BGP hijacking occurred in April 2010 in China [319]. In this event, China Telecommunications Corporation (AS 23724) announced approximately 37,000 prefixes that it did not own. This corresponded to approximately 11% of the total number of prefixes worldwide [320]. The prefixes announced by AS 23724 were further propagated

by the upstream provider China Telecom (AS 4134) [319, 321]. According to a US government report, 15% of the global traffic could have been eavesdropped during the 15 minute propagation [322]. However, further analysis showed that the actual traffic that might have been intercepted by China is estimated to be around 0.015% [323].

In August 2010 an Internet research experiment caused instability in the global routing tables [266]. RIPE NCC was announcing a valid optional transitive BGP attribute to propagate certificate information for secure BGP research. The BGP announcements were made for approximately 30 minutes and caused a significant increase in the number of update and withdrawn messages. At the peak, 4,500 prefixes, which account for 1.4% of the total global routes, were affected [266]. Further examination of the issue yielded a software bug in Cisco routers [324].

In February 2009, Czech ISP SuproNET (AS 47868) was prepending its AS path information to its prefix (94.125.216.0/21) announcements for traffic engineering [267, 325]. In this case, SuproNET's primary peer was the CD-Telematika a.s. (AS 25512) and they were configuring the back-up paths through Sloane Park Property Trust, a.s. (AS 29113) [267, 325]. The routers used by SuproNET were manufactured by MikroTik and during the configuration instead of prepending the *number* of AS-paths, SuproNET prepended its AS number (47868) [267]. In modulo 256, this corresponded to 252 and indeed the announcements during the instability had AS path 47868 252 times [267]. Cisco routers that couldn't handle the long AS-path information started tearing down the BGP sessions, causing global instability. The first issue was that SuproNET operators were prepending incorrect information [267]. The second problem was that unpatched Cisco routers were still vulnerable to buffer allocation for long AS paths [267, 326]. The third issue was that the AS-path information was not filtered by the upstream provider [267]. For two hours, an up to 100-fold increase was observed in the number of update messages [267, 325]. The long AS path prepending issue was seen in other instances as

well [327].

In another BGP prefix hijacking event, Google services were affected [328]. Between 7 and 9 May 2005, one particular prefix (64.233.161.0/24) that Google (AS 15169) owned started to be announced from Cogent (AS 174) [328]. According to Google, the outage was due to DNS misconfiguration. However, the network operations group at Google were surprised to learn that BGP prefix hijacking is a possible means of attack [328]. Redirection from www.google.com to search.msn.com was suggestive of malicious activity [328].

In February 2012 incorrect BGP filtering caused an approximately 30 minute of outage impacting many Australians [329]. In this case, Telstra (AS 1221) provides transit service for Dodo (AS 38285) and Dodo re-announced all Internet routes it learned from another ISP to Telstra. Since Telstra did not filter the Dodo's announcements, Telstra preferred Dodo as the best way to access the Internet [329]. In this cascading failure scenario, approximately 1400 IPv4 network prefixes were impacted, whereas none of the IPv6 prefixes were impacted.

DNS is the other protocol that *bloats* the narrow waist of the Global Internet. DNS associates the name of a destination (URL or e-mail) with an IP address. It has become an essential service, particularly for web and e-mail applications. It is based on an hierarchy in which the root servers reside at the top of this hierarchy [330]. There are a total of 13 distributed root servers designated A through M. DNS service disruptions can occur due to misconfigurations [331, 332] or due to attacks [333–335].

DNS attacks can be categorised as cache poisoning, denial of service, domain hijacking, and compromised data [11, 336–338]. In October 2002 all 13 root name servers were exposed to a DDoS attack. During the 75-minute assault, attack traffic volume was seen to be between 0.9–2 Gb/s [334, 339]. 5% of the root server queries went unanswered due

to congestion [334] and DNS response messages were delayed [339,340]. Further analysis showed that delays were also observed in prior weeks of the attack demonstrating an attack preparation [340].

In February 2006, TLD (top level domain) name servers came under DDoS attack on two separate days in which each attack duration was approximately 14 minutes [341]. The DDoS attack used bots to launch the attack through DNS recursive name servers [342, 343], which sent small size queries impersonating the IP addresses of the TLD name servers as the source address. Upon receiving the small size query, the compromised open DNS recursive name server replies with a large response. In this case the amplification factor of these attacks were 72:1, since the query was 56 B and the response was 4000 B [344]. The attack traffic seen at the TLD name servers was around 2 Gb/s with an estimated 35,000 DNS recursive name servers sending responses [344].

In February 2007, root-name servers were attacked twice in an 11-hour window [334,345]. The primary target of this DDoS attack was the four root name servers F, G, L, and M [346]. The G and L root servers were impacted the worst [345], since they were not using anycast [347, 348]. The 5,000 bots that launched the attack were distributed all over the world [346]. In this case the attacks had only limited impact on the end users [345, 346].

DNS cache poisoning can occur when a DNS response is incorrect. In March 2010 for almost three weeks, the responses from the I-root name server in China were incorrect for facebook.com, youtube.com, and twitter.com [322, 349]. The same issue surfaced in June 2010 [350] and November 2010 [351]. The I-root name server is managed by the Swedish Netnod. To increase service resilience, the I-root name server is geographically located in 36 different locations, including one location in China [352]. When a request is made to a root server, no matter what the instance is, the response is supposed to be

same. However the requests for facebook.com, youtube.com, and twitter.com got invalid IP address responses from the Chinese instance of the I-root name server resulting in *cache poisoning* [351]. 57% of network prefixes worldwide were affected by the censorship of China [353]. However, it is not only Chinese government imposing censorship on the Global Internet, countries all over the world also have low rating for the Internet censorship [354, 355].

ICANN, which administers the L-root name server, announced in October 2007 that it would change its IP address from 198.32.64.12 to 199.7.83.42; however, the old IP address would continue to work for the next six months [356]. Despite the announcement for transition of the new IP address, CommunityDNS (AS 42909), ep.net (AS 4555), and Diyixian.com (AS 9584) continued to announce the old IP address of the L-root name server [357–360]. Following pressure from ICANN, the announcements were terminated two weeks after the deadline of the use of the old IP address [360]. Despite no visible impact to end users due to the bogus announcements since the announcers replied to queries with legitimate responses, the event shows a deviation from the *normal operation*. The experience also caused significant *privacy concerns* since it was believed that end-user traffic might have been surveilled and stored for further analysis [359–362].

Domain name hijacking is another form of attack to which the DNS can be exposed, in which the rightful owner of the domain losses the control to the domain [363]. The consequences can be financial loss[1], damaged reputation, and disruption to network services. One form of phishing is domain name hijacking [363,365] in which users are tricked into thinking they are using a legitimate web site. Some of the high profile incidents have been analysed in ICANN reports [363, 366]. During December of 2009, twitter.com was hijacked and unavailable for almost 2 hours due to twitter.com's DNS service provider having its password compromised [367].

[1]Cybersquatting can also cause financial losses; however, this is primarily a policy and legal issue [364].

Social, Political, and Economical Events

We presented challenges caused by natural phenomenological events, malicious actions, and accidental misconfigurations that can disrupt communication network services. However, disruptions can also occur as a result of social, political, and economical events. Terrorist attacks, political unrest, peering disputes, and censorship are examples of such events. Moreover, cyber attacks can be motivated by a combination of social, cultural, political, and economical reasons [368].

Terrorist attacks occurred on 11 September 2001 in the US. The terrorists did not directly target the communication infrastructure; however, *collateral* damage was done to communication networks as a result of the terrorist act, such as the collapse of the Verizon building. The impact on the Global Internet was not significant; however, news websites and telephone networks were overloaded [369] in the resulting flash crowd [54]. Network reachability of less than 1% of the global routes was lost for several hours, primarily due to cable cuts and power losses [370]. Although there was little global impact, some European ISPs whose access to the Internet was served through NY experienced disconnectivity problems [369].

Political unrest is a type of challenge motivated by social and political factors. The first known government-induced network challenge was targeted against satellite networks in Bangladesh in 2007 and in the same year the Internet in Myanmar was *disconnected* [371, 372]. The recent political events in North Africa and the Middle East caused disruptions in both regions. In some cases this extended to eventual disconnected nations from the Global Internet [373]. Beginning 27 January 2011 Thursday approximately at 22:30 UTC, Egyptian networks were withdrawn from the global routing tables by the order of the Egyptian authorities [374–377]. Approximately 3,500 networks among the Egypt's ISPs were withdrawn. The only ISP, Noor Group (AS 20928) that provided connectivity

for the government, financial, and educational institutions also withdrew from the global routing tables on 31 January 2011 around 21:00 UTC [375, 378]. On 2 February 2011 9:30 UTC Egyptian networks appeared again in the global routing tables [373, 379–381].

In February 2011, 13 globally routed networks from Libya were withdrawn, disconnecting Libya from the rest of the Global Internet [377,382]. However, unlike the Egypt case, the withdrawn prefixes were not long lived. On the other hand, traffic in and out of Libya was blocked [372,383]. In August 2011, 6 to 11 networks out of a total 16 networks being advertised by Libya had sporadic outages [384, 385].

The 40 out of a total of 59 network prefixes were withdrawn on 3 June 2011 for almost 24 hours in Syria [386, 387]. The traceroute measurements suggested that during the Syrian Internet blackout, some networks were still up (e.g. government), and the delay measurements were lower than usual, indicating backbone networks were operational; however, access networks were brought down [388]. Later, in August 2012 as the conflict in Syria continued, there were two instances in which one of them resulted that all prefixes from Syria were withdrawn from the global routing table for 17 minutes, and the other event resulted withdrawal of 20 prefixes sporadically over a five hour period [389]. More recently, on 29 November 2012 at 10:26 UTC all 84 IP prefixes that belong to Syria were withdrawn from the global routing table [390–393].

Depeering caused by business and economical factors is another challenge to the Internet [394]. The Internet is interconnected through ISPs that are hierarchically related. While tier-1 ISPs peering with each other use a financial *settlement-free mode*, these transit providers charge customer ISPs. The *peering tactics* in BGP are several and they are complex for a variety of situations [395]. The disputes among the ISPs resulted in disconnection of networks from the Internet. In October 2005, the peering dispute between Level 3 (AS 3356) and Cogent (AS 174) resulted in almost 2 days of disconnection

between single-homed customers of both providers interconnected to one another [396]. During that time, 4.3% of the total prefixes from the global routes were isolated from each other [397]. The peering dispute between Cogent (AS 174) and Telia (AS 1299) in March 2008 resulted in the disconnection of stub AS networks of each provider from each other [398, 399]. The depeering between Cogent and Telia lasted for 15 days and it is estimated that the Telia customers were more severely impacted [400, 401]. At the end of October 2008, Cogent (AS 174) and Sprint (AS 1239) depeered for three days [402, 403]. At the time, there were 214 single-homed ASes downstream of Sprint and 289 single-homed ASes behind Cogent. An estimated 3,500 networks did not have full Internet connectivity [402].

Net neutrality fueled by economic interests is a challenge to the Internet. Net neutrality arguments primarily focus on two polar ideas. Proponents of net neutrality argue that an open Internet model is the reason behind innovation in the Internet ecosystem. On the other hand, opponents of net neutrality argue that providers should have total control of their traffic including the ability to discriminate based on the *source* or *provider* of content, an argument that goes far beyond QoS on metered pricing and traffic shaping [404, 405]. A third model discourages market dominance and allows QoS in the Internet to exist [406]. In 2005, the FCC ruled in favor of Vonage, who argued that its traffic should not be blocked by an ISP [407]. Other high profile cases include CAIP (Canadian Association of Internet Providers) vs. Bell Canada [408] and Comcast vs. BitTorrent in 2008 [409]. In 2012, Verizon Wireless settled with the FCC for 1.25 million dollars and allowed third party 802.11 tethering applications to run on smartphones used in its network [410]. We should remember that the entire point of building a network infrastructure is to support the distributed applications that need it [411].

Censorship can be triggered by political [412] or economic [413] causes. Either way it impacts the *availability* of information to the end users. Censorship mechanisms in-

clude IP address filtering, URL keyword filtering, and DNS redirection mechanisms [414]. Furthermore, filtering mechanisms can be deployed deep inside a network to achieve censorship. For example, in China, filtering is not only being applied on the border routers, but as deep as 13 hops inside of the Chinese border routers [415]. While censorship has local effects (e.g. within a country, a peer-to-peer application), as we mentioned previously, misconfiguration of Chinese instance of the I-root name server had a global impact [322, 349–351, 353, 354]. Even though China is a canonical example of state-sponsored censorship, other examples such as North Korea, Iran, and Syria exist around the world [355, 415]. The level of network censorship in authoritarian countries is severe. For example, North Korea's highly-regulated network infrastructure is not part of the Global Internet, but rather it is considered to be a national intranet [416, 417].

Privacy is an aspect of trustworthiness and entails authorised communication between end systems [179]. The privacy of communication can be threatened by economy-induced factors such as web cookies [418–420] or policy-induced factors such as surveillance [421, 422]. The 2005 FCC ruling that VoIP (Voice over Internet Protocol) must comply with CALEA (Communications Assistance for Law Enforcement Act) requires that wiretapping capabilities be deeply embedded in the Internet protocol stack. The result is that the Internet is *less secure* [421]. Additionally, wiretapping capabilities are against the end-to-end arguments [272, 273], since successful wiretapping requires wiretapping capabilities to be deeply implemented in the protocol stack [421]. Furthermore, the provision of surveillance capabilities of the network can violate the privacy of communications of the innocent and can be used as an attack mechanism by outside actors [421, 422].

3.1.2 Traffic Tolerance

In this section we provide challenge examples relevant to the *traffic tolerance* discipline. Traffic tolerance is the ability to withstand abnormal traffic conditions. Traffic engineering is a discipline that studies performance of networks and traffic anomalies [423, 424]. The traffic anomalies can be caused by network failures, flash crowds, and attacks [425].

Traffic Tolerance to Legitimate Traffic

Flash crowds are events that are sudden and are due to simultaneous access requests from multiple clients to a target [54, 56, 426]. This leads to service denial as a result of network resource exhaustion. The network resources of interest are bandwidth, memory, and processing power [411]. Flash crowd traffic is legitimate, whereas anomalies such as DDoS (Distributed Denial of Service) attacks are malicious [425]. On the day of the 9/11 terrorist attacks on 11 September 2001, major news websites became unresponsive after the second plane crash into the WTC (World Trade Center) [369]. The demand for the CNN.com website increased by an order of magnitude on 9/11 [427]. Pages were simplified to text only, many pages were removed, and additional Akamai CDNs were employed to deal with the increased demands [369, 427].

London was struck by terrorist acts on 7 July 2005; the cellular telephony networks were overloaded with up to 250% more calls being made that morning [428]. Access Overload Control (ACCOLC) was activated on a cellular network provider in an area with a radius of 1 km, which restricts who can access the network and requires special phones be used by the authorities. ACCOLC was not activated in other parts[2] of the impact zone with the assumption that lack of mobile phone service might cause more public panic [428].

[2]The decision by City of London Police to activate ACCOLC by overriding the command chain gained considerable attention detailed in the report [428].

Telephone networks were overloaded in the aftermath of 9/11 terrorist attacks [369], Hurricane Katrina [189], and the Wenchuan Earthquake [182] as well.

Traffic Tolerance to Attacks

DDoS attacks involve multiple zombies managed by an attacker to overwhelm the target by exhausting its resources. DDoS attacks, classification of these attacks, and defense mechanisms are covered in comprehensive surveys and the references therein [176–178]. Furthermore, the US-CERT (US Department of Homeland Security Computer Emergency Readiness Team) database [429] and NIST vulnerability database [430] provide extensive coverage of vulnerabilities, and it is therefore outside the scope of this work to cover every DDoS attack and its impact to the network. However, due to geographic scope in which nations are impacted, we provide the following examples of traffic tolerance.

Estonian networks came under a DDoS attack in April 2007 due to a political decision made by the Estonian government [431]. DDoS attacks against the websites eventually forced Estonian network providers to disconnect themselves from the Global Internet. Recent Stuxnet attacks [432, 433] that targeted industrial control systems, Operation Aurora, and Titan Rain were some of the politically motivated cyberwars [434–436]. Furthermore, hactivist groups such as Anonymous are responsible for politically motivated cyber attacks against government websites [437–439].

During 2009, government-induced *traffic engineering* in Iran resulted in the throttling of its own bandwidth [440]. Following the elections on 12 June 2009, global route instabilities involving 400 prefixes were observed [441]. The state run Data Communications Iran (DCI) provides Iran's Internet connectivity [440]. The day after the elections, more than 180 prefixes were withdrawn from the global routing tables for an hour; moreover, the routing instabilities continued for the next couple of days [441]. Further analysis

demonstrated that after the prefix withdrawals, DCI (AS12880) outbound traffic reduced from about 5 Gb/s to 1 Gb/s [440]. It became clear that filtering was occurring as evidenced by the impact on web, video, and interactive applications [442]. The overall carried load in DCI and local ISPs showed significant variability in the months of June and July [443, 444]. The citizens of Iran used proxy web servers to overcome the filtering [445].

3.1.3 Disruption Tolerance

The Internet is a collection of heterogeneous components [2, 120, 446]. Fundamentally, the data in communication networks are carried via wired and wireless media. Wireless technologies that link nodes pose a variety of different challenges. Data is transported via an *open channel* and nodes may be *mobile* in wireless networks. Furthermore, mobile nodes can be deployed without a dependency on an infrastructure, thus limiting their resources.

Wireless networks are broadly categorised based on the area that they cover (e.g. wide area network, local area network). Some common examples of these type of networks are satellite networks, cellular telephony networks, and wireless sensor networks. For the networks that have relatively small coverage, the impact of challenge *scope* is *local*. However, for metropolitan and wide area networks, since coverage area is greater, the scope is *regional*. We present some past challenges that involve satellite networks and the PSTN. Since the network outage reports are not publicly available particularly for cellular networks [23], we are unable to present any further specific examples for these networks. However, our focus is to present challenges in the wireless networking environment that cause disruptions.

Mobility

The *anytime and anywhere* paradigm requires non-trivial changes to the network architectures [447]. While end-users enjoy the freedom of untethered access to the network, mobility of the nodes bring extra burden for designers [448]. MSC (mobile switching center) failures in a cellular network lead to roaming failures [449] and BTS (base transceiver station) failures can lead to handoff failures [450]. In MANETs (mobile ad hoc networks), self-organisation and self-configuration is needed to set up the network [53, 109]. However, episodic connectivity due to the mobility of the nodes in MANETs results in routing challenges. Several routing algorithms are proposed to solve such problems [451–454]. Highly-dynamic environments also present identical problems for topology management, location management, and routing management [455], particularly for supersonic aeronautical nodes with relative speeds of up to 7 Mach, due to short contact duration and limited connectivity of the nodes [456, 457].

Connectivity

Since the communication channel is open and the medium is shared, the cost of attacking a wireless network is smaller than that of attacking wired networks. Attacks against WSNs (wireless sensor networks) may be easier since they are often deployed in hostile environments [458, 459]. Jamming the RF (radio frequency) signal is one way of carrying out denial of service attacks in wireless networks [460, 461]. Since WSNs are constrained by limited resources, attack cost against WSNs are even lower [462, 463].

Wireless channels are lossy, meaning the channel has a lower SNR (signal-to-noise ratio) compared to wired channels. Signal levels attenuate as the distance between the source and the destination increases [464, 465]. This distance has a significant effect particularly for satellite networks [466, 467]. However, usage of wireless mesh networks and multihop

routing can alleviate requirement for direct LOS (line-of-sight) communication [468]. The signal quality of wireless links degrades due to meteorological events, such as rain storms [204]. On the other hand, routing algorithms can predict the increased BER (bit error rate) using weather radar information and route packets around storms [206]. Therefore, disconnected operation of links can be masked by the higher layers to a certain extent. Signals can be distorted when they are diffracted, reflected, and scattered from an object on its path from source to destination.

The erroneous channel in the wireless environment particularly impacts the transport layer. TCP cannot differentiate between corruption and congestion, and it assumes losses are due to congestion [469, 470]. When the packets are lost in a high error rate channel, TCP throttles the sender transmit rate unnecessarily; hence, ELN (explicit loss notification) mechanisms are required to differentiate between congestion- and error-based losses [471, 472]. ETEN (explicit transport error notification) can help improve TCP performance by discriminating between congestion- and corruption-based losses, so that sender does not reduce its transmission rate needlessly [467]. A variety of solutions have been proposed for TCP modifications in lossy wireless environments [464, 465, 473, 474]. Note that while mechanisms at the transport layer can request retransmission of original data they cannot recover the corrupted data. Therefore, error correction schemes such as FEC (forward error correction) at the link layer can mask the effects of a lossy channel.

Delay

The TCP/IP protocol suite is designed with the assumption that the end hosts are continuously connected and the propagation delay is relatively low. The assumed delay on wired links is on the order of milliseconds; however, this assumption does not hold for satellite networks in which propagation delay approaches seconds [475, 476]. Moreover,

the propagation delay can increase to the order of minutes or hours for interplanetary networks [477–479]. TCP utilises the three-way handshake mechanism to establish a connection and it uses acknowledgment mechanism to reliably deliver the segments. Short RTT (round trip time) is needed for both of these mechanisms since TCP retransmits for reliable delivery after *timeouts*. Retransmissions due to long propagation delays can consume the valuable bandwidth unnecessarily. Therefore, environments in which long propagation delays exist challenge the transport layer and a number of protocols for delay-tolerant networking have been proposed [474]. Among the most prominent protocols are: SCPS-TP (Space Communications Protocol Standards-Transport Protocol) [480, 481], Bundle Protocol [482], and the LTP (Licklider Transmission Protocol) [483, 484]. Furthermore, unlike terrestrial networks in which the delay is *low* and interplanetary networks in which the delay is *predictably high*, delay characteristics can be *unpredictably high* in habitat monitoring using wireless sensor networks. An example is the MULEs project in which sensed data in a sparse network is transported to a base station opportunistically [485].

Energy

Energy is an essential resource for communication networks. Mobile phones, satellite stations, and MANET nodes require energy to do necessary computations and transmission of communication signals [486]. While the *power grid* directly supplies the necessary energy for desktop computers and network routers, smart phones and laptops use *replaceable batteries* when they are not connected to the power grid. Furthermore, it is difficult, if not impossible, to recharge batteries of energy-constrained sensor nodes deployed in hostile and remote locations [458]. Thus, energy is a necessary resource for communication networks and several mechanisms such as energy scavenging have been proposed to deal with such energy-constrained networks [487].

3.2 Challenge Models

In the previous section, we identify the known and potential challenges to communication networks. In this section, we provide the challenge models and a taxonomy we have developed [161]. First, we review the challenge → fault → error → failure chain and its relationship with the ResiliNets strategy. Then, we discuss the spatial and temporal impact of challenges. Based on the challenges we identify, we provide a taxonomy of challenges. Finally, we present how these challenges are correlated with our taxonomy.

3.2.1 Challenge → Fault → Error → Failure Chain

A *challenge* is an event that impacts normal operation of the network [2]. A challenge triggers *faults*, which are the hypothesised cause of *errors*. Eventually, a fault may manifest itself as an error. If the error propagates it may cause the delivered services to *fail* [13]. The fault → error → failure chain relationship has been extensively studied by the IFIP 10.4 working group [13,14]. We note that while the IFIP 10.4 taxonomy focused on *faults* in computer systems, our focus in this work is taxonomy of *challenges* in communication networks. Challenges to the normal operation of networks include unintentional misconfiguration or operational mistakes, malicious attacks, large-scale disasters, environmental, and deliberate human actions driven by social, political, and economic agendas [2,19–23,53]. The challenge, fault, error, failure chain relationship and ResiliNets strategy (Figure 2.2) is shown in Figure 3.2.

Challenges have primary impact on the defence and detection aspects of the ResiliNets D^2R^2+DR strategy. We can defend against challenges passively by building diverse structural components and technologies, as well as using redundant components [3]. Furthermore, we can strengthen networks by installing active defence mechanisms such as firewalls. However, building a 100% resilient system is not practical due to cost con-

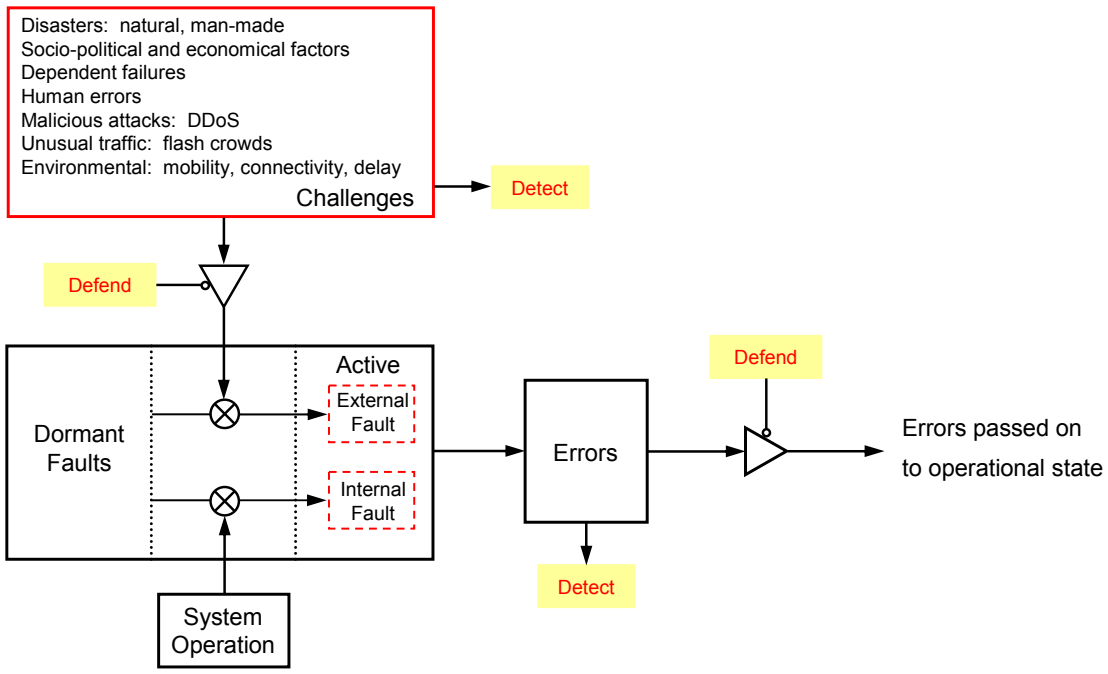

Figure 3.2: Challenge → fault → error → failure chain

straints. For example, while a full mesh interconnection provides maximum robustness to link failures, it is prohibitively expensive to deploy. As a result, defences may be penetrated and challenges activate dormant faults in the system. We note that system operation can also activate the faults; for example, a particular input pattern can activate faulty software code [13]. Some challenges can be detected by using in-network mechanisms, such as signature-based detection against known attacks or behaviour-based detection against flash crowds. Out-of-network detection includes mechanisms that are outside the boundary of the network system. For example, weather storm tracking or an early warning system against an EMP (electro magnetic pulse) weapon can be used as an input to a predictive algorithm to utilise alternative paths [205–207]. On the other hand, it is very difficult to detect some challenges, such as operator mistakes, before they result in failure.

Faults are hypothesised causes of errors [13], and once activated, result in errors that

can be detected using the network management and monitoring systems. Moreover, we can defend against errors by redundant components and cross-layer techniques. An example is having FEC (forward error correction) at the link layer to protect against wireless challenges. By using a cross-layer mechanism, the transport layer can request retransmission of original data if it cannot recover the corrupted data. Thus, we can defend against errors before they are passed to the operational state.

3.2.2 Spatial and Temporal Impact of Challenges

It is important to understand the spatial and temporal characteristics of challenges in order to model them realistically. For survivable operation against threats, a certain geographic distance between data centers has been proposed [488]. For example, a minimum of 32 miles and a maximum of 151 miles have been designated in data center topologies for site separation against several threats [488]. We provide order-of-magnitude temporal and spatial characteristics of some challenges in Table 3.1. For example, the geographic scope of a devastating earthquake can be on the order of 100 km^2 and its impact region on networks might be the same. On the other hand, a fire's geographic scope in a key network node might be on the order of 100s m^2; however, the impact to the communication networks can be larger. The duration of an earthquake can be on the order of seconds whereas recovery of the communication networks can take days. Another example is a policy decision taken by a governing body in which the spatial region and temporal duration of the challenge might not be accurately known. While the impact of a challenge may be only on a nation or a service that impacts users globally, it might take years to revise a policy.

Table 3.1: Spatial and temporal characteristics of network challenges

Challenge Examples	**Spatial Region**		**Temporal Duration**	
	challenge	**impact**	**challenge**	**impact**
earthquake	100s km^2	100s km^2	seconds	days +
fire	100s m^2	10s km^2	hours	days
hurricane	100s km^2	100s km^2	hours	weeks +
solar storm	1000s km^2	1000s km^2	minutes	days +
misconfiguration	node	global	seconds	minutes
malicious attack	node	global	hours	hours
terrorism	100s m^2	global	hours	hours +
policy related	N/A	regional +	N/A	years
depeering	N/A	global	seconds	days
pandemic	global	global	days	months
power blackout	100s km^2	regional	minutes	hours

3.2.3 Challenge Taxonomy

Network challenges can be categorised based on the phenomenological cause, system boundary, target, objective, intent, capability, dimension, domain, scope, significance, persistence, and repetition they impose on the communication networks, as shown in Figure 3.3. Our *challenge* taxonomy is based on the IFIP 10.4 working group studies on *fault* taxonomy [13]. We note that the previously developed taxonomy of challenges focused on an aspect that only considers a resilience discipline such as security [104,176] or a specific functionality such as emergency management [489]. However, our framework and challenge taxonomy is more generalised and covers multiple aspects of resilience disciplines in network systems, including disruption tolerance and dependencies among critical infrastructures.

The taxonomy developed by the IFIP 10.4 working group has focused on *computer systems*. We expand and cover the challenge taxonomy with an emphasis on *network systems*. In accordance, we keep the system boundaries, objective, intent, and capability classes the same as the IFIP 10.4 fault taxonomy [13, 14, 16]. We remove the phase of

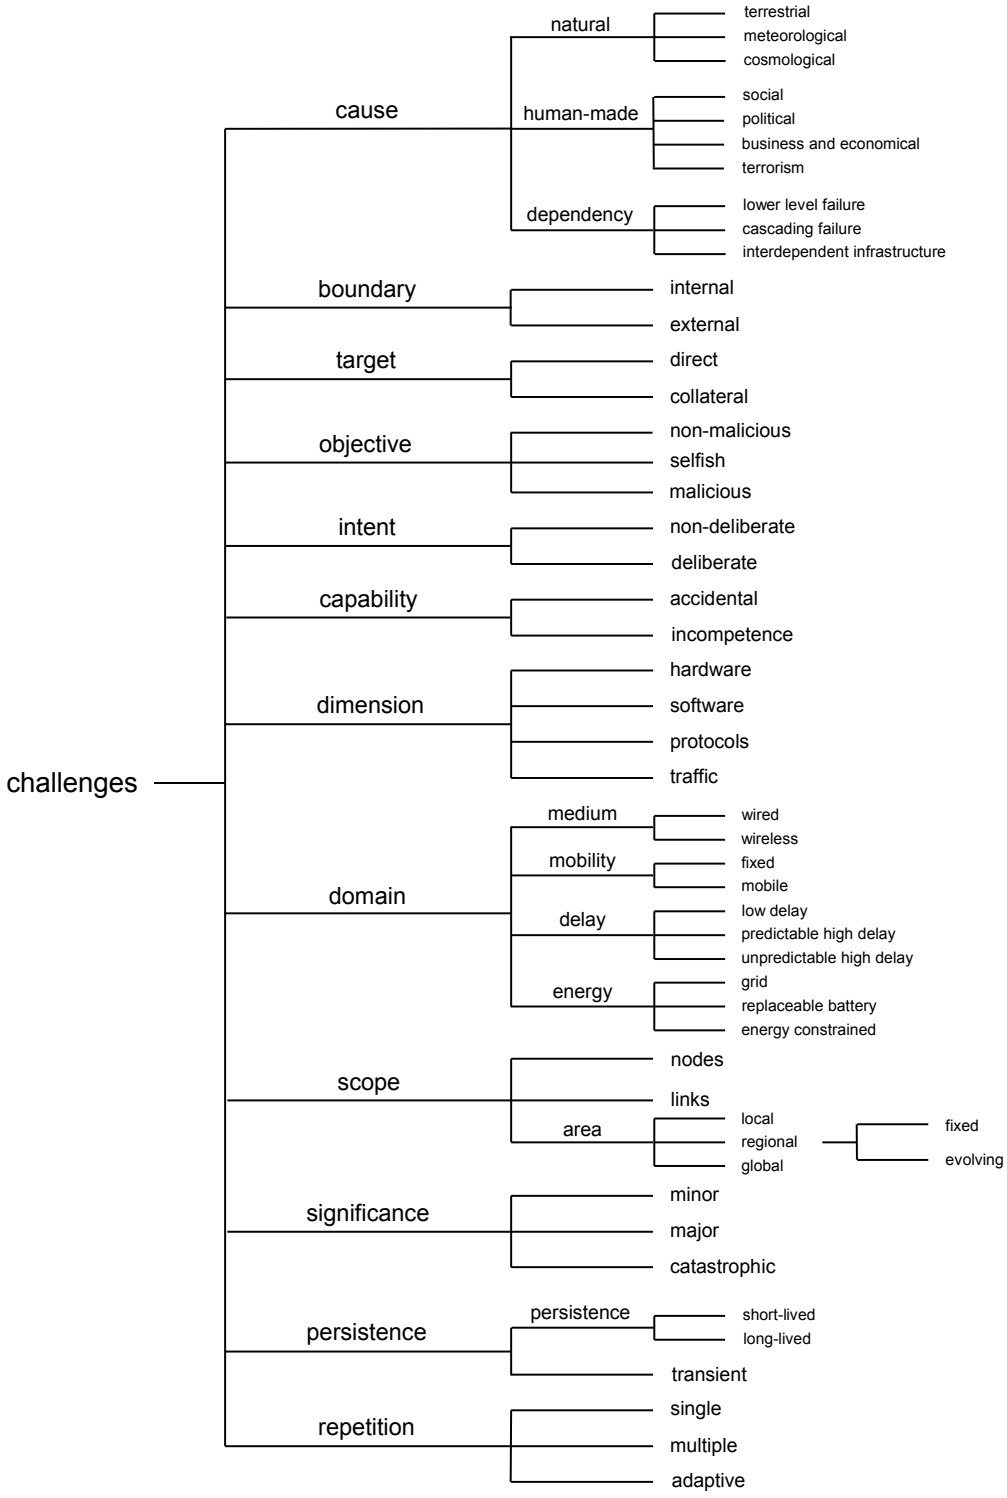

Figure 3.3: Taxonomy of network challenges

occurrence class since it is applicable to faults only (as opposed to a challenge to the existing network). We add target, domain, scope, significance, and repetition classes to our challenge taxonomy. We modify the phenomenological cause class to include a dependency subclass, add a protocols subclass to dimension, and modify persistence to cover challenges that might be long-lived and short-lived. The permanent subclass is eliminated for challenge scenarios. Next, we elaborate on each of these classes.

1. **Phenomenological cause:** The cause of a challenge can be further classified based on ***natural*** causes, ***human-made*** causes, and ***interdependencies*** between the infrastructures. Natural phenomena can occur ***terrestrially*** (e.g. earthquake, fire), ***meteorologically*** (e.g. hurricane, ice storms), or be caused by ***cosmological*** events (e.g. solar storm, space debris). Human-made challenges can be due to decisions driven by ***social***, ***political***, and ***economic*** causes, as well as causes related to ***terrorism***. Examples of such events include recreational crackers, government decisions to block Internet access to nations, and depeering for some financial gain or to increase market share. Finally, phenomenological causes can be due to dependencies within or between the different infrastructures. A failure within the system, at a ***lower level*** can impact the services provided at the higher levels since the services at the higher levels are dependent on the services of lower levels. For example, end-to-end transport is dependent on the lower level hop-by-hop links. Propagation of incorrect BGP announcements is an example of a ***cascading*** failure across the same level within a system. Finally, a power blackout can impact the communication network due to ***interdependencies*** between the power grid and the Internet infrastructures.

2. **System boundary:** The system of interest in which it interacts with its environment can be a single system or a system of systems. For example, while a single

AS (autonomous system) can be considered as a single system, the Global Internet, which is a collection of ASes, must be considered as system of systems. The challenges can be ***internal*** as in the case of BGP cascading failures, and ***external*** to the system in the case of natural disasters. Moreover, defensive mechanisms developed for external threats falls short for threats coming from inside a system.

3. **Target:** The challenges can be ***directly*** targeted to communication infrastructure (e.g. malicious worm) or the network can suffer ***collateral*** damage as a result of a challenge, such as a terrorist activity not directly targeting the network as in the US 9/11 and UK 7/7 attacks.

4. **Objective:** The objective of a challenge can be ***non-malicious*** such as misconfigurations or ***malicious*** such as attacks. Furthermore, a ***selfish*** node or AS can limit network resources in its own interest without a malicious objective.

5. **Intent:** The intent of the actions taken by humans can be ***non-deliberate*** such as misconfiguration errors or ***deliberate*** such as attacks.

6. **Capability:** The challenges caused by humans can be ***accidental*** or due to ***incompetence***. We note that while incompetence refers to lack of professional competence, accidents generally occur as a result of an inadvertent action by humans. For example, BGP prefix hijackings have occurred due to misconfigurations and incompetence of the operator. In the case of the 2003 blackout in the US, one of the causes of the blackout was contact of the power lines with overgrown trees. If the power lines had been laid underground, the catastrophic event could have been prevented.

7. **Dimension:** Challenges can affect the ***hardware***, ***software***, ***protocols***, or the ***traffic*** within a network. For example, random hardware failures fall under the

hardware sub-class of the dimension class, software bugs fall under the software sub-class of the dimension class, and attacks exploiting a vulnerability in a protocol fall under the protocol sub-class of the dimension class. Furthermore, legitimate traffic can impact the services being offered by the network such as the case of flash crowds. We note that DDoS attacks also impact the legitimate user traffic.

8. **Domain:** Challenges vary depending on the domain in which communication network operates. ***Medium***, ***mobility***, ***delay***, and ***energy constraints*** impose different mechanisms to be considered when dealing with challenges. The medium in which nodes communicate can be using ***wired*** or ***wireless*** links. The nodes can be at ***fixed*** locations or ***mobile*** in which topology control mechanisms are fundamentally different. Delay characteristics in which the networks operate also vary: in a terrestrial network a ***low*** delay, in interplanetary communication a ***predictable high*** delay, and in the case of sensor networks for habitat monitoring ***unpredictably high*** delay occurs. Moreover, energy resources are different for networks operating in different domains: while a desktop computer that is connected to ***power grid*** has unlimited energy, a laptop with a ***rechargeable battery*** face different challenges than an ***energy-constraint*** sensor node in a hostile area in which it might not be feasible to replace its battery.

9. **Scope:** The scope of a challenge can impact the ***nodes*** within a network, the ***links*** within a network, and some parts or the entire ***geographic area*** of the network. Geographic scope in which a challenge might impact the network can be ***local***, ***regional***, or ***global***. Moreover, the geographic scope of regional challenges can be ***fixed*** (e.g. earthquake) or ***evolving*** (e.g. hurricane).

10. **Significance:** A challenge's significance can be ***minor***, ***major***, or ***catastrophic***. In the case of the PSTN, the number of lost customer minutes provides a good

measure of the significance of an event. Large-scale disasters such as Hurricane Katrina and the Fukushima Earthquake that caused human and financial losses were catastrophic in significance. Depeering of ISPs in which some customers cannot reach each other is a challenge with major significance, whereas a jammer preventing communication between two individuals may be a challenge with minor significance.

11. **Persistence:** Persistence captures the continuation property of a challenge. The ***persistent*** challenges such as BGP misconfigurations can be ***short-lived*** or ***long-lived***. The majority of BGP misconfigurations are considered short-lived, meaning that minutes after discovery of the mistake, remediation takes place. An example of a long-lived challenge would be a pandemic that affects communication services for months. A challenge can be ***transient*** such as a lightning strike taking down power equipment.

12. **Repetition:** Challenges can occur in ***single*** instances or ***multiple*** instances. While natural disasters are single instance events, malicious attacks might be repetitious. Furthermore, a repeated instance of a challenge that ***adapts*** to failures can cause worse harm.

3.2.4 Correlation of Challenges

In the previous section, we categorise challenges to the network. In this section, we present major challenge groupings and demonstrate the applicability of our taxonomy. The challenges can be broadly listed as follows: Large-scale disasters, socio-political and economic challenges, dependent failures, human errors, malicious attacks, unusual traffic, and environmental challenges.

Table 3.2: Correlation of network challenges

Challenge categories			Large-scale disasters	Soc.-pol. & eco. challenges	Depend. failures	Human errors	Malicious attacks	Unusual traffic	Environ. challenges
cause	natural	terrestrial	×						
		cosmological	×						×
		meteorological	×						×
	human-made	social	×	×		×	×	×	×
		political	×	×			×	×	×
		business & economical	×	×			×	×	×
		terrorism	×	×			×	×	×
	dependency	interdependent infrastructure	×		×				
		lower-level failure	×		×				
		cascading failure	×		×				
boundary	internal			×	×	×	×	×	×
	external		×	×	×		×		×
target	direct		×	×	×	×	×	×	×
	collateral		×	×	×	×			
objective	non-malicious		×	×	×	×		×	×
	selfish		×	×				×	
	malicious		×	×			×		
intent	non-deliberate		×		×	×		×	×
	deliberate		×	×		×	×		
capability	accidental		×		×	×	×		×
	incompetence		×	×	×	×	×	×	×
dimension	hardware		×		×	×	×		
	software		×		×	×	×		
	protocols		×	×	×	×	×		×
	traffic		×	×			×	×	×
domain	medium	wired	×	×	×	×	×	×	
		wireless	×	×	×	×	×	×	×
	mobility	fixed	×	×	×	×	×	×	×
		mobile	×	×	×	×	×	×	×
	delay	low	×	×	×	×	×	×	×
		high	×	×	×	×	×		×
		unpredictable	×			×	×		×
	energy	grid	×	×	×	×	×	×	×
		replaceable	×	×	×	×	×	×	×
		constrained	×				×		×
scope	nodes		×		×	×	×	×	
	links		×		×	×	×	×	×
	area	local			×	×	×		×
		regional	×	×	×	×	×		×
		global	×	×		×	×		
significance	minor					×	×	×	×
	major			×	×	×	×	×	×
	catastrophic		×	×	×	×	×		
persistence	persistence	short-lived	×	×	×	×	×	×	
		long-lived	×	×	×		×		×
	transient		×			×			
repetition	single		×		×	×	×	×	×
	multiple			×	×	×	×		
	adaptive			×			×		

We note that these coarse groupings of challenges overlap with each other partially. For example, a DDoS attack can be categorised under malicious attack as well as under the unusual traffic category. Next, we correlate the challenge taxonomy with the challenge grouping as shown in Table 3.2. In this case, we list the challenge categories from our taxonomy in Figure 3.3 in the first three columns and the major challenge groupings in the last seven columns. We mark a given (category, grouping) cell with an × if that particular challenge group may occur within that challenge category. Furthermore, not all the binary combinations are possible. For example, a malicious attack is caused

by humans, but not by natural phenomena. Such a cross-correlation matrix can be beneficial for correct threat modelling [161]. Next, for each major challenge listed above, we describe its relation to our challenge taxonomy. Note that a comprehensive list of challenges are presented in Section 3.1; therefore, only a select few examples are presented for illustration of each category in this section.

Large-scale disasters can be caused by natural phenomena, human actions, and dependencies among infrastructures. Target, objective, intent, capability, dimension, domain, and persistence aspects of the challenge categories can take any value. On the other hand, the scope of large-scale disasters are not *local* and large-scale disasters are non-repetitive catastrophic events that cause human and financial losses.

Socio-political and economical events are caused by humans challenging communication networks. In the case of nationwide Internet outages these occurred within the nation, thus the system boundary was internal (e.g. Iran blocking its own traffic [442]), whereas DDoS attacks against Estonia due to a political decision was launched from outside of Estonia [431]. While the target and objective category of these challenges can take any value, the socio-political and economical events fall into deliberate intent and incompetence capability of our challenge category. Such social, political, and economic events impact the protocol and traffic dimensions across the wired and wireless domains of challenge categories. In the case of a nationwide Internet outage, the impact of the challenge scope is regional, whereas a policy decision can have global impact on networks with a major or catastrophic significance. During the Arab spring, Syria's network prefixes were withdrawn from the global routing table *multiple* times (3 June 2011 [386,388], 19 July 2012 [490], 18 August 2012 [389], 29 November 2012 [390–393]). Furthermore, in the case of political unrest in Egypt, social networks were initially blocked on 25 January 2011 [491] along with suspension of the mobile telephony service in certain areas [492]. This was followed by the withdrawal of most network prefixes from the global

routing table on 27 January 2011, except the prefixes that belong to financial institutions [374,378]. Eventually, all network prefixes in Egypt were withdrawn on 31 January 2011 [375, 376, 378], showing an *adaptive* challenge. After more than a week, network services in Egypt returned to normal on 2 February 2011 [379].

Dependent failures occur as a result of the failure of one system that provides service to another one. For example, critical infrastructures such as the power grid and the Internet are becoming more dependent on each other. If the power fails, communication networks can halt as a *collateral* result. The power grid increasingly requires the Internet to transport its SCADA (Supervisory Control and Data Acquisition) [493]. On the other hand, a service failure at a lower level is a *direct* challenge against higher layers. BGP cascading failures are also a direct target against communication networks. The capability of dependent failures are due to accident or incompetence. They impact the hardware, software, and protocol dimensions of the network system across the wired and wireless domains. While the dependent failure's scope can impact nodes, links, and areas, the significance of this challenge can be major or catastrophic. Dependent failures are persistent and repetitious, but not adaptive.

Human errors can directly impact the networks or can cause collateral damage. These are non-malicious activities and occur as a result of non-deliberate or deliberate intent. Operational mistakes occur accidentally or due to incompetence. The dimension, domain, scope, and significance of these challenges vary. Operational mistakes are generally short-lived or transient. There can be a single occurrence or multiple repetitive occurrences.

Malicious attacks are caused by humans directly targeting networks with a malicious objective and deliberate intent. For example, a bot can exploit the vulnerabilities if the host is not properly secured, and this lack of secure perimeter can be accidental or due to incompetence. Dimension, domain, scope, and significance properties can take any value.

Malicious attacks can be short-lived or long-lived. Moreover, they can be single, multiple, and adaptive. We note that the system boundary for attacks can be internal in which most attacks come from insiders [494] or external. On the other hand, a non-malicious user writing her password on a sticky note and attaching it next to her computer monitor is a human error with incompetence capability in which an insider or outsider can exploit this to attack the network [495].

Unusual traffic, such as flash crowds, is caused by humans. These events target networks directly with a non-malicious or selfish objective. The intention of users who want to access information is deliberate; however, their intent is not to consume all of the network resources. Therefore, this is a non-deliberate event. In the case of a flash crowd event, the network is overwhelmed with requests by users who do not cease trying to access the network resources. If the users understand the situation in a flash crowd and back off, then the resources may be available after a time period; however, the network resources may still not be available at the instant users request. Therefore, we designate this case as incompetence, since users do not know how the network operates and continue trying to access network resources. The impact is on the traffic dimension of the challenge categories. Unusual traffic impacts network resources on nodes and links. This kind of challenge has minor and major significance, since the network might be operational; however, network services can be limited.

Environmental challenges are inherent in the wireless communication medium, such as rain storms and CMEs (coronal mass ejections), therefore the cause can be natural with a non-malicious objective and non-deliberate intent. Moreover, connectivity on a wireless link can be disrupted by a malicious jammer driven by socio-political and economical reasons. As explained in malicious attacks, capability can be due to accidental or incompetence. Their impact is on the traffic and protocol dimension of the challenges. They only impact the wireless medium, impacting links, and have a local and regional

area scope; however, in the case of interplanetary communication, the scope of disruption is larger. Environmental challenges have minor or major significance with long-lived and non-repetitious characteristics.

The taxonomy of challenges along with correlation table presented in this dissertation is a methodology to gain deeper understanding of past and potential challenges, as well as designing future networks. In summary, a wide range of past and potential challenges exist and we describe with examples of how the challenges correlate with our taxonomy in Table 3.2. By considering the dimension, scope, significance, and persistence challenge categories, large-scale disasters and malicious attacks can cause the worst harm to networks. Their impact can be global in scope, and they can be long-lived, resulting in catastrophic service failures. Moreover, an attack that adapts to defensive measures can be even more harmful. Environmental challenges, such as delay, mobility, and connectivity are only applicable to the wireless domain, and these challenges should be considered during the design phase. In other words, wired networks can be strengthened using redundancy and diversity; however, the same is more difficult for wireless networks. The capability category is primarily applicable to human errors and not applicable for most of the challenge examples. Incompetence and accidental challenges can be avoided by proper training of the operations personnel. Among the social, political, and economical challenges, nationwide Internet outages are the worst, since a country can be disconnected from the Global Internet. As presented, there exists a wide spectrum of challenges, and we cannot avoid them; however, with careful planning, the consequences can be alleviated.

3.3 Summary

Networks face a variety of challenges that disrupt normal operation. Understanding these challenges is necessary for developing correct threat models to design resilient networks that are cost-efficient. Based on past and potential challenges that are summarised, we present a taxonomy of challenges that can be beneficial to evaluate network design choices. Furthermore, we describe how these challenges correlate with our taxonomy.

Chapter 4

Modelling Complex Networks

The Internet has evolved to become a multilevel infrastructure critical to the functioning of society. The multilevel behaviour emerged in part due to the fact that protocols interact in multiple levels and in part because of the ways in which players operate, provide, and use the services of the Internet. Over the years, studies by the research community investigating the resilience of the Internet have suggested controversial findings [496], one such being that an attack on a few central nodes could bring the entire Internet down. But this claim was dismissed by other researchers [497,498] based on the mesh-like structure of actual service-provider backbones. Therefore, *realistic* models are required to mathematically understand the properties of the Internet and improve its resilience.

The work presented in this chapter has resulted in several publications. We showed the structural similarities between transportation and communication networks using the normalised Laplacian spectra [75]. We developed a formal multilevel graph model and a framework to analyse flow robustness of a multilevel graph [499]. Furthermore, we extended this to include multiprovider graphs that capture logical IXP (with exchange providers) links [76]. Finally, we showed that physical level topologies can be modelled by Gabriel graphs, since both are grid-like structures [76,500]. The rest of this chapter is organised as follows: We present the topological dataset we use in Section 4.1. We present

structural similarities between critical infrastructures using graph spectra in Section 4.2. The multilevel and multiprovider graph model is presented in Section 4.3. Lastly, we present how well synthetic graph models capture the structural properties of physical level networks in Section 4.4. Section 4.5 concludes with a summary of this chapter.

4.1 Topological Dataset

We study real networks (i.e. transportation and communication) that are geographically located within the continental United States. Therefore, we only include the 48 contiguous US states, the District of Columbia, and exclude Hawaii, Alaska, and other US territories. Furthermore, we have developed the KU-TopView (KU Topology Map Viewer) [5] using the Google Map API and JavaScript to visually present and assist in analysis of these topological maps. Unlike other visualisation tools, KU-TopView makes raw data conveniently available in the universal form of an adjacency matrix along with the node coördinates and permits their manipulation. We have made these topologies publicly available [51].

4.1.1 Transportation Network

We have generated the freeway topology to represent the transportation network. Our starting point is the American Association of State Highway and Transportation Officials (AASHTO) data, which lists *control cities* and their sequential listing along each interstate highway. A control city is a major population center or destination on or near the interstate highway system determined by each state [501]. However, while generating the transportation topology, we realised that the existing list of control cities was not sufficient to represent the graph accurately. For example, there is no control city at some

interchanges between interstate highways. Therefore, we add 6 additional cities[1] in those cases after verifying the crossing on Google Maps, as well as two that are needed to correspond to physical fibre junctions[2]. There are also a few important newer freeways that are not listed in the 2001 AASHTO document that we add to reflect current connectivity[3]. This US freeway graph with 411 nodes, 553 links, and an average degree of 2.7 is shown in Figure 4.1. We note that in a previous study of US interstate highway system, the authors used GIS (geographic information system) databases from the year 2000 (unfortunately there is no reference to the source of data), and the resulting interstate freeway network consisted of 1337 links and 935 nodes with an average degree of 2.86 [502]. We note that the number of nodes and degree distribution in this *geographic graph* is highly dependent on the number of control cities used for geographic representation and that a number of cities are degree-2 vertices in between higher degree nodes at interchanges. We will discuss a uniform solution to this problem in Section 4.4.

4.1.2 Communication Networks

The Internet is a complex and large-scale network for which collective analysis is non-trivial. Therefore, we restrict this study to include physical fibre and logical level topologies. We note that throughout this work we refer to IP router, PoP, and AS level graphs as *logical* level graphs or L3, whereas fibre level topologies as physical level topologies and denote them as L1. We use Rocketfuel-inferred AT&T, Level 3, and Sprint PoP-level topologies [33, 503] to study logical level topologies. We note that international links, as well as links crossing over Pacific and Atlantic Oceans, are removed intention-

[1]Benton Harbor MI, Country Club Hills IL, Effingham IL, Gary IN, Joilet IL, Lake Egypt IL
[2]Blaine WA, Hannibal MO
[3]I-335 Kansas Turnpike, I-86 East, I-97, I-68, I-495 in NY, and the important non-Interstate US-101 in California between Los Angeles and San Francisco

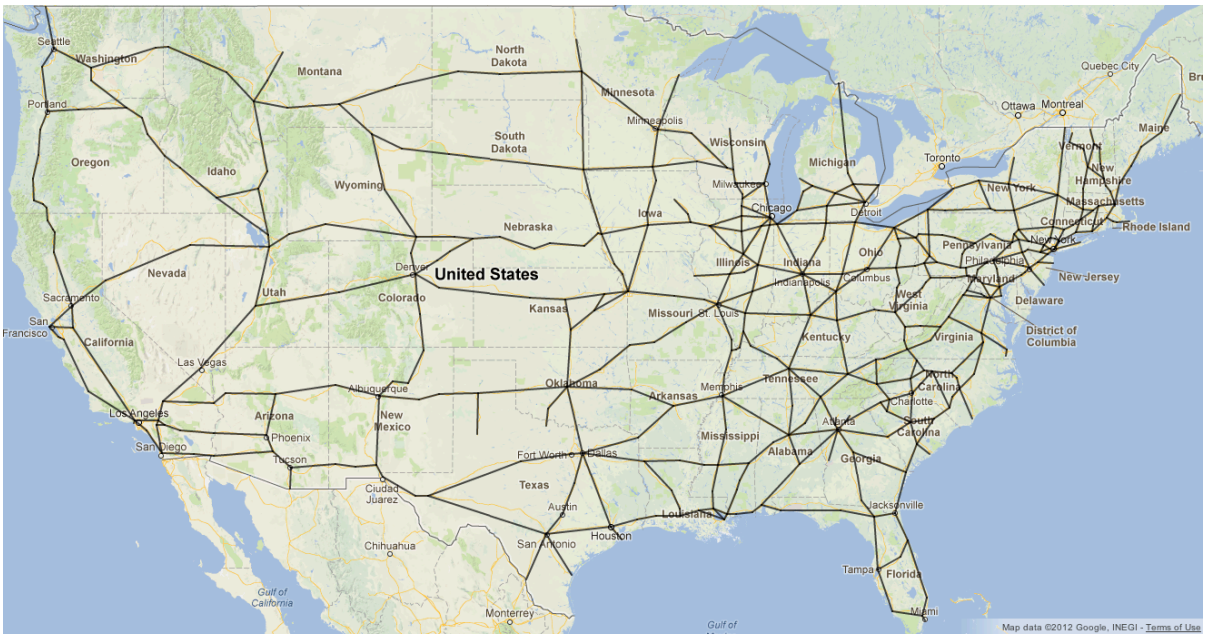

Figure 4.1: Visual representation of US freeways

ally to compare the logical level topologies against the US fibre deployments and freeway topologies.

We then use a US *long-haul fibre-optic routes* map data to generate physical topologies for AT&T, Sprint, and Level 3[4] [505]. In this map, US fibre-optic routes cross cities throughout the US and each ISP has a different coloured link. We project the cities to be physical node locations and connect them based on the map, which is sufficiently accurate on a national scale. We use this data to generate adjacency matrices for each individual ISP. To capture the geographic properties as well as the graph connectivity, cities are included as nodes even if they are merely a location along a link between fibre interconnection. As with the freeway graph, we will further discuss this in Section 4.4. Finally, we also make use of the publicly available TeliaSonera[5] network [506], Internet2 [507], and CORONET [508, 509] topologies. CORONET is a synthetic fibre

[4]We also utilised the Level 3 network map in an effort to reflect the data as accurately as possible [504].

[5]TeliaSonera physical graph has a link between Houston and Miami that appears to cross over the Gulf of Mexico. This is because TeliaSonera does not provide intermediate geographic path information.

topology designed to be representative of service provider fibre deployments, and this does not have a corresponding logical topology as shown in Figure 4.2.

Figure 4.2: Visual representation of CORONET fibre network

The physical and logical commercial service provider networks are shown in Figures 4.3, 4.4, 4.5, and 4.6. The Internet2 research network at the physical and logical level is shown in Figure 4.7. Initial visual inspection suggests that the physical topologies are similar to the freeway topology. The relation of the physical level topology and other physical infrastructures has been stated before [5,510]; however, to best of our knowledge, we are not aware of any previous work that quantitatively demonstrates the correlation between these different infrastructures rigorously.

4.1.3 Properties of Networks

Although topology viewing is a powerful tool, it does not suffice for rigorous analysis of topologies [511]. We therefore calculate the graph metrics of regular networks (shown in

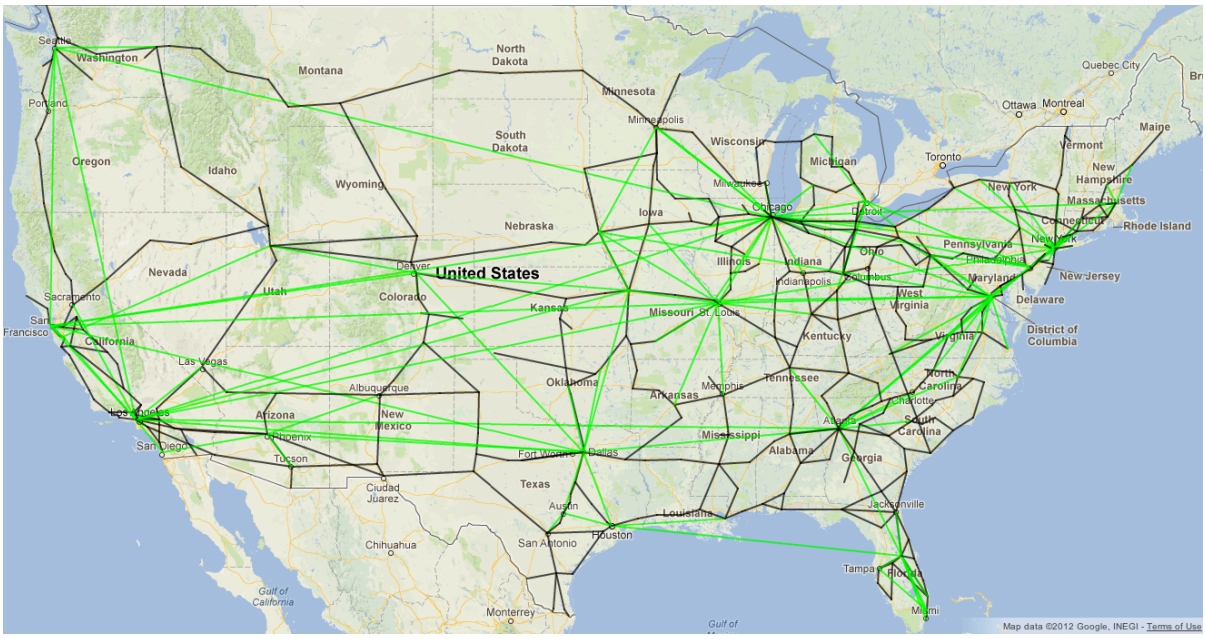

Figure 4.3: Visual representation of AT&T physical and logical level networks

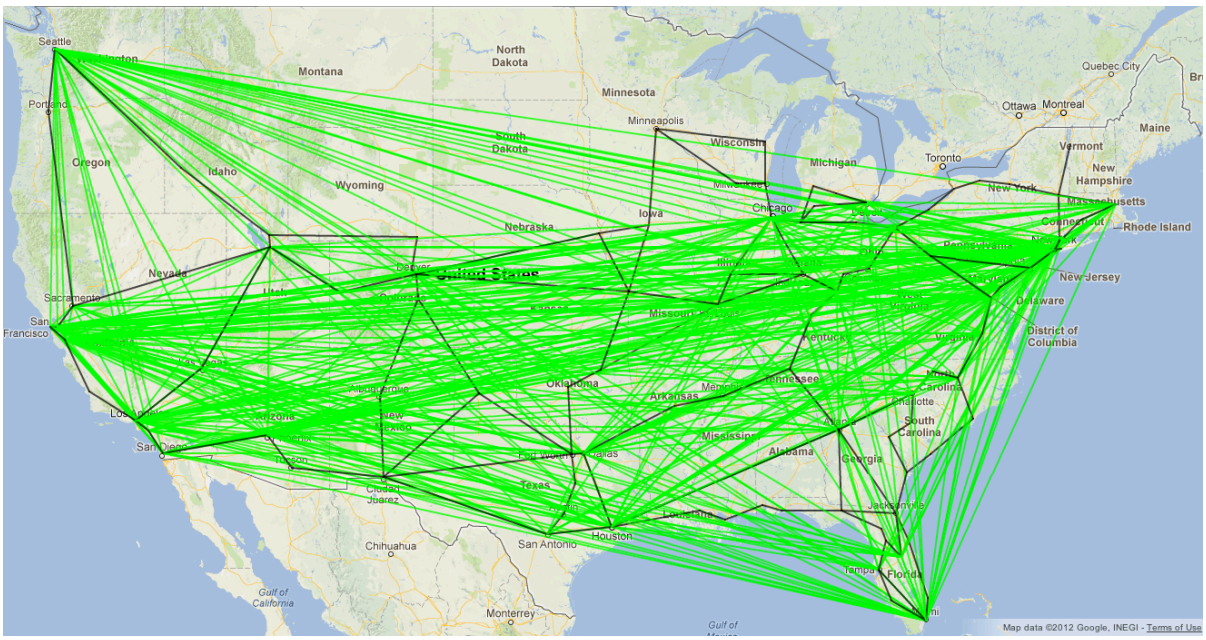

Figure 4.4: Visual representation of Level 3 physical and logical level networks

Table 4.1) and critical infrastructures as shown in Table 4.2 using the Python NetworkX library [512].

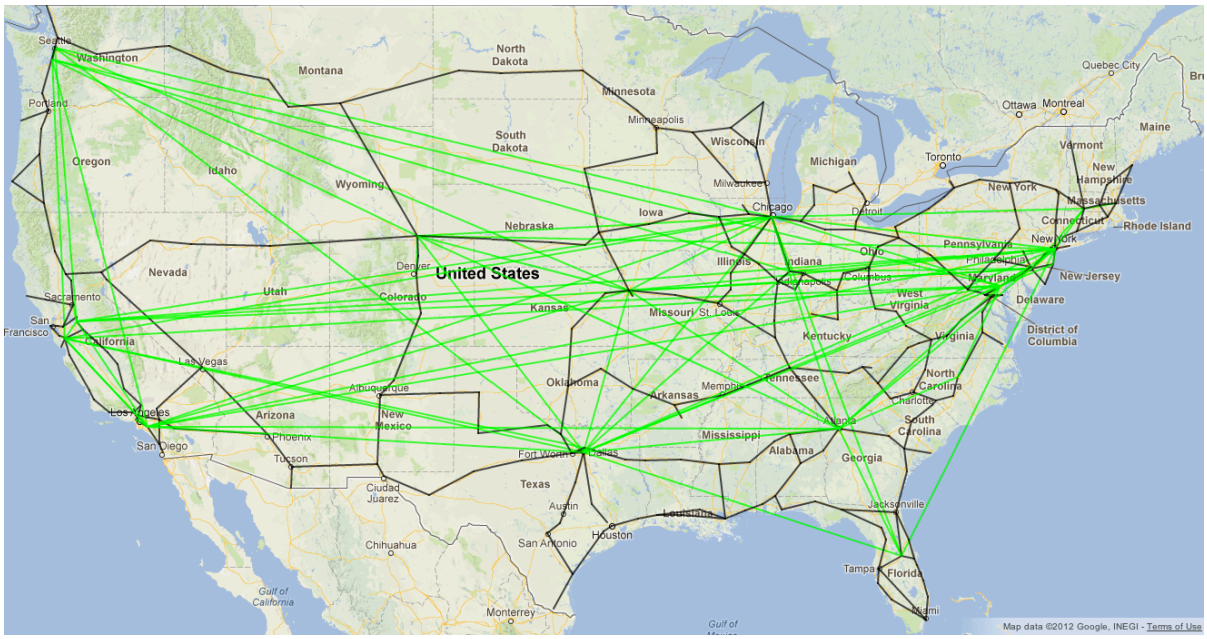

Figure 4.5: Visual representation of Sprint physical and logical level networks

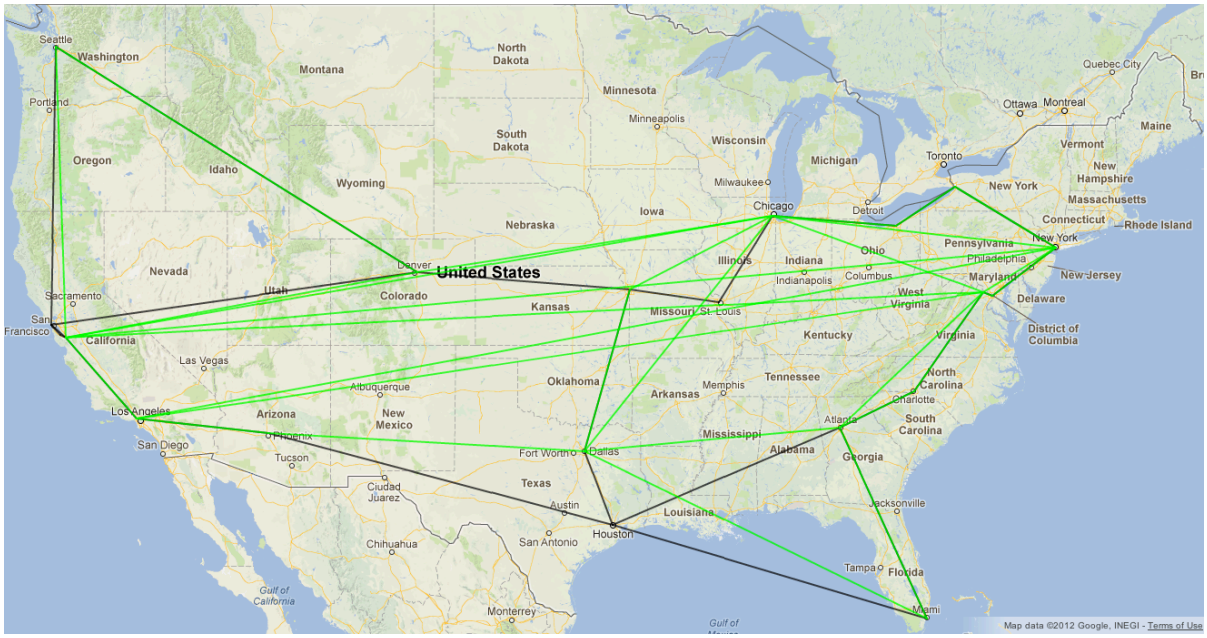

Figure 4.6: Visual representation of TeliaSonera physical and logical level networks

Graph Metrics

Some of the well-known metrics provide insight on a variety of graph properties, including distance, degree of connectivity, and centrality. Network diameter, radius, and

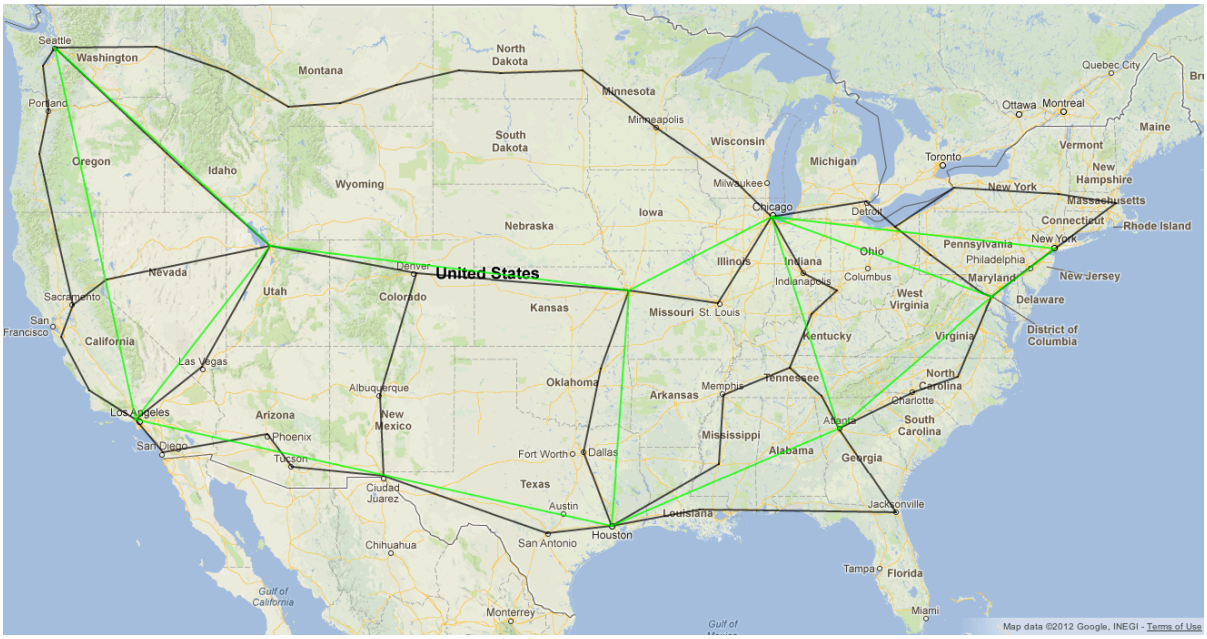

Figure 4.7: Visual representation of Internet2 physical and logical level networks

average hop count provide distance measures [35]. Eccentricity of a node is the longest shortest path from this node to every other node; the largest value of eccentricity among all nodes is the diameter and the smallest eccentricity is the radius. Betweenness is the number of shortest paths through a node or link and provides a centrality or importantness measure [513,514]. Clustering coefficient is a centrality measure of how well a node's neighbours are connected [35]. Closeness centrality is the inverse of the sum of shortest paths from a node to every other node [515]. Assortativity provides a measure of degree variance in a network [516]. Algebraic connectivity, $a(G)$, is the second smallest eigenvalue of the Laplacian matrix [72]. For the graphs of the *same order* (number of vertices), algebraic connectivity provides a very good measure of how well the graph is connected and it indicates robustness of networks against node and link failures [74, 517, 518].

Graph Properties of Baseline Networks

We start our metrics-based analysis on seven regular graphs: star, linear, binary tree[6], ring, grid, toroid, and full mesh. We investigate the effect of an increase in the order (number of nodes) from $n = 10$ to $n = 100$ for the baseline topologies as shown in Table 4.1. We note that the values are rounded to the nearest tenth decimal. The number of edges (links) are increased as necessary for each topology to scale to the number of nodes. Some metrics yield the same values for graphs of the same order (e.g. average degree for star, linear, tree), and others yield the same values for graphs of differing sizes and orders (e.g. same $a(G)$ for 10 node linear and 100 node grid), therefore relying on a single metric for graph analysis is clearly not sufficient.

Table 4.1: Topological characteristics of baseline networks

Topology	**Star**	**Linear**	**Tree**	**Ring**	**Grid**	**Toroid**	**Mesh**	**Star**	**Linear**	**Tree**	**Ring**	**Grid**	**Toroid**	**Mesh**
Nodes	10	10	10	10	10	10	10	100	100	100	100	100	100	100
Links	9	9	9	10	13	15	45	99	99	99	100	180	200	4950
Max. degree	9	2	3	2	3	3	9	99	2	3	2	4	4	99
Avg. degree	1.8	1.8	1.8	2	2.6	3	9	2	2	2	2	3.6	4	99
Deg. assort.	−1	−0.1	−0.5	1	0.3	1	1	−1	0	−0.3	1	0.6	1	1
Closeness	0.6	0.3	0.4	0.4	0.4	0.5	1	0.5	0	0.1	0	0.2	0.2	1
Clust. coeff.	0	0	0	0	0	0	1	0	0	0	0	0	0	1
Algeb. con.	1	0.1	0.2	0.4	0.4	1.4	10	1	0	0	0	0.1	0.4	100
Diameter	2	9	5	5	5	3	1	2	99	12	50	18	10	1
Radius	1	5	3	5	3	3	1	1	50	6	50	10	10	1
Hopcount	1.8	3.7	2.8	2.8	2.3	1.9	1	2	33.7	7.8	25.3	6.7	5	1
Max. Node betweenness	36	20	26	8	11	4	0	4851	2450	3068	1201	616	200	0
Max. Link betweenness	9	25	24	13	12	6	1	99	2500	2496	1250	341	200	1

Graph Properties of Real Networks

We investigate the graph-theoretic properties of the logical and the physical topologies of four commercial ISP networks (AT&T, Level 3, Sprint, TeliaSonera) and the Internet2 research network, as well as the fibre-link level of the CORONET synthetic topology.

[6]We note that not all leaves are binary as needed for a given order.

We also study the US Interstate Highway graph. Our results are shown in Table 4.2. In general, the metrics for the *logical* topologies differ from the *physical* topologies in that the physical topologies have more nodes and links compared to logical topologies.

Table 4.2: Topological characteristics of communication and transportation networks

Network	Nodes	Links	Avg. Node Degree	Clust. Coeff.	Diam.	Radius	Avg. Hopcount	Close.	Max. Node Between.	Max. Link Between.
AT&T L1	383	488	2.6	0	39	20	14.1	0.1	17011	14466
AT&T L3	107	140	2.6	0.1	6	3	3.4	0.3	2168	661
Level 3 L1	99	130	2.6	0.1	19	10	7.7	0.1	1628	1046
Level 3 L3	38	376	19.8	0.8	3	2	1.5	0.7	59	37
Sprint L1	264	312	2.4	0	37	19	14.8	0.1	11275	9570
Sprint L3	28	76	5.4	0.4	4	2	2.2	0.5	100	27
TeliaSonera L1	21	25	2.4	0.2	9	6	4.1	0.3	75	61
TeliaSonera L3	16	29	3.6	0.5	4	2	2.1	0.5	34	17
Internet2 L1	57	65	2.3	0	14	8	6.7	0.2	630	521
Internet2 L3	9	13	2.9	0.4	4	2	2	0.5	9	11
CORONET L1	75	99	2.6	0	17	9	6.5	0.2	1090	704
US freeways	411	553	2.7	0.1	42	21	13.7	0.1	23872	19785

The maximum degree of each provider's physical topology is less than that of its corresponding logical topology. This is due to the ability of logical topologies to arbitrarily overlay virtual links. The average degree of each provider's physical topology is less than that of its corresponding logical topology, in particular for the Level 3 topology in which the average degree for the logical level graph is a relatively highly meshed 19.8. Physical topologies have a higher value of network diameter, radii, and average hopcount than that of logical topologies. Betweenness values also differ for physical and logical topologies, showing a difference of one or two orders of magnitude higher for physical topologies. Clustering coefficient and closeness centrality metrics are also higher for the logical topologies compared to physical topologies.

From a distance metrics (as discussed earlier) perspective, clearly physical topologies have higher values. We observe that the values of degree-based metrics also differ between physical and logical topologies. This can be attributed to the ease with which nodes can be connected in a logical topology as compared to the difficulty involved in connecting

node in a physical topology, in which one must physically lay down fibre between nodes. Long links are added to logical topologies to reduce the forwarding overhead of multihop paths. From a centrality metrics perspective, we can see that physical topologies are not as clustered and have more homogeneous degree distributions.

We can also see that US freeway graph metrics are closer to those of the physical topologies. This is not surprising: both the US Interstate Highway system and the physical level of the Internet are physical infrastructures rather than logical overlays, and they frequently share the same paths since freeways (and railways) provide inexpensive right-of-way along which to lay fibre.

Distinction Between Structural and Geographical Physical-Level Graphs

The physical level topologies consist of a number of degree two intermediate nodes for accurate geographic representation that are necessary for modelling area-based challenges on the network, such as power failures and severe weather. However, these intermediate nodes artificially change the graph theoretic properties of the networks, in particular artificially skewing the degree distribution toward degree-2 nodes. Therefore, we modify the existing *geographical* physical level graphs by removing nodes with a degree of two, as long as there is not a logical level node at that location for which the physical node provides service to upper layers. The topological characteristics of these structural physical-level communication networks are shown in Table 4.3. The cost of structural physical-level networks will be explained in Section 4.4.2.

Structural physical graphs have fewer nodes and links than their corresponding geographical physical-level graphs. However, with the exception of TeliaSonera, each structural graph has a larger average degree than its corresponding physical level graph. For example, the structural graph of Internet2 has 16 nodes, 24 links, and an average degree

Table 4.3: Topological characteristics of structural physical-level networks

Network	Nodes	Links	Avg. Node Degree	Clust. Coeff.	Diam.	Radius	Avg. Hopcount	Close.	Max. Node Between.	Max. Link Between.
AT&T	162	244	3.0	0.1	28	14	9.2	0.1	3592	2936
Level 3	63	94	3.0	0.2	14	7	5.7	0.2	655	568
Sprint	77	114	3.0	0.1	16	9	6.5	0.2	743	602
TeliaSonera	18	21	2.3	0.2	7	5	3.6	0.3	54	43
Internet2	16	24	3.0	0.1	6	3	2.6	0.4	40	33
CORONET	39	63	3.2	0.1	9	5	4.1	0.3	173	133

of 3 whereas the original Internet2 physical graph has 57 nodes, 65 links, and an average degree of 2.28. We believe that the structural graph of TeliaSonera has a smaller average degree than the original graph of TeliaSonera due to the latter's small order and size.

Finally, collective analysis of graph metrics provides a good indication of resilience of different topologies; however, it is difficult to infer sensible conclusions about the structure of a network or how similar two different networks are. Therefore, we redirect our attention to the spectra of these graphs.

4.2 Spectrum of Networks

Let $G = (V, E)$ be an unweighted, undirected graph with n vertices and m edges. Let $V = \{v_0, v_1, \ldots, v_{n-1}\}$ denote the vertex set and $E = \{e_0, e_1, \ldots, e_{m-1}\}$ denote the edge set. A graph can be represented by several methods including an adjacency matrix, incidence matrix, Laplacian matrix, and normalised Laplacian matrix [519, 520]. $A(G)$ is the symmetric adjacency matrix with no self-loops where $a_{ii} = 0$, $a_{ij} = a_{ji} = 1$ if there is a link between $\{v_i, v_j\}$, and $a_{ij} = a_{ji} = 0$ if there is no link between $\{v_i, v_j\}$. The Laplacian matrix of G is: $L(G) = D(G) - A(G)$ where $D(G)$ is the diagonal matrix of node degrees, $d_{ii} = \deg(v_i)$. Given degree of a node is $d_i = d(v_i)$, the normalised Laplacian matrix $\mathcal{L}(G)$ can be represented:

$$\mathcal{L}(G)(i,j) = \begin{cases} 1, & \text{if } i = j \text{ and } d_i \neq 0 \\ -\dfrac{1}{\sqrt{d_i d_j}}, & \text{if } v_i \text{ and } v_j \text{ are adjacent} \\ 0, & \text{otherwise} \end{cases}$$

Let M be a symmetric matrix of order n and I be the identity matrix of order n. Then, *eigenvalues* (λ) and the *eigenvector* ($\mathbf{x}$) of M satisfy $M\mathbf{x} = \lambda\mathbf{x}$ for $\mathbf{x} \neq 0$. In other words, eigenvalues are the roots of the characteristic polynomial, $\det(M - \lambda I) = 0$. The set of eigenvalues $\{\lambda_1, \lambda_2, \ldots, \lambda_n\}$ together with their multiplicities (number of occurrences of an eigenvalue λ_i) define the *spectrum* of M. Spectral graph theory has been extensively covered in several monographs [519–523]. The spectrum of the AS-level topology of the Internet has been analysed based on the k largest values of the adjacency matrix [524]. The IP-level topology of the Internet has also been investigated and its Laplacian spectrum compared against synthetically generated topologies [525]. The normalised Laplacian spectrum of AS-level topologies has been shown to differ significantly from that of synthetically generated topologies [526]. Recently, a weighted spectral distribution metric has been proposed and has shown that synthetically generated graphs can be fine-tuned using spectral properties [66]. While previous studies utilised graph spectra to analyse *logical* level topologies, in this study we focus on *physical* networks and how they relate to each other structurally, as well as to their logical overlays.

4.2.1 Spectral Analysis of Networks

The normalised Laplacian spectrum provides insight into the structure of networks that are different in order (number of nodes) and size (number of links). The eigenvalues of the $\mathcal{L}(G)$ reside in the $[0, 2]$ interval and take values $\{0 = \lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_n\}$. The algebraic multiplicity of $\lambda = 0$ indicates the number of connected components. Hence,

there is always at least one eigenvalue equal to 0. Furthermore, matrices which resemble one another may have similar eigenvalues and multiplicity. The spectrum of $\mathcal{L}(G)$ is *quasi-symmetric*[7] around 1, which means a large algebraic multiplicity for the eigenvalue $\lambda = 1$ may indicate *duplications* in a network [527]. In other words, two separate nodes $\{u, v\}$ might have all or some of their neighbours being same. For example in a star graph with all other nodes connecting to the single central node, the leaves will all have the same neighbour, which is the central node. Likewise, while in a full mesh all nodes have the same neighbours, a partial mesh will have partial duplications. The presence of many small eigenvalue multiplicities may indicate that there are many components within a graph and these components are loosely connected to each other [527]. An eigenvalue of 2 indicates the graph is bipartite; eigenvalues close to 2 indicates the graph is nearly bipartite [527]. A bipartite graph is a graph in which its vertex set can be divided into two groups in such a way that there will be no edges between the vertices within each group. Once the *discrete* and *deterministic* eigenvalues are calculated for a given graph, the *relative frequency* of eigenvalues yield valuable information about the structure of a network. Moreover, spectra can be presented in *relative cumulative frequency* as well, and we describe our choice in the next section. For the rest of this work we abbreviate relative frequency as RF and relative cumulative frequency as RCF.

Spectra of Baseline Networks

The RF (relative frequency) of the normalised Laplacian eigenvalues for baseline topologies (star, linear, ring, tree, grid, toroid, full mesh) of order $n = 100$ is shown in Figure 4.8. Since *most* of the eigenvalues have very small multiplicities, the RF of eigenvalues has a floor that is too noisy to be able to gather useful information. Because of the noisy

[7]We use the term quasi-symmetric to represent almost symmetric graph spectra. For example, a finite full-mesh graph is quasi-symmetric, since all eigenvalues except the first (which is equal to 0) are equal to a value close to 1. We will detail those graphs in the next section.

floor in representing multiple RFs, we use the RCF (relative cumulative frequency) for the baseline graph analysis and for the rest of the work. Furthermore, we note that while some researchers use RFs that they term density of eigenvalues to represent the spectra [527] and others use RCF that they term normalised index of eigenvalues to represent the spectra [526]. Since we show multiple curves in a plot to compare different graphs, our preference is to show spectra using RCFs since it is more informative.

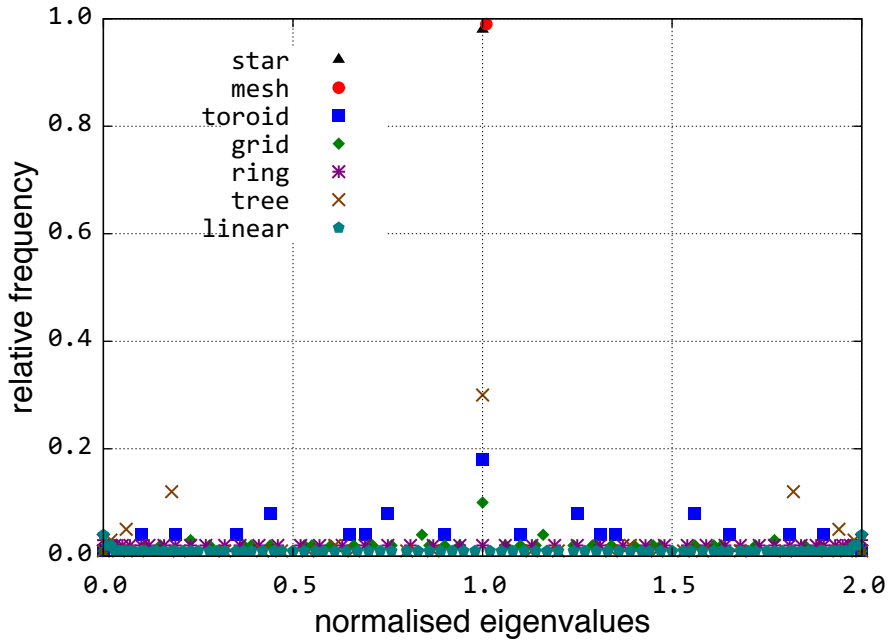

Figure 4.8: Spectra of baseline topologies, RF for $n = 100$

The RCFs of the eigenvalues for these baseline topologies are shown in Figure 4.9. The star topology has its eigenvalues fixed, independent of the graph order: $\{0 = \lambda_1 \leq 1 = \lambda_2 = \ldots = \lambda_{n-1} \leq \lambda_n = 2\}$. The spectrum of a 100 node full mesh looks similar to a star, except that it does not have an eigenvalue of 2 and the eigenvalues are fixed at 1.0101 (we comment on that later). An interesting observation is that the spectrum of these two baseline topologies look very similar. Indeed, at a micro level we can think of each individual node in a mesh as a star motif. Furthermore, the algebraic connectivity of a star is 1 [72]. However, since node centrality measures are largest for a star topology, the

central node in a star can be the target of an attack or the single point of failure from a network engineering perspective. An attack against the root node of a binary tree is also the worst case scenario, however, this partitions the network into two components, in which nodes in each component can communicate with each other whereas this is impossible for a star topology. The spectrum of linear and ring topologies look almost identical, since a ring has an additional link compared to a linear topology, and both linear and ring topologies have the lowest algebraic connectivity values. Likewise, multiplicities of grid and toroid topologies look very similar, since a toroid has additional links to connect the nodes on the edges of a grid. We also observe that since a Manhattan grid is a combination of linear topologies, its spectrum looks similar to a linear topology. Multiplicities of a tree topology lie somewhere between the two extremes of mesh and linear.

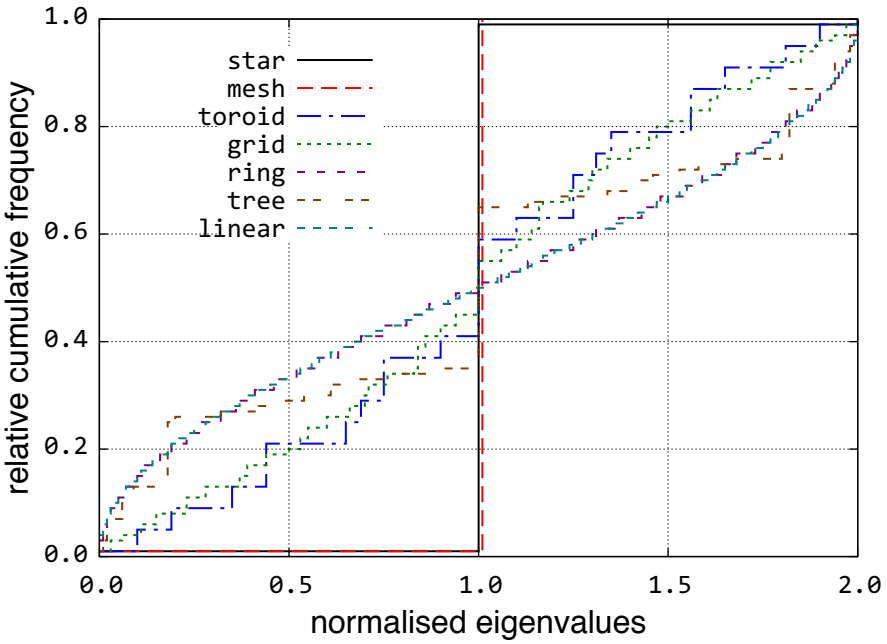

Figure 4.9: Spectra of baseline topologies, RCF for $n = 100$

We show the spectra of five different full-mesh *complete* graphs in Figure 4.10. The eigenvalues of a n-order complete graph are: $\{0 = \lambda_1 \leq \frac{n}{n-1} = \lambda_2 = \ldots = \lambda_n\}$. The

multiplicity of the eigenvalue equal to $n/(n-1)$ for complete graphs is $n-1$. Moreover, as the order of the graph approaches infinity, the eigenvalues will converge to a value of 1 since $\lim_{n\to\infty} \frac{n}{n-1} = 1$. However, eigenvalues λ_2 through λ_n are never exactly equal to 1 in a *finite* full mesh topology. Furthermore, the algebraic connectivity is equal to the order of a complete graph $a(G) = n$.

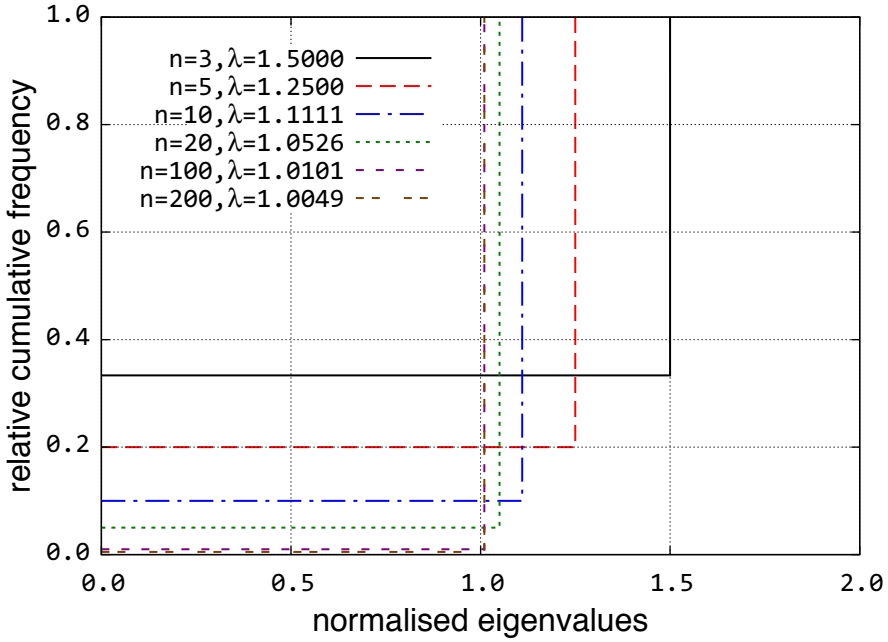

Figure 4.10: Spectra of complete graphs

Spectra of Real Networks

We plot the RCFs of eigenvalues of US freeways against physical[8] and logical level topologies in Figure 4.11 and Figure 4.12 respectively. Clearly, the spectra of the logical and physical topologies differ. Furthermore, the spectra of the physical topologies resemble the spectra of the US Interstate Highway graph as shown in Figure 4.11. This confirms our supposition that the properties of networks are similar since fibre is laid along right-of-ways, such as freeways.

[8]These plots use the geographic version of the physical graphs; this will be explained in Section 4.4.

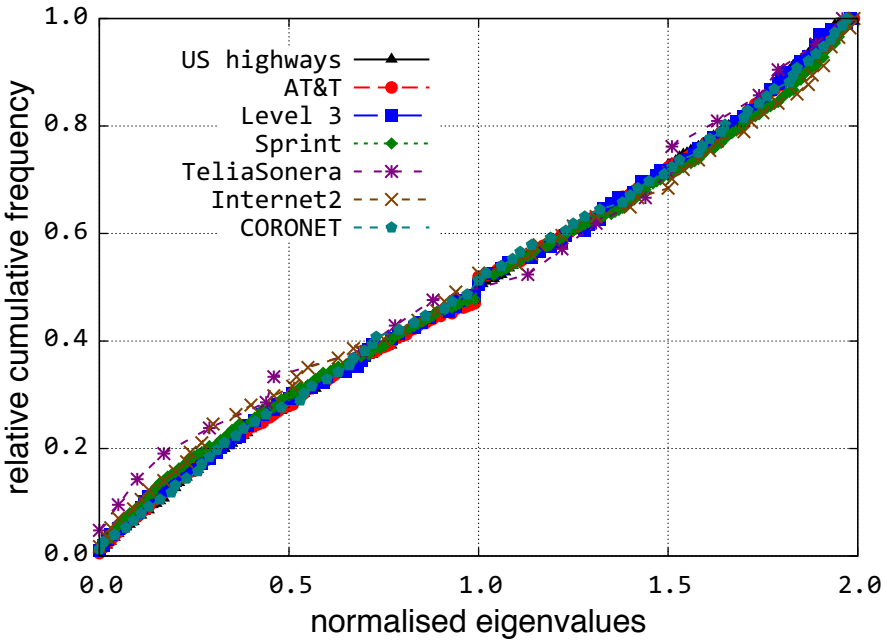

Figure 4.11: Spectra of geographical physical networks

The spectra of logical level topologies along with the US Interstate Highway graph is shown in Figure 4.12. We intentionally include the transportation graph to compare it against the logical level topologies, which clearly shows the spectra do not match to that of freeways. The algebraic multiplicity for the eigenvalue $\lambda = 1$ is largest for the AT&T logical topology, indicating that this topology contains the largest number of node duplications. In other words, this topology has the most star-like components, as is evident by visually inspecting it on KU-TopView [51]. The largest eigenvalues indicate to what degree a graph is bipartite [527]. The largest eigenvalues of the physical topologies and the largest eigenvalues of the freeways graph are the eigenvalues closest to 2. Hence, the physical topologies and the freeways topology are the most nearly bipartite graphs.

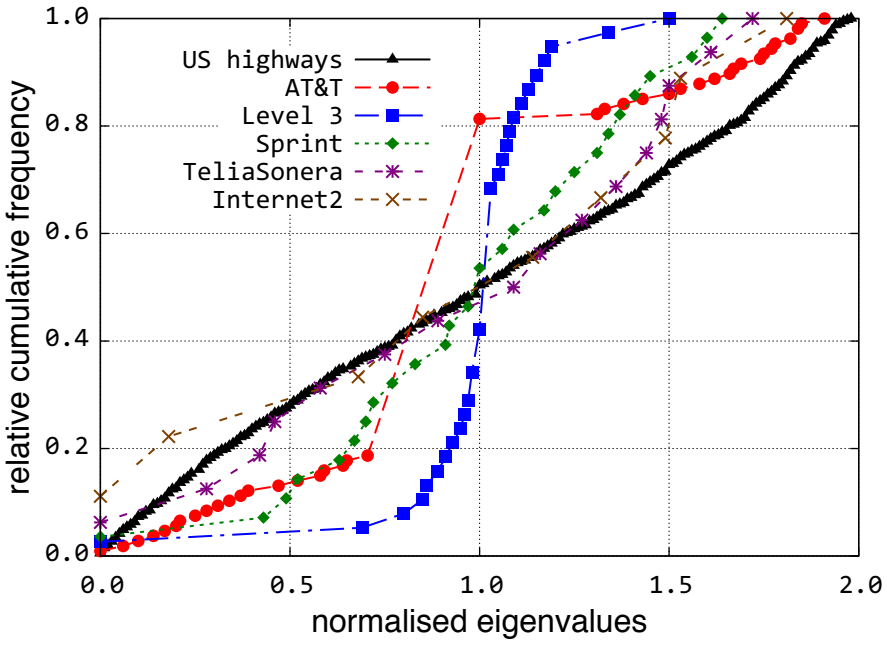

Figure 4.12: Spectra of logical networks

4.2.2 Flow Robustness and Spectral Properties

Previously, we presented how physical communication topologies match the structure of the right-of-way of freeways using normalised Laplacian spectra. In this section, we compare the flow robustness against spectral properties of the networks we study. The two important spectral properties we are interested are algebraic connectivity $a(G)$ and the spectral radius $\rho(\mathcal{L})$. We present these metrics of five logical level topologies, six physical level network topologies, and US freeway graph in Table 4.4.

We measure the resilience of graphs in terms of flow robustness [71, 73], which is the ratio of connected node pairs to the maximum number of node pairs ranging [0,1]. If the graph is not partitioned the flow robustness value is 1, if the graph has no links flow robustness is 0, and if it is partitioned, its value is calculated by adding the number of connection pairs in each component. We approximate the average flow robustness of a given network by averaging the flow robustness over its 10,000 link-failure sets

Table 4.4: Ranking of flow robustness and spectral properties

Network	**Avg. Flow Robustness**	**FR Rank**	$a(G)$	$a(G)$ **Rank**	$\rho(\mathcal{L})$	$\rho(\mathcal{L})$ **Rank**
Level 3 L3	0.9413	1	0.9758	1	1.5037	1
Sprint L3	0.6503	2	0.6844	3	1.6361	2
TeliaSonera L3	0.5963	3	0.7669	2	1.7237	3
Internet2 L3	0.4779	4	0.4885	4	1.8091	4
AT&T L3	0.2996	5	0.1324	5	1.9127	5
TeliaSonera L1	0.1615	6	0.1178	6	1.9642	6
CORONET L1	0.0958	7	0.0401	7	1.9688	7
Level 3 L1	0.0721	8	0.0261	9	1.9811	9
Internet2 L1	0.0626	9	0.0386	8	1.9858	11
US freeways	0.0323	10	0.0055	10	1.9752	8
AT&T L1	0.0222	11	0.0055	11	1.9892	12
Sprint L1	0.0164	12	0.0053	12	1.9840	10

drawn uniformly and randomly from the pool of all of its link-failure sets. Next, we consider the algebraic connectivity $a(G)$ of these topologies. Algebraic connectivity is the second smallest eigenvalue of the Laplacian matrix and is well-suited for measuring graph connectivity and for comparing the connectivities of graphs with the *same order* [72]. Finally, we consider the spectral radius of these 12 topologies. The spectral radius ρ is the absolute value of the maximum eigenvalue, $\rho = |\lambda_{\max}|$. Moreover, if $\rho(\mathcal{L}) = 2$, then the graph is bipartite, and the closer the spectral radii to 2, the closer the graph is to bipartite. We calculate the spectral radius of the normalised Laplacian matrix $\rho(\mathcal{L})$ shown in column 6 of Table 4.4. We note that we previously studied spectral radii of Laplacian matrices $\rho(L)$ and adjacency matrices $\rho(A)$, but did not observe any pattern for $\rho(L)$ and $\rho(A)$ [75].

We rank the flow robustness of networks in descending order in columns 1 and 2. The logical topologies have higher values compared to the physical topologies and the US freeway graph. When we rank the topologies according to descending values of $a(G)$,

we observe a similar ranking order. In this case, only the rankings of the Internet2 and Level 3 physical topologies are swapped. Finally, we rank the spectral radii of these 12 topologies in ascending order. The ranking according to the spectral radii of the first seven topologies matches the rankings of the flow robustness and $a(G)$. Our conclusion from this ranking comparison is that flow robustness, algebraic connectivity, and spectral radii are suitable metrics for the resilience analysis of networks.

4.3 Multilevel and Multiprovider Graph Model

The Internet infrastructure can be examined at the physical, IP router, PoP (point of presence), and AS (autonomous system) level from a topological point of view [31]. An abstract view of different levels of the Internet is shown in Figure 4.13. At the lowest level we have the physical topology, which consists of network elements such as fibre and copper cables, ADMs (add drop multiplexers), cross-connects, and layer-2 switches. The router level consists of devices operating at the IP-layer. A PoP is a collection of routers in a geographic location, and PoP-level topology can be seen as an aggregated view of the routers. At the AS-level, different provider networks peer with each other at the IXPs (Internet eXchange Points) and private peering points [32]. Finally, end users communicate with each other using this multilevel and multiprovider graph. The E2E (end-to-end) level graph depends on users' interactions and requests for information they want to access. For example, users and applications that reside in different ISPs may communicate with each other using the client/server paradigm, or they can form a P2P (peer-to-peer) network to exchange information among themselves in which the E2E level topology resembles a full-mesh structure; examples of this are shown in Figure 4.13. Intuitively, a richly connected lower level can improve the survivability [528] and resilience [2] of a service at higher levels.

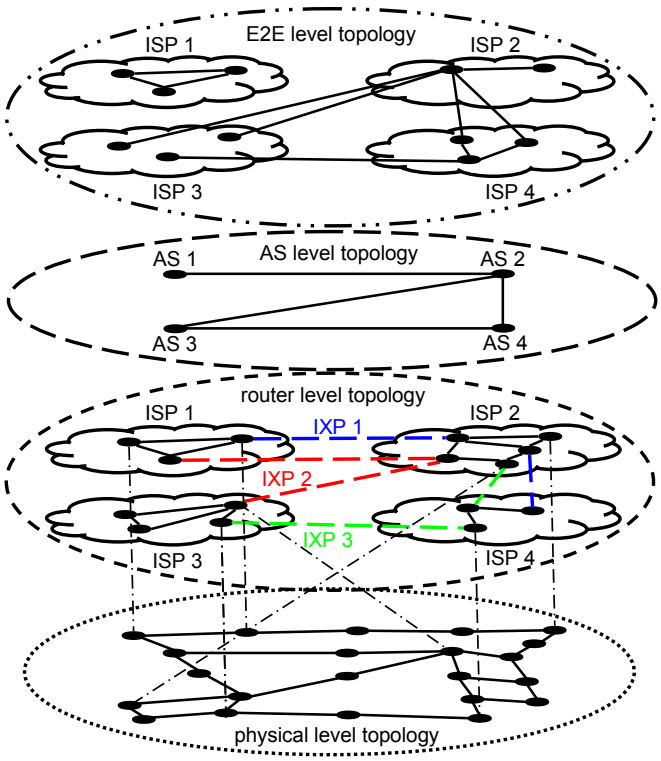

Figure 4.13: An abstract view of Internet graph

A holistic graph analysis that considers the multilevel and multiprovider nature of the Internet is non-trivial and does not exist to the best of our knowledge. Understanding the evolution of the Internet from a multilevel point of view is more realistic than examining its properties at individual levels. Therefore, we develop a formal multilevel and multiprovider graph model and a framework to analyse the *flow robustness* of multilevel and multiprovider networks. When designing a resilient network, our main goal is providing *resilient service* to the users in a cost-efficient manner. Hence, it is extremely important that we ensure connectivity between pairs of end systems. It is for this reason that we use the flow robustness metric as described in Section 4.2.2. Furthermore, we categorise networks in the following four groups [499]:

1. **Single level, single provider:** These networks consist of the physical or logical

level of a single provider. Most previous studies analysed this type of graph [23, 71, 75, 134, 497].

2. **Multilevel, single provider:** These networks consist of multilevel graphs within a single provider. There are a few studies examining multilevel graphs for a single provider [37–39].

3. **Single level, multiprovider:** These networks consist of AS-level graphs that include several provider networks, but as a single adjacency matrix in which each provider is a single vertex of the graph. While several studies analyse AS-level graphs [529], they treat multiprovider graphs at an abstract level (i.e. AS number), and they fail to capture how the ASes peer one another via IXP links.

4. **Multilevel, multiprovider:** This type of model and analysis most realistically captures the complexity of the Internet. To the best of our knowledge, there are no other studies that evaluate the resilience of the Internet from multilevel and multiprovider perspectives simultaneously.

We begin our multilevel analysis of flow robustness of a 3-level graph and a 2-level graph in which the top two level graphs are the same to demonstrate the difference in performance analysis of multilevel networks. We show that the two multilevel graphs exhibit different performance and using fewer levels of graphs obscures accurate resilience evaluation of the top level of a multilevel graph. We then analyse the flow robustness of a number of two-level graphs constructed from real-world communication networks. Next, we analyse a multiprovider graph, which is constructed by aggregating four different ISP networks into a single adjacency matrix. Our results confirm that it is difficult to partition the tier-1 ISP connectivity using attacks targeted at logical links.

4.3.1 Multilevel Graph Model

In an effort to further understand the structure of a number of communication networks, we employ a framework for studying multilevel graphs. A multilevel graph $\mathcal{G}$ is a sequence of graphs, $\mathcal{G} = (G_{\ell_0}, G_{\ell_1}, ..., G_{\ell_{L-1}})$, ordered from lowest-level graph to highest-level graph where:

1. L is the number of levels
2. G_{ℓ_i} is the graph corresponding to level ℓ_i, where ℓ_i can be any desired label, given by $G_{\ell_i} = (V_{\ell_i}, E_{\ell_i})$
3. For all non-negative integers i and j such that $i \leq j$, $V_{\ell_j} \subseteq V_{\ell_i}$
4. For all non-negative integers i and j such that $i \leq j$ and all nodes u and v such that $u, v \in V_{\ell_j}$, if $\text{conn}_{\ell_i}(u, v) = \text{false}$, then $\text{conn}_{\ell_j}(u, v) = \text{false}$, where the function conn_{ℓ_m} takes as its two parameters nodes in V_{ℓ_m} and returns true if the two nodes are connected in G_{ℓ_m} and false otherwise.

In other words, a multilevel graph consists of multiple graphs, one for each level, arranged such that for any pair of levels, the set of all nodes in the higher level is a subset of the set of all nodes in the lower level, and such that nodes that are not connected in a lower level are not connected in a higher level. In this work, we only consider unweighted and undirected graphs. A connected multilevel graph is depicted in Figure 4.14, and when a link is removed at the bottom level, this does not impact the higher level graphs if dynamic routing is utilised as shown in Figure 4.15. Note that in Figure 4.16, the removal of links (1, 6) and (3, 4) in the lowest level partitions the graph and necessitates the removal of all links between the disconnected clusters in the above levels as well.

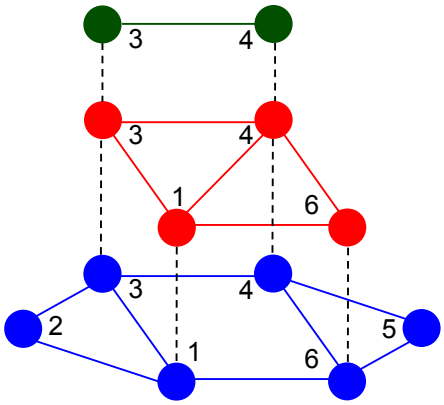

Figure 4.14: Connected multilevel network

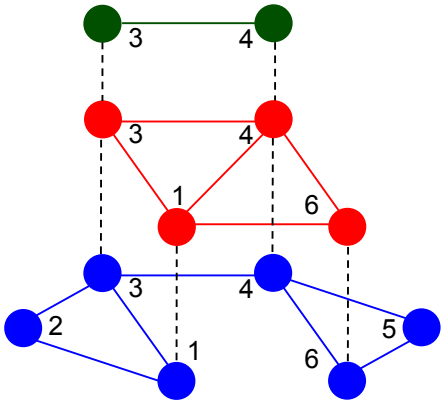

Figure 4.15: Disconnected multilevel network

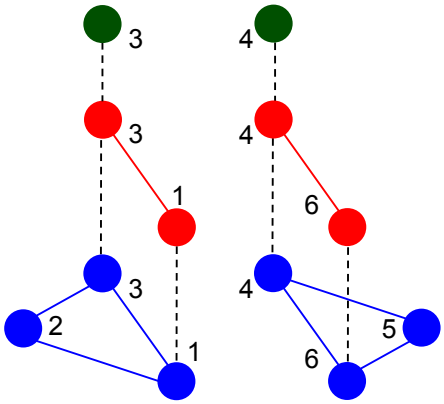

Figure 4.16: Partitioned multilevel network

A number of authors have discussed the importance of multilevel graphs as a means of further studying the resilience and survivability of the Internet [3, 37–39, 115, 530]. Some have developed multilevel graph frameworks of their own [37, 38]. One study made use of

a multilevel framework in order to study railway, peer-to-peer, brain, and random graph topologies [38]. Each topology was subjected to random and loaded [37] link deletions, which were used to simulate errors and attacks, respectively. The robustness of each topology was then quantified in two different ways: as the fraction of logical link weight remaining and as the size of the largest connected component, both as a function of the number of link deletions. In our work, we study challenges [161] on multilevel networks by subjecting topologies to deletions drawn from a far more extensive group of graph metrics. Moreover, rather than treating robustness as the fraction of remaining logical link weight or as the size of the largest connected component, we consider the quantity *flow robustness* as described in Section 4.2.2.

We implement our multilevel model in Python. Our code takes as input a collection of adjacency matrices – one for each level – and stores them in a single multilevel graph data structure in memory, with the following requirements:

1. For any pair of levels, the set of all nodes in the level above are required to be a subset of the set of all nodes in the level below.

2. For any pair of levels, nodes that are disconnected from one another in the level below are also required to be disconnected from one another in the level above.

If the above requirements are met, we can then perform node and link deletions at any level and calculate any number of graph metrics with the help of the Python NetworkX library [512]. When node and link deletions are performed within a given level, the effects of the deletion are propagated to the higher levels to ensure that requirement 2 remains satisfied.

4.3.2 Multilevel Graph Analysis

We first employ our multilevel graph analysis framework to demonstrate the effect using multiple levels of graphs on the *service resilience* [2] at the top level. For this demonstrative analysis, we use a 3-level graph (US freeways, geographical physical, and logical-level topology of Internet2 research network) and a 2-level graph (physical- and logical-level topology of Internet2 research network) in which the top two levels are identical both for 3-level and 2-level graphs. We emphasise that the lowest graph in the 3-level graph is the freeways graph, and it does not provide a service in the conventional sense to the physical topology other than the provision of right-of-way. This is an example to show the impact of using multiple levels of graphs on evaluating the service resilience of the top level. For both the 3- and 2-level network, we perform random node and link deletions at the lowest level and observe how these deletions affect the highest level. Moreover, we consider the effects of these deletions under two separate scenarios – dynamic routing and static routing. Under perfect dynamic routing, we allow any pair of nodes in a given level to remain connected so long as there exists some path between them in the level below. Under static routing, which we show for worst-case baseline comparison, we immediately sever the connection between two nodes within a given level the moment that the shortest path between them in the level below is disrupted.

The results of this experiment are shown in Figure 4.17 for node deletions and in Figure 4.18 for link deletions. For both networks, the average flow robustness of the topmost level is plotted against the number of random deletions performed at the lowest level. For a given number of deletions, the average flow robustness was computed by averaging the flow robustness over 1000 failure sets, each of which was generated by performing the specified number of random deletions. For each value of average flow robustness on the curve, we also plot the 95% confidence interval. We note that the 3-level network

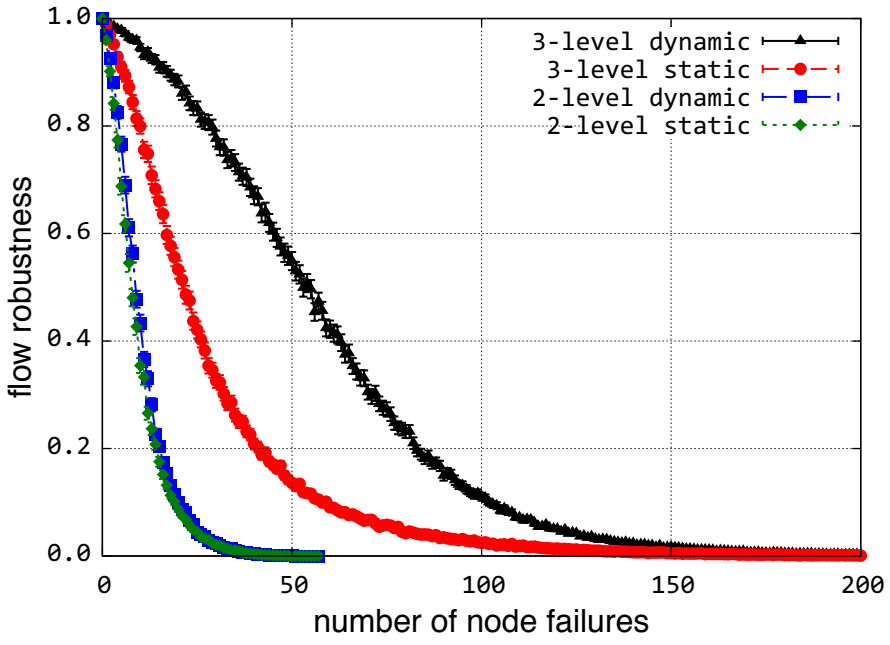

Figure 4.17: Robustness of multilevel network for node deletions

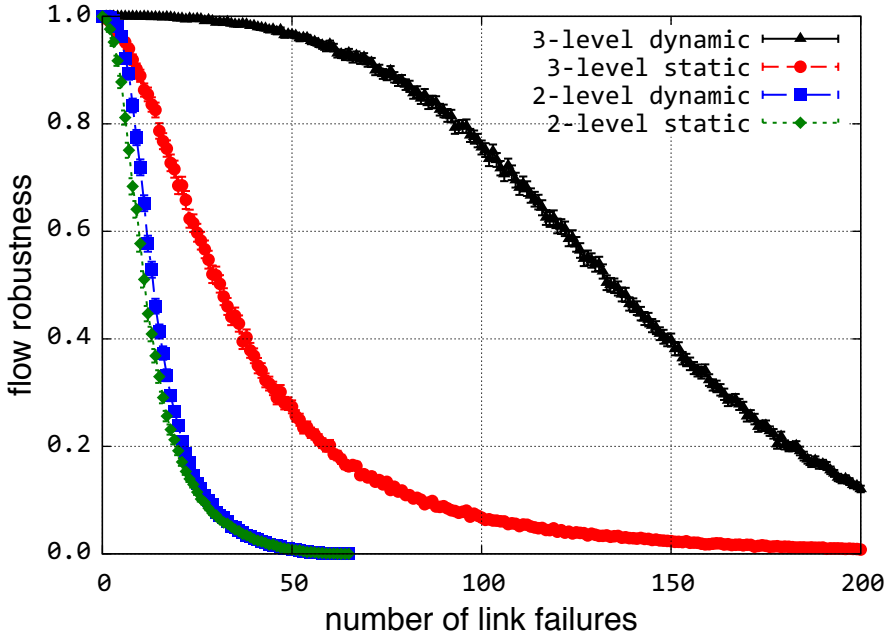

Figure 4.18: Robustness of multilevel network for link deletions

has higher values of average flow robustness for any given number of deletions than the 2-level network. For example in Figure 4.17, when we delete 50 random nodes in the

lowest topology of the 3-level graph (in the freeways graph), the flow robustness at the top level is approximately 0.55, whereas in a 2-level graph when we delete random 50 nodes in the lowest topology (in the physical topology), the flow robustness at the top level is approximately 0. This shows that adding multiple levels of graphs in resilience analysis impacts the outcome significantly. The difference when considering multiple levels is due to the fact that the bottom level graph has *nodes* that are a superset of the top 2 levels. We also note that if the US freeway topology was less connected (e.g. instead of a grid-like, it was linear) then the flow robustness would be lower. However, it is outside scope of this work to analyse different connected graphs at the lower layers, and it will be part of our future work. Moreover, both the 3-level and 2-level network have higher values of average flow robustness under dynamic routing than under static routing. Finally as expected, average flow robustness diminishes more severely with node deletions than with link deletions since a single node deletion results in the deletion of all of its incident links.

Our framework can handle graphs with any number of levels. Part of the reason behind the experiment given above was to demonstrate the ability of our framework to handle multilevel graphs with more than two levels, in particular, the 3-level graph with the Internet2 physical and logical topologies in the two upper levels and the freeway right-of-way graph in the lowest level. We focus on 2-level communication networks for the rest of our multilevel analysis. To that end, we use the geographical physical and logical level adjacency matrices for each of AT&T, Level 3, Sprint, TeliaSonera, and Internet2 to create multilevel graphs for each network, and then perform node and link deletions within each multilevel graph at the physical level. Finally, we calculate the resulting flow robustness in the logical level for every failure set. The results of the experiments involving node deletions for Sprint are shown in Figures 4.19 through 4.22, while the results of link deletions for Sprint are shown in Figures 4.23 through 4.24. Plots showing

the flow robustness results for all 6 topologies are presented in Appendix A.

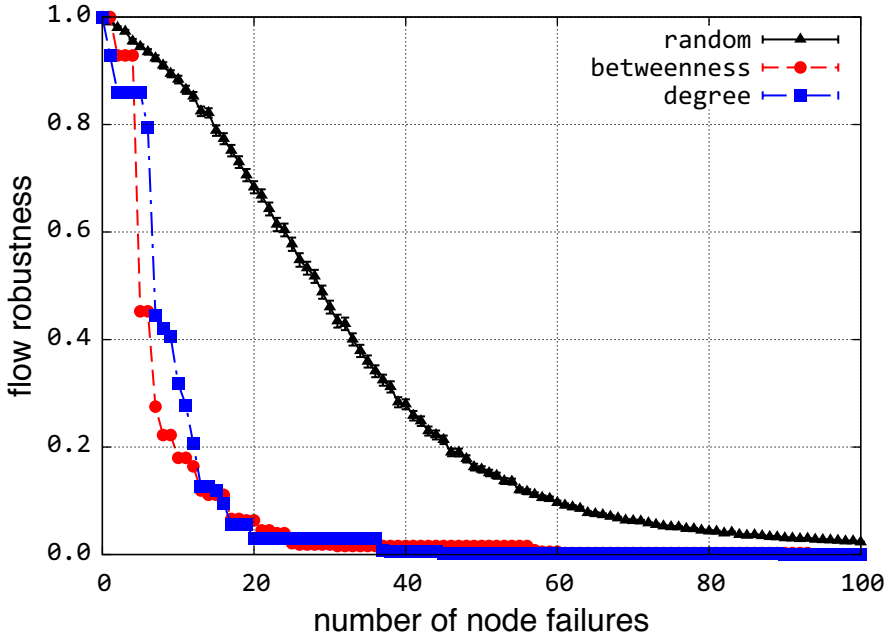

Figure 4.19: Robustness for dynamic routing during adaptive node deletions

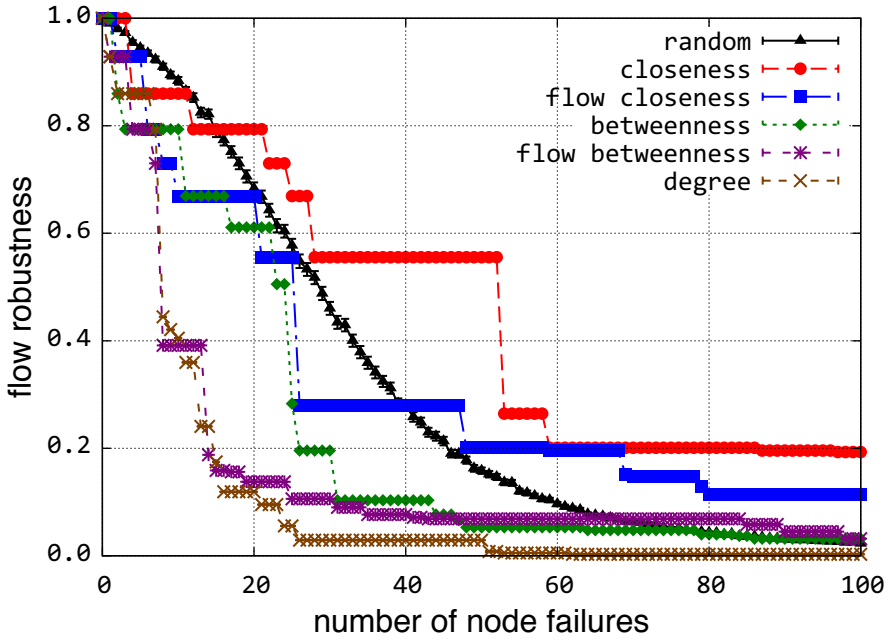

Figure 4.20: Robustness for dynamic routing during non-adaptive node deletions

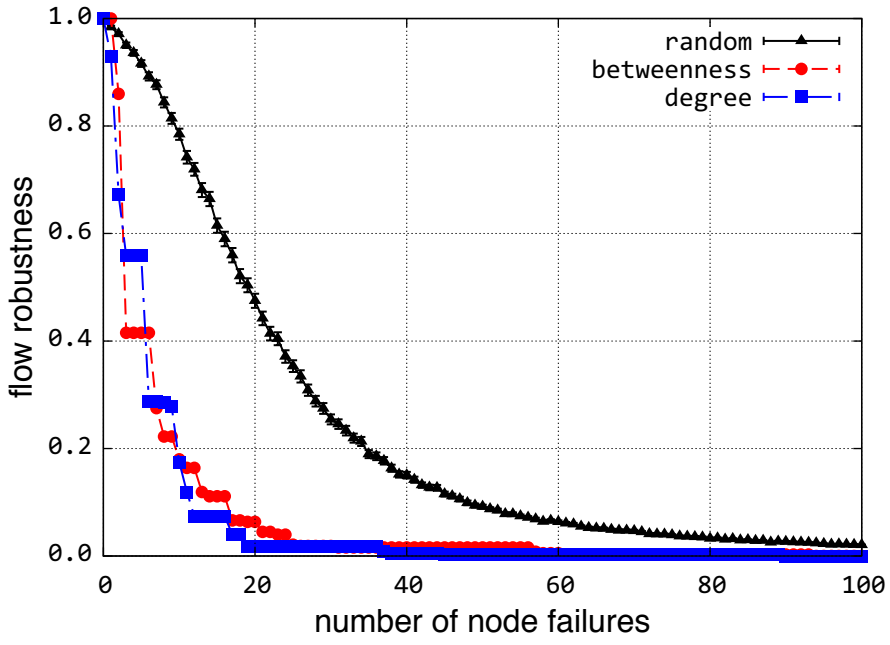

Figure 4.21: Robustness for static routing during adaptive node deletions

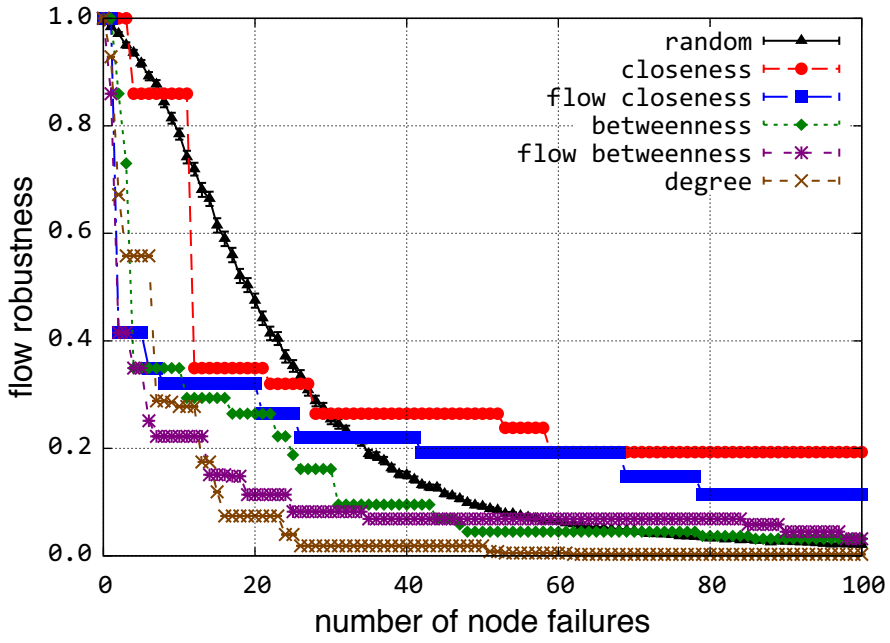

Figure 4.22: Robustness for static routing during non-adaptive node deletions

In some of cases we delete nodes and links at random while in others we delete nodes and links with very specific properties. The former experiments serve as a baseline for

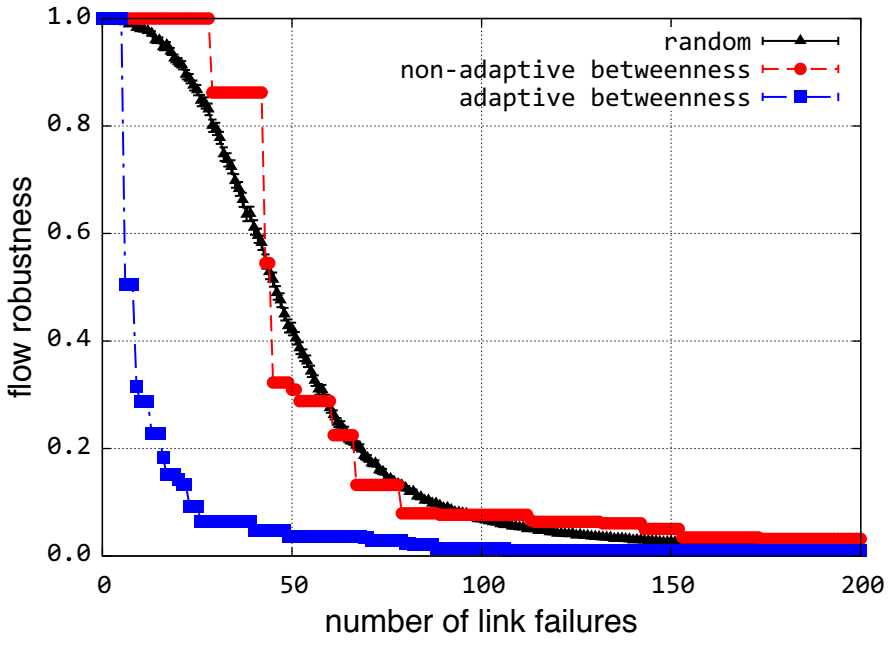

Figure 4.23: Robustness for dynamic routing link deletions

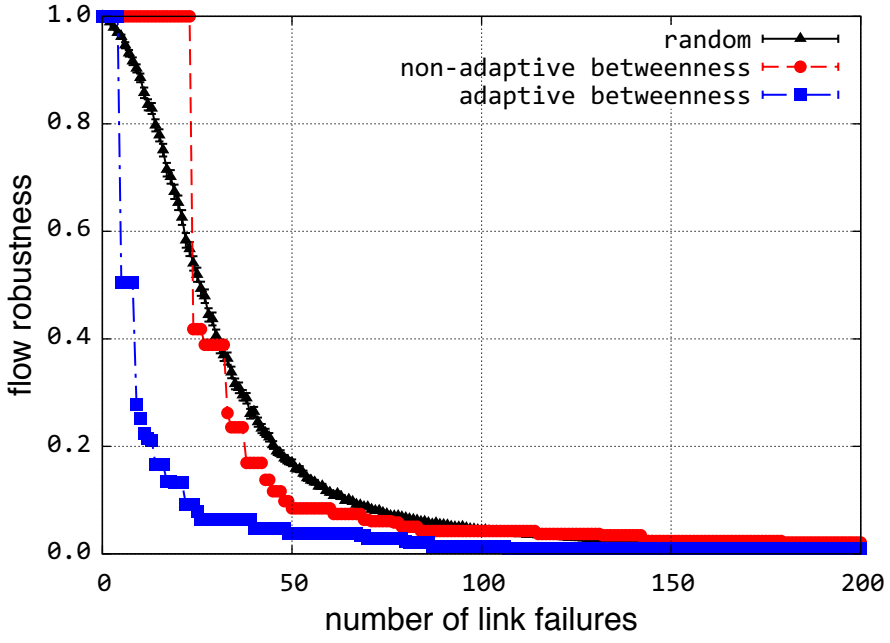

Figure 4.24: Robustness for static routing during link deletions

comparison against the latter, which focus on those nodes and links with large values of certain forms of centrality – in particular, betweenness, closeness, degree, link between-

ness, current-flow betweenness, and current-flow closeness. We discussed the first four metrics earlier in Section 4.1.3; here, we define current-flow betweenness and current-flow closeness [531].

Unlike conventional betweenness and closeness that measure a node's centrality based on the shortest paths going through that node, current-flow betweenness and current-flow closeness are both ways of measuring a node's centrality based on information flow alone. To understand these two measures, we must first view the graph under consideration as an electrical network into which one unit of current enters from a node known as the *source* and from which one unit of current exits through another node known as the *sink*[9]. The locations of the source and sink suffice to specify a unique current for each link in the network, as argued in Lemma 1 of [531]. Moreover, once each link is assigned a current, it is possible to assign absolute potentials to each node throughout the network, as argued in Lemma 2 of [531][10].

The current-flow betweenness of a node in a graph is simply the average of the total current passing through that node (from all of its incident links) over all possible electrical networks resulting from different possible (source, sink) pairs. The current-flow closeness of a node in a graph is the inverse of the average over all other possible nodes of the potential difference between that node when it is treated as the source and the other node when it is treated as the sink. If we view "current" as information, then in essence, current-flow betweenness is a measure of the amount of information that can pass through a given node, while current-flow closeness is a measure of the ease with which information can be sent out from one node into the rest of the network.

[9]Note that the concept of an electrical network – and therefore the measures of current-flow closeness and current-flow betweenness – make sense only if the graph is simple and connected. That is why these measures, along with closeness, are employed only for *non-adaptive deletions*, explained in the subsequent paragraph.

[10]In order to compute these potentials, we assign each link one unit of resistance. In other words, we employ the standard practice of assigning each link of an unweighted graph a length of one. This provides the ability to capture the link capacity in future analysis.

We use all of these measures (betweenness, closeness, degree, link betweenness, current-flow betweenness, and current-flow closeness) as a means to study what sorts of deletions at the physical level have the most disruptive effect at the logical level. Furthermore, we consider two different categories of deletions: adaptive deletions and non-adaptive deletions. A non-adaptive deletion is defined as a deletion performed based on the initial node or link centrality rankings that existed prior to the occurrence of any deletion. An adaptive deletion is defined as a deletion performed based on centrality rankings that are recomputed after the most recent deletion. This can result from an attacker that has real-time access to internal network management and operations information.

Finally, note that for centrality-based deletions we compute flow robustness, while for random deletions we compute *average* flow robustness in the same manner as before, that is by averaging the flow robustness over 1000 failure sets, each of which was generated by performing the number of random deletions. We also plot the 95% confidence intervals on each of the points located on the random curves.

As before, flow robustness diminishes more severely under static routing than under dynamic routing, and node deletions have a greater impact on flow robustness than link deletions. Furthermore, adaptive deletions have a more severe impact on the network than non-adaptive deletions. The reason for this should be clear: an adaptive deletion is always selecting from the pool of existing nodes or links the one with the highest centrality value, whereas a non-adaptive deletion will select from the pool of one that used to – but may no longer – have the highest centrality value. Hence, adaptive deletions have a far greater tendency to select the most important nodes or links than non-adaptive deletions, which results in a more severe impact on the flow robustness of the logical level.

Given a sufficiently small number of deletions, random deletions tend to have less effect on flow robustness than any other type of deletion. This is unsurprising, since deletions

based on centrality metrics have a greater tendency to delete more "important" nodes and links than random deletions. What is surprising, however, is that, given a sufficient number of deletions, the flow robustness resulting from non-adaptive deletions based on closeness and current-flow closeness surpasses the average flow robustness resulting from random node deletions. This holds true for all five of the networks under study. For example in Figure 4.20, with 40 random node deletions the flow robustness of the Sprint network is about 0.3, whereas the flow robustness for closeness is about 0.55. Similarly in Figure 4.20, for 60 random node deletions the flow robustness is about 0.1 and for flow closeness the flow robustness is about 0.2. We speculate that since these are *non-adaptive* challenges, by the time network arrives in a state in which several nodes are deleted, initially calculated rankings are no longer accurate. However, why this happens *only* for closeness and current-flow closeness centrality metrics is not known. The reasons for the occurrence of this phenomenon, and an investigation into the types of multilevel graphs to which it is restricted, will be the subject of future work.

4.3.3 Multiprovider Graph Model

We introduce a new graph-theoretic model in which we define the concept of a *multi-provider graph*. Within our framework, a multiprovider graph is an ordered pair $(G_{\mathrm{L3}}, G_{\mathrm{AS}})$, in which L3 represents PoP-level topology and AS represents the interprovider AS topology, where $G_{\mathrm{L3}} = (V_{\mathrm{L3}}, E_{\mathrm{L3}})$ and $G_{\mathrm{AS}} = (V_{\mathrm{AS}}, E_{\mathrm{AS}})$ are graphs such that:

1. the vertices in V_{AS} are mutually disjoint connected subgraphs of G_{L3} that, when taken together, contain all of the vertices in V_{L3}. More specifically, if $V_{\mathrm{AS}} = \{v_1, v_2, ..., v_n\}$, then

   (a) any two distinct vertices $v_i, v_j \in V_{\mathrm{AS}}$ will be connected subgraphs of G_{L3} given by $v_i = (V_i, E_i)$ and $v_j = (V_j, E_j)$ such that $V_i \cap V_j = \emptyset$

(b) if we let $v_i = (V_i, E_i)$ for all integers i such that $1 \leq i \leq n$, then $\bigcup_{i=1}^{n} V_i = V_{\text{L3}}$.

2. there exists some function $f : E_{\text{AS}} \to 2^{E_{\text{L3}}}$ such that for any pair of distinct vertices $v_i, v_j \in V_{\text{AS}}$ given by $v_i = (V_i, E_i)$ and $v_j = (V_j, E_j)$, if $\{v_i, v_j\} \in E_{\text{AS}}$, then $f(\{v_i, v_j\}) = V_{ij} \cap E_{\text{L3}}$ where V_{ij} is the set of unordered pairs $\{u_i, u_j\}$ such that $\{u_i, u_j\} \in V_{ij}$ if and only if $u_i \in V_i$ and $u_j \in V_j$. More explicitly, the mapping f is used to identify edges between specific AS peer routers that serve to connect two ASes $v_i, v_j \in V_{\text{AS}}$ that share a given AS-edge $\{v_i, v_j\} \in E_{\text{AS}}$.

To study multiprovider graphs, first we combine the PoP-level topologies of four commercial ISPs (AT&T, Level 3, Sprint, TeliaSonera). We treat each ISP as a single AS, and the resulting AS-level abstract graph is a full-mesh with 4 nodes, in which each AS is connected to the other through a logical IXP (Internet exchange point) link. We select Atlanta NAP [532], Equinix [533], Terremark [534], and MAE-East [535] as the IXPs in which 4 ISPs are connected. The reason we select these 4 IXPs is that we analysed a number of IXP websites and found that these IXPs do provide service to the 4 commercial ISPs. We do not claim that this is an exhaustive list of IXPs, however, it was sufficient to generate a full-mesh AS-level graph for those tier-1 ISP providers. The 4 IXPs are distributed across the US in 17 different cities and there are 51 logical links that connected the four ISPs.

In Figure 4.25, the flow robustness of a multiprovider graph is shown. In this case we delete all inter-AS IXP links in a *city*, ranked based on betweenness. As expected, adaptive attacks inflict more harm than non-adaptive attacks, which, in turn, inflict more harm than randomly-placed attacks. The sharp reductions of flow robustness due to targeted attacks indicate the disconnection of an AS from the AS-level graph following such attacks. Note that several cities must be deleted in order to disconnect a single AS. In contrast, the flow robustness values in random scenarios decrease at a smoother rate

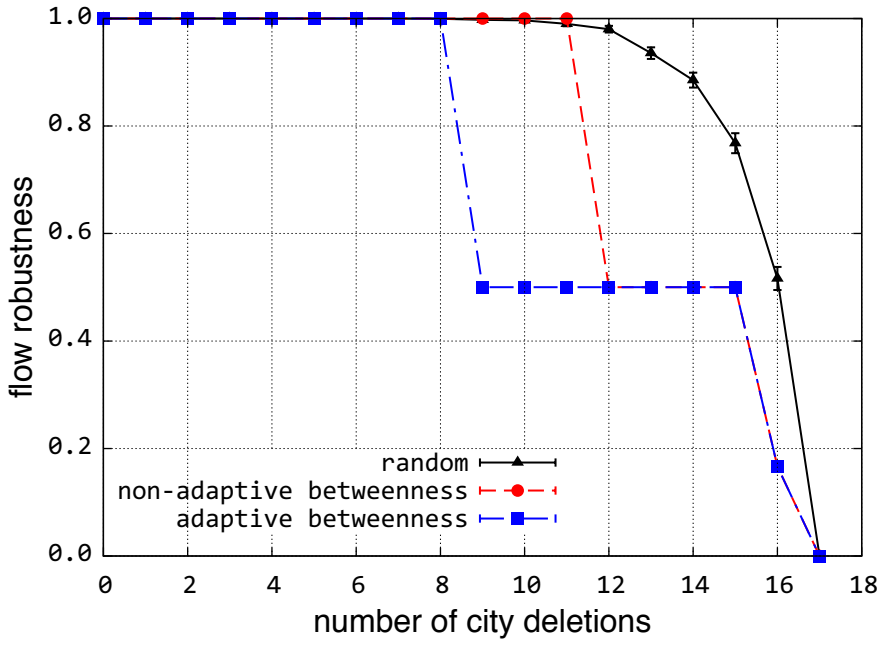

Figure 4.25: Robustness of multiprovider network

because the flow robustness is averaged over 1000 failure sets. For example, the flow robustness values indicate that a very high percentage of the failure sets following the twelfth city deletion *did not* partition the network in any manner. Furthermore, our results indicate that it is very difficult to partition the tier-1 ISP connectivity, which is a full-mesh, given that it requires at least 9 cities and all the IXP links in a city to be destroyed. If we had included all IXPs in more than 17 cities, intuitively it would have been even more difficult to partition the AS-level graph.

Next, we analyse flow robustness of two provider graphs and their connections via IXPs. In this case, we investigate the flow robustness of provider duos as shown in Figure 4.26. Since these graphs are constructed by only two providers, the impact of random deletions and centrality-based attacks result in the same flow robustness. The connectivity between these provider duos breaks when the nth IXP link is broken. For example, AT&T and Level 3 peer with each other in 16 cities, and flow robustness remains at 1 until the 16th

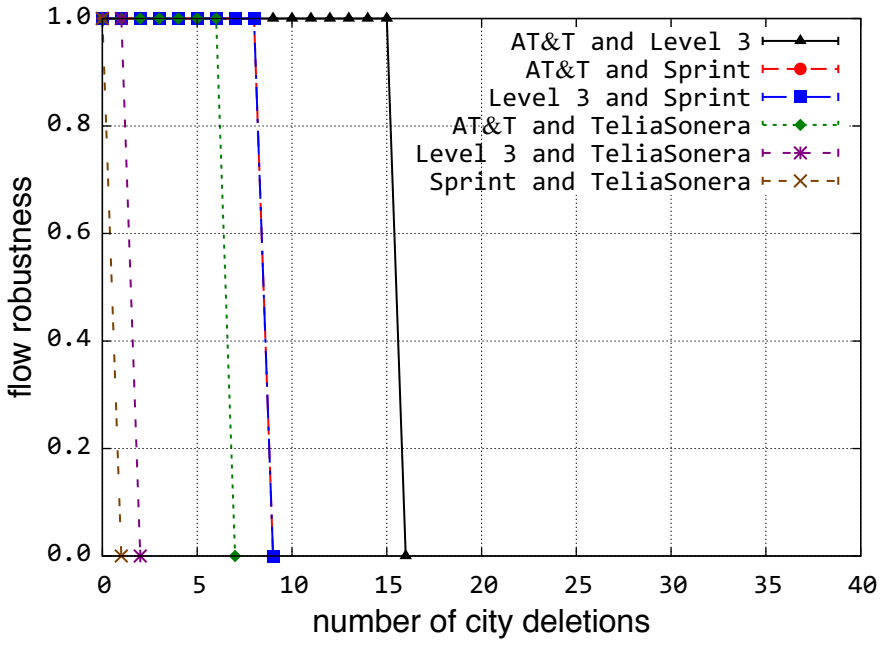

Figure 4.26: Robustness of provider duos

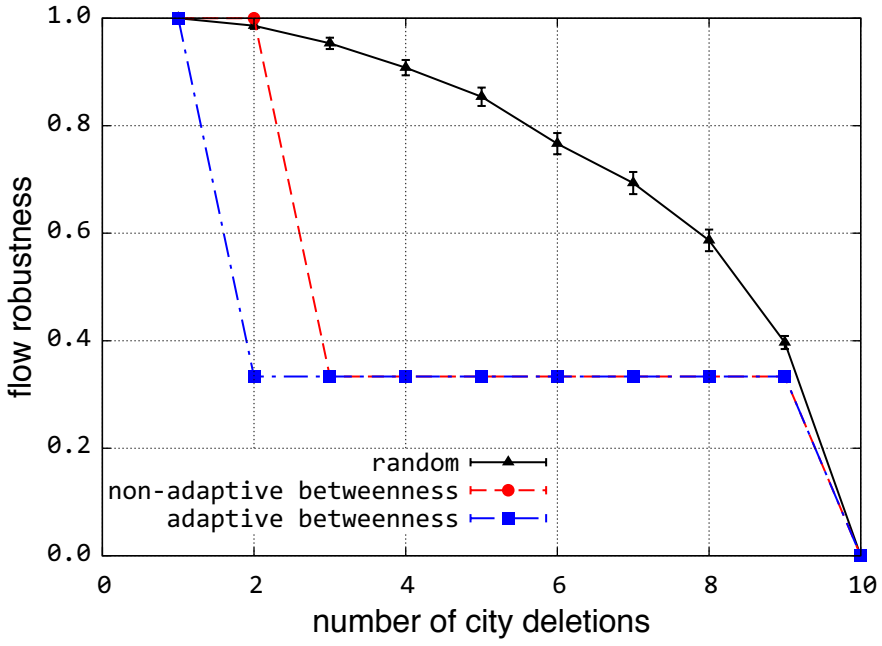

Figure 4.27: Robustness of a provider trio (Level 3, Sprint, TeliaSonera)

IXP link breaks, in which flow robustness drops to 0.

Finally, we analyse the flow robustness of a provider trio combination (Level 3, Sprint, TeliaSonera). The network performance when challenged by random failures and targeted attacks against IXP links are shown in Figure 4.27. The number of IXP connections is reduced from 17 to 10 when we consider only three providers. As expected, the flow robustness depends on the number of IXP links between the provider trio.

4.3.4 Multilevel and Multiprovider Analysis

In this section we present the analysis of multilevel and multiprovider networks. We use the multiprovider graph of 4 service providers explained in Section 4.3.3. Additionally, we combine the *structural physical-level* topologies of each provider using a similar methodology in which we connect each provider's physical topology in 17 cities where there is an exchange point.

2-Level Multiprovider Analysis

For our multiprovider analysis we first examine the 2-level multiprovider graph in which the lowest level is the combined logical level topology of the 4 providers and the topmost level is the E2E level topology of aggregated flows between 9 nodes. Since the combined multiprovider graph might have more than one node in a given city, we assign the given 9 nodes (Atlanta, GA; Chicago, IL; Dallas, TX; Denver, CO; Los Angeles, CA; New York, NY, San Jose, CA; Seattle, WA; Washington, DC) in the following manner. We assign Atlanta, GA, Chicago, IL, and Dallas, TX, to AT&T; Denver, CO and Los Angeles, CA, to Level 3; New York, NY and San Jose, CA, to Sprint; Seattle, WA and Washington, DC, to TeliaSonera. This assignment was purely arbitrary and was based on the alphabetical

order of the cities. Our objective is to have a same graph on top of the combined multiprovider graph.

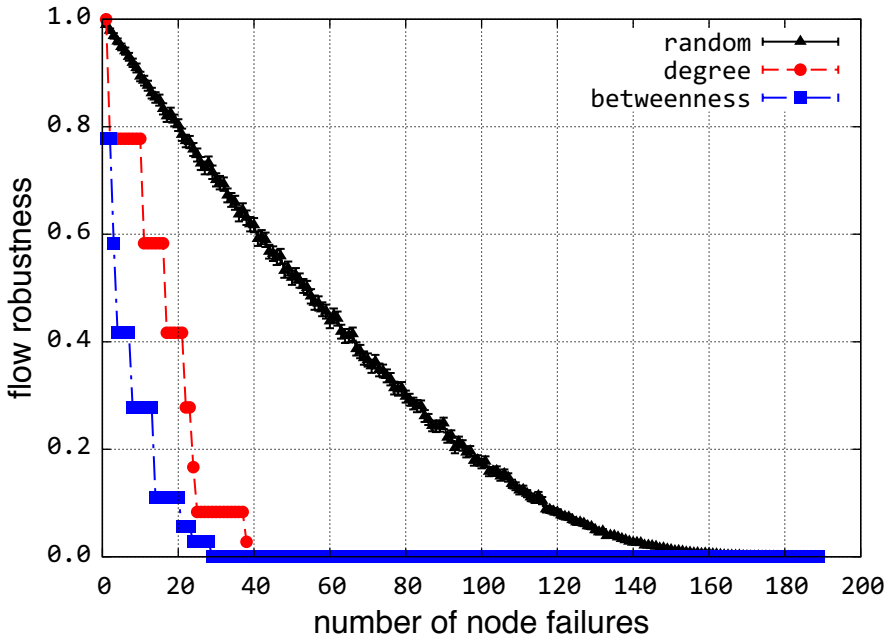

Figure 4.28: Robustness of 2-level multiprovider graph

The results for the 2-level multiprovider graphs are shown in Figure 4.28. Compared to single providers it takes more effort to achieve a flow robustness of 0 when considering multiprovider networks. For example, even if we consider the best case scenario of Level 3 in which it requires 12 nodes to be deleted for a flow robustness of 0, it takes 28 nodes to be deleted to have a flow robustness of 0 for the 2-level multiprovider graphs. It requires 38 nodes to be deleted when considering degree of connectivity attacks, whereas random failures resulting in a flow robustness of 0 requires 185 nodes to be deleted.

3-Level Multiprovider Analysis

Finally, we consider a 3-level multiprovider graph. In this case, the 3 levels from the lowest level to the highest level are given by the combined structural physical level graph,

combined multiprovider graph, and E2E level aggregated flows. The total number of nodes in the combined physical level is 320 and total number of links is 524.

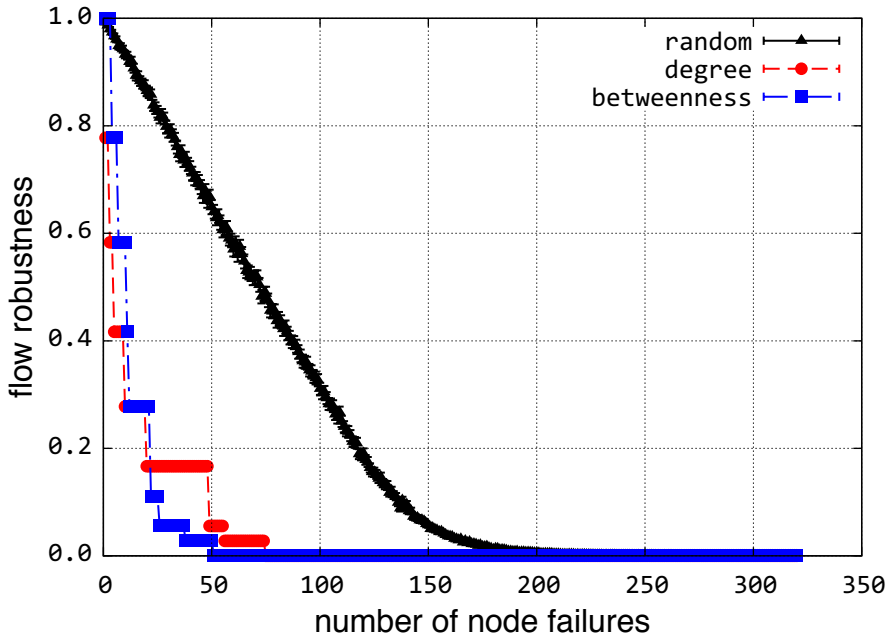

Figure 4.29: Robustness of 3-level multiprovider graph

The flow robustness resulting from the various challenges for the 3-level multiprovider graph is shown in Figure 4.29. The flow robustness falls to 0 when 290 out of 320 (90%) nodes are deleted in a random fashion, 51 out of 320 (16%) of the nodes are deleted based on betweenness centrality, and 74 out of 320 nodes (23%) are deleted based on degree of connectivity.

4.4 Physical Level Network Modelling

Physical level topologies are necessary to study the structure of the Internet realistically. They provide the means of connecting nodes in the higher levels [499] and they are needed to model area-based challenges on networks [23]. In an effort to maintain security and

competitiveness, service providers generally do not share fibre topology of their networks. The applicability of existing geographic graph generators for modelling physical-level networks is lacking in the literature and are an important area of research [34–36].

In this section, we describe the synthetic geographical graph models and provide structural- and cost-based comparisons of geographic graph models applied to graphs with node locations that are constrained to those of actual physical graphs. Furthermore, we discuss how one might develop a better alternative geographical graph model to capture a graph's structural properties.

4.4.1 Network Cost Model

Structural properties impact the connectivity and *cost* of building networks. While at the logical level the cost is captured by the number of nodes and the capacity of each node (i.e. bandwidth and number of ports available in a router [497,498]), at the physical level, the length of the fibre dominates the cost. Previously, we provided a network cost model as:

$$C_{i,j} = f + v \times d_{i,j} \tag{4.1}$$

where f is the fixed cost associated with link (including termination), v is the variable cost per unit distance for the link, and $d_{i,j}$ is the length of a link [5, 536, 537]. Based on the assumption that the variable cost dominates in long haul fibre networks, we ignore the fixed cost associated with links, and simplify network cost as:

$$C = \sum_{i} l_i \tag{4.2}$$

where l_i is the length of the i-th link [76, 500, 538]. We calculate the total link length for each provider with this simplified network cost model as shown in 5th column in Table 4.5. We note that, the detailed graph properties of these graphs are presented in Section 4.1, Table 4.2. The total link length of each physical topology is somewhere between 14,000 to 50,000 km. For these topologies, the smaller the size of the network, the smaller the total length link of the fibre.

Table 4.5: Cost of physical-level and full-mesh networks

Network	**Geographical**				**Full mesh**	
	Nodes	**Links**	**Avg. Node Degree**	**Tot.** l **[km]**	**Links**	**Tot.** l $\times 10^6$ **[km]**
AT&T	383	488	2.55	50,026	73,153	116.8
Level 3	99	130	2.63	28,538	4,851	7.5
Sprint	264	312	2.36	33,627	34,716	57.8
TeliaSonera	21	25	2.38	14,190	210	0.4
Internet2	57	65	2.28	19,050	1,596	2.7
CORONET	75	99	2.64	28,325	2,775	4.6

Next, for each physical level topology, we consider as an upper baseline the full-mesh topology whose vertex set is identical to that of the original topology. We then calculate the total link count and length of each full-mesh topology as shown in column 6 and 7 in Table 4.5, respectively. Note that the total link lengths are given in millions of km for a hypothetical full-mesh physical level topology, emphasising that real networks cannot have unlimited resilience due to cost constraints.

4.4.2 Structure of Physical-Level Graphs

In Section 4.1.3 we describe the distinction between geographical and structural physical-level graphs. In this section, we show an example to emphasise the difference in representing geographic and structural physical level graphs in Figure 4.30. In this case, the

map we have [505] has a path from Spokane, WA (dark coloured pin on upper left corner of Figure 4.30) to Billings, MT (cyan coloured pin on lower right corner of Figure 4.30) crossing five cities (Coeur d'Alene, ID; Thompson Falls, MT; Missoula, MT; Helena, MT; Bozeman, MT), forming a zigzag shaped path that captures geography of the path. This geographic physical graph is *necessary* to *accurately* study area-based challenges such as severe weather. On the other hand, this physical path from Spokane, WA to Billings, MT can be represented as a single link *structurally*. After all, it does not matter where the link is cut between Spokane, WA and Billings, MT since there is no logical PoP that is a traffic source or sink between these cities. We note that, the nodes that have a degree 3 or higher are kept to capture the *physical layer structure* even if there is not a logical PoP in these locations. Moreover, stub nodes are removed since we are interested in *backbone networks* and in the geographic representation these stub nodes might as well represent access networks.

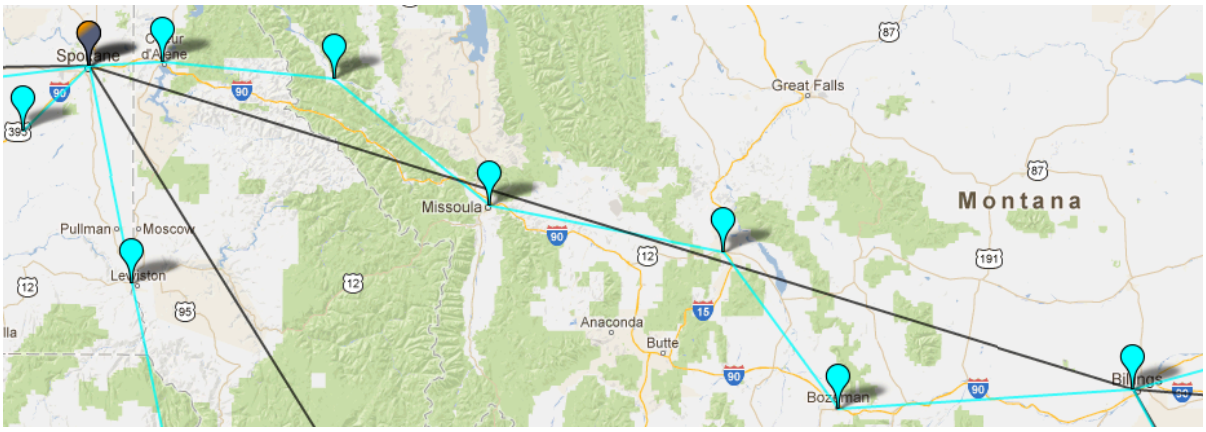

Figure 4.30: Geographical vs. structural graphs

The number of nodes, links, and average degree of the *structural* graphs are shown in Table 4.6. The topological characteristics of structural graphs are detailed in Section 4.1.3, Table 4.3. Each structural graph has fewer nodes and links than its corresponding physical level graph. However, we note that the total fibre length of the structural graphs (cf.

Table 4.6) is close to that of the geographical physical graphs (cf. Table 4.5).

Table 4.6: Cost of structural graphs

Network	**Nodes**	**Links**	**Avg. Deg.**	**Tot. l [km]**
AT&T	162	244	3.01	40,985
Level 3	63	94	2.98	27,597
Sprint	77	114	2.96	28,069
TeliaSonera	18	21	2.33	14,040
Internet2	16	24	3.00	18,146
CORONET	39	63	3.23	27,579

4.4.3 Synthetic Graph Models for Physical-Level Networks

In this section we present four different geographic graph models. The Gabriel graph model is a parameterless model that uses only node locations as input, while the geometric, geographical threshold, and Waxman models all require at least one parameter. The geometric graph model uses a single threshold parameter, while the geographic threshold model and the probabilistic Waxman model use two parameters. We apply each of these graph models to graphs with node locations constrained to those of actual physical topologies. Given the diverse nature of these models, we believe the following sections represent a fairly comprehensive analysis of geographic graph models applied to physical topologies.

Gabriel Graphs

Next, we generate Gabriel graphs of the six service provider networks. Gabriel graphs are useful in modelling graphs with geographic connectivity that resemble grids [539,540]. We would expect the Gabriel graph to be one of the best ways to model physical topologies for this reason. In a Gabriel graph, two nodes are connected directly if and only if there

are no other nodes that fall inside the circle whose diameter is given by the line segment joining the two nodes. The number of links and the total link length of Gabriel graphs of six networks are shown in Table 4.7.

Table 4.7: Cost of Gabriel graphs

Network	**Links**	**Tot. l [km]**
AT&T	686	66,157
Level 3	170	33,991
Sprint	474	57,104
TeliaSonera	26	12,111
Internet2	94	27,786
CORONET	127	33,265

Geometric Graphs

A 2-dimensional geometric graph is a graph in which nodes are placed on a plane or surface and any pair of nodes is connected if and only if:

$$d(u, v) \leq d_\theta \tag{4.3}$$

where $d(u, v)$ is the Euclidean distance between the two nodes $\{u, v\}$, and d_θ is a distance threshold parameter [541]. In the conventional random 2-dimensional geometric graph model, nodes are distributed randomly on a plane.

Using the physical level node locations of six provider networks, we generate four different geometric graphs based on four different d_θ distance threshold values. For the first set of graphs, we use the maximum link length of the actual physical graph as the d_θ value. For the second set of graphs we select the largest possible values of d_θ such that the total link lengths of these graphs are less than the total link lengths of the original physical level

graphs. Using this methodology, we find that all of the synthetically generated graphs are disconnected. For the third set of graphs, we select the smallest value of d_θ such that the graphs are connected. It turns out that none of these graphs are biconnected. For the fourth set of graphs we select the smallest values of d_θ such that the graphs are biconnected: that is, such that the graphs will remain connected after the failure of any one node or link. This is a basic requirement for basic network resilience and survivability [2, 53]. The link lengths l of the actual graphs as well as the synthetically generated geometric graphs are shown in Table 4.8.

Table 4.8: Cost of geometric graphs based on a threshold value

Network	Actual			Threshold Optimised			Cost Optimised			Cost & Con. Optimised			Cost & Bicon. Optimised		
	Links	Tot. l [km]	Max. l [km]	Links	Tot. l $\times 10^3$ [km]	d_θ [km]	Links	Tot. l [km]	d_θ [km]	Links	Tot. l [km]	d_θ [km]	Links	Tot. l $\times 10^3$ [km]	d_θ [km]
AT&T	488	50,026	629	15,062	5,719	629	783	49,937	99	4,916	918,353	302	8,343	2,169	424
Level 3	130	28,538	1,063	2,107	1,326	1,063	209	28,358	226	749	234,721	528	1,104	449	683
Sprint	312	33,627	602	6,478	2,328	602	466	33,573	112	3,417	804,197	390	4,261	1,159	452
TeliaSonera	25	14,190	1,592	106	88	1,592	37	13,757	614	56	27,842	859	93	68	1,425
Internet2	65	19,049	910	442	246	910	83	18,997	334	131	37,532	424	258	105	616
CORONET	99	28,325	943	922	506	943	156	28,144	280	512	188,663	604	613	254	691

To further explain the data in Table 4.8, consider the AT&T physical graph with the given node locations. The number of links, total link length, and maximum link length of the actual AT&T physical graph are shown in columns 2, 3, and 4, respectively. For the case of AT&T, when we assign $d_\theta = \max(l_i)$ (where $\max(l_i) = 629$ km in this case), the synthetically generated geometric graph has 15,062 links and the total length of the graph is approximately 5.7×10^6 km. Using this *threshold optimised* methodology we obtain the number of links, total link length, and d_θ as shown in columns 5, 6, and 7, respectively. With the second *cost optimised* methodology we generate synthetic geometric graphs such that the total link length is less than that of the actual physical topology. In the case of AT&T, the generated graph has a total link length of 49,937 km, which is less than that of the actual AT&T graph whose total link length is 50,026 km. We note that the cost optimised geometric graphs of all service providers are disconnected graphs. The

number of links, total link length, and d_θ for cost-optimised graphs are shown in columns 8, 9, and 10, respectively. Since the cost-optimised geometric graphs are disconnected graphs, we increase the value of d_θ until we obtain connected graphs. Applying this *cost and connectivity optimised* methodology to AT&T, the total number of links is 4,916, the total length of the links is 918,353 km, and $d_\theta = 302$ km, as shown in columns 11, 12, and 13, respectively. While cost and connectivity optimised graphs are connected, none of them are biconnected. Therefore, we increase d_θ so that the resulting geometric graphs are biconnected. Applying this *cost and biconnectivity optimised* methodology to AT&T, we obtain a synthetically generated geometric graph with 8,343 links, 2.2×10^6 km of total link length, and a d_θ value of 424 km, as shown in columns 14, 15, and 16, respectively. The rest of the service provider data is shown in the consecutive rows in Table 4.8.

Population-weighted Geographical Threshold Graphs

A threshold graph is a type of graph in which links are formed based on node weights [542]. Two nodes $\{u, v\}$ with node weights $\{w_u, w_v\}$ are connected if and only if:

$$w_u + w_v \geq t \tag{4.4}$$

in which t is a threshold value that is a non-negative real number. A modified version of a threshold graph is a *geographical* threshold graph that includes geometric information about the nodes [543]. In this case, two nodes $\{u, v\}$ with node weights $\{w_u, w_v\}$ are connected if and only if:

$$w_u + w_v \geq \psi d(u, v)^\phi \tag{4.5}$$

where ψ and ϕ are model parameters and $d(u, v)$ is the distance between nodes $\{u, v\}$. We assign the node weights to be the population estimates of cities for year 2011, which are taken from the US Census Bureau [544]. The population statistics for each provider are given in Table 4.9. For the AT&T physical graph, the total of population of all of the cities (e.g. 383 cities) is about 76 million, and the average city population is about 197,000. The most populous city (NYC for all networks) has about 8.2 million people, and the least populated city has 182 people. These statistics are shown in columns 2, 3, 4, and 5 in Table 4.9 respectively for each provider network.

Table 4.9: Population statistics of cities as node weights

Network	**Total**	**Average**	**Maximum**	**Minimum**
AT&T	75,753,034	197,789	8,244,910	182
Level 3	53,221,035	537,586	8,244,910	12,695
Sprint	67,794,208	256,796	8,244,910	448
TeliaSonera	27,944,279	1,330,680	8,244,910	65,397
Internet2	40,980,611	718,958	8,244,910	8,438
CORONET	49,559,726	660,796	8,244,910	33,395

Using city populations as node weights, we generate synthetic graphs for each provider network. We choose $\phi = 1$ so that we can manipulate only ψ. Moreover, by choosing $\phi = 1$, we find that the righthand side of inequality (4.5) varies linearly with distance. Hence, as the distance increases between two nodes they are less likely to be connected. Having fixed $\phi = 1$, we first choose ψ so as to minimise cost while ensuring connectivity, and then choose ψ so as to minimise cost while ensuring biconnectivity. More specifically, for each network, we select the largest value of ψ rounded to the nearest tenth such that the graph is connected, and then select the largest value of ψ rounded to the nearest tenth such that the graph is biconnected. The results of both methodologies are shown in Table 4.10. For AT&T, we find that the largest value of ψ such that AT&T is connected is 3.1, yielding a link number of 1670 and a total link length of 690,941 km. Additionally,

the largest value of ψ such that AT&T is biconnected is 2.4, which yields a link number of 2,336 and a total link length of 1,036,747 km.

Table 4.10: Cost of population-weighted geographic threshold graphs for $\phi = 1$

Network	**Connectivity Optimised**			**Biconnectivity Optimised**		
	ψ	**Links**	**Tot. l [km]**	ψ	**Links**	**Tot. l [km]**
AT&T	3.1	1,670	690,941	2.4	2,336	1,036,747
Level 3	3.4	324	158,316	2.4	526	304,696
Sprint	3.0	1,164	500,678	2.4	1,532	717,311
TeliaSonera	3.4	43	31,099	2.3	62	58,492
Internet2	3.2	151	98,733	2.3	233	194,938
CORONET	3.3	244	127,387	2.4	374	233,360

Location-constrained Waxman Graphs

The Waxman model provides a probabilistic way of connecting nodes in a graph [545]. Given two nodes $\{u, v\}$ with a Euclidean distance $d(u, v)$ between them, the probability of connecting these two nodes is:

$$P(u, v) = \beta e^{\frac{-d(u,v)}{L\alpha}} \tag{4.6}$$

where $\beta, \alpha \in (0, 1]$ and L is the maximum distance between any two nodes. Increasing β increases the link density and a large value of α corresponds to a high ratio of long links to short links.

In the Waxman model nodes are uniformly distributed in the plane. We modify the Waxman model so that it is constrained by the node locations. The resulting link properties of the location-constrained Waxman model, along with the β and α parameters, are shown in Table 4.11. For each network, we choose β and α such that the resulting graph is a connected graph with the smallest possible total link length. For example, in

the AT&T graph, using the node geographic locations we use β and α values of 0.1 and run the experiments 10 times, which results graphs that are disconnected. Then, we keep β at a value of 0.1 and increase α to a value of 0.2, which results in connected graphs but with a mean of 1.6 million km total link length. We calculate total link length by averaging 10 runs with increments of 0.1 for β and α parameters until we find connected graphs that result in least total length. The β and α parameters for each provider are shown in columns 2 and 3 in Table 4.11. The average number of links for each topology resulting from 10 runs is shown in column 4, whereas the standard deviation σ of the number of links resulting from 10 runs is shown in column 5. The average total link length of 10 runs is shown in column 6, and the standard deviation σ of the total link of length resulting from 10 runs is shown in column 7.

Table 4.11: Cost of location-constrained Waxman graphs

Network	$\boldsymbol{\beta}$	$\boldsymbol{\alpha}$	**Avg. No. of Links**	σ **Links**	**Avg. Tot.** $\boldsymbol{l}$ **[km]**	σ **Tot.** $\boldsymbol{l}$
AT&T	0.2	0.1	1,981	54	1,044,856	29,509
Level 3	0.6	0.1	392	14	205,036	7,896
Sprint	0.2	0.1	904	43	475,943	24,271
TeliaSonera	0.6	0.2	31	3	24,498	4,743
Internet2	0.6	0.1	102	10	62,100	7,723
CORONET	0.5	0.1	174	15	91,002	10,062

4.4.4 Analysis of Physical-Level Graphs

In this section we present the cost incurred using different graph models, as well as visually present the structure of the synthetic models for the Internet2 network.

Cost Analysis of Graphs

We presented the total link lengths of the synthetically generated graphs in the previous section. However, in order to see the big picture we summarise them again in Figure 4.31. The y-axis shows the cost incurred in terms of total link length in units of m for each graph and x-axis shows six provider networks for different graph models. We use the graphs that provide minimal connectivity with the least cost. For the Waxman graph (as discussed in Section 4.4.3), among the set of ten connected graphs we generated, we choose the graph with the smallest total link length to present in Figure 4.31.

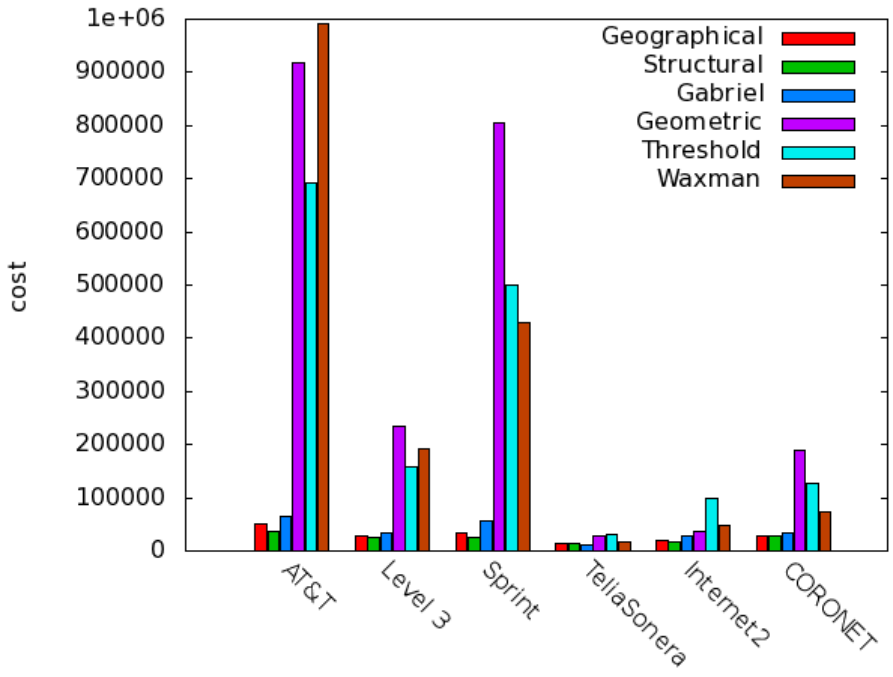

Figure 4.31: Cost analysis of physical graph models

Cost analyses of the graphs indicate that the cost of each synthetically generated graph depends on the order of the network. For example, the cost incurred for TeliaSonera is the smallest and TeliaSonera also has the lowest number of nodes. Second, we can infer that geographical, structural, and Gabriel graphs incur about the same cost for all providers.

The cost of geometric, population-weighted geographical threshold, and Waxman graphs are higher than the previous three models. However, the cost difference between different graph models for TeliaSonera is not as drastic as larger size networks due to its smaller order. In other words, the difference between the first three and last three graph models differs more as the number of nodes increase. The location-constrained Waxman model is probabilistic in nature and the cost values are shown for a sample generated graph using this model with $\beta = 0.6$ and $\alpha = 0.1$. The cost incurred with the Waxman model is generally higher than that of the original geographic physical level graphs across all providers.

A graph's connectivity of can be improved by adding links; however this adds additional cost to achieve resilience. By examining the synthetically generated topologies using the geometric graph model and geographical threshold graph model in Tables 4.8 and 4.10 respectively, we observe that it incurs about 90% or more additional cost to result in biconnected graphs. For example, applying the geometric graph model on Internet2 topology yields a total link length of about 37,000 km for a minimal connected graph. However, for the same node locations of Internet2, when we generate a biconnected synthetic graph, the total link length is about 105,000 km, which is more than double the cost of the uniconnected version. Similar conclusions can be also observed for the geographic threshold model. When we compare these cost values against the upper bound of the Internet2 graph, which is 2.7 million km, we observe that they are far less than the upper bound. From these results, we conclude that all synthetic graph models discussed in this work—with the exception of the Gabriel graph model—result in a total link length that is not feasible to model physical level topologies.

Visual Analysis of Graphs

We inspect all the synthetically generated topologies of all the providers using KU-TopView (KU Topology Map Viewer) [5, 51]. We find the results to be similar across all providers. We discuss the Internet2 graph here because its smaller order makes it easier to visualise, and thus more informative for demonstrating the fitness of each synthetic graph model on this topology. The geographical, structural, Gabriel, geometric, population weighted geographical threshold, and location-constrained Waxman model of the Internet2 physical-level graphs are shown from Figure 4.32 through Figure 4.37.

Figure 4.32: Visual representation of Internet2 geographical topology

The geographic physical-level Internet2 topology with 57 nodes and 65 links is shown in Figure 4.32. In Section 4.2 we showed using graph spectra that geographic physical-level graphs resemble a grid-like structure [75]. The structural physical-level Internet2 topology in which degree-2 nodes are removed is shown in Figure 4.33. The synthetically generated Gabriel graph of the geographic Internet2 graph is shown in Figure 4.34. While the Gabriel graph preserves the grid-like structure of the geographic physical-level topol-

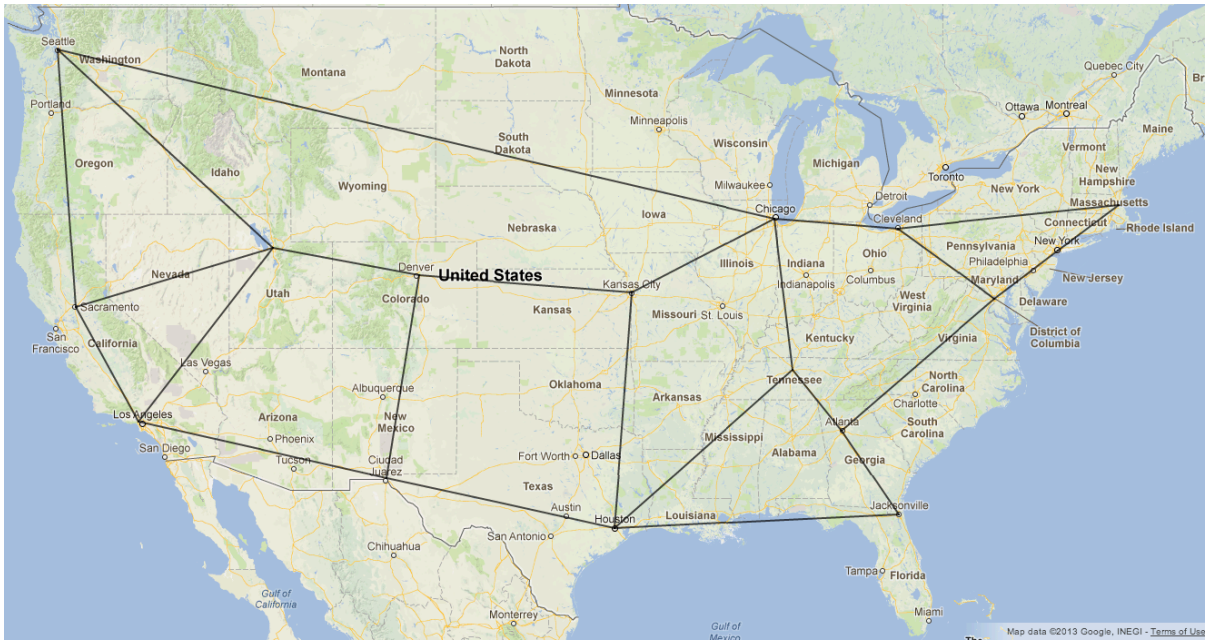

Figure 4.33: Visual representation of Internet2 structural topology

Figure 4.34: Visual representation of Internet2 Gabriel topology

ogy, it omits some of the links at the periphery of the actual geographic physical-level graph (e.g. link between Baton Rouge, LA and Jacksonville, FL) and adds links that are infeasible to deploy due to terrain. The synthetically generated geometric graph based on

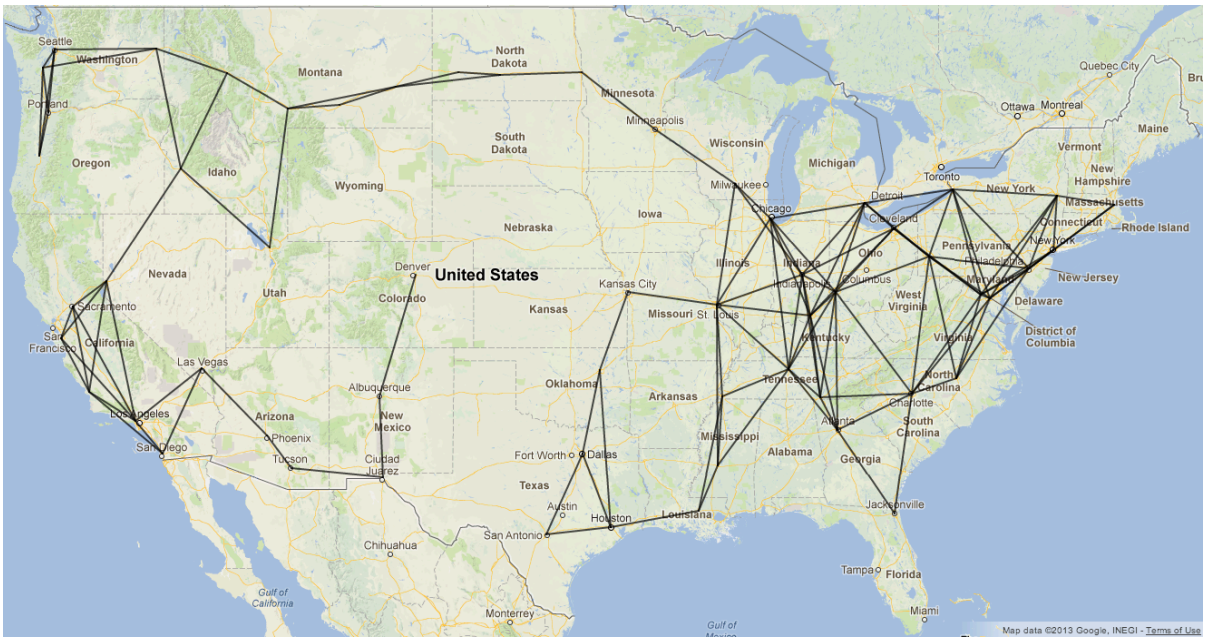

Figure 4.35: Visual representation of Internet2 geometric topology

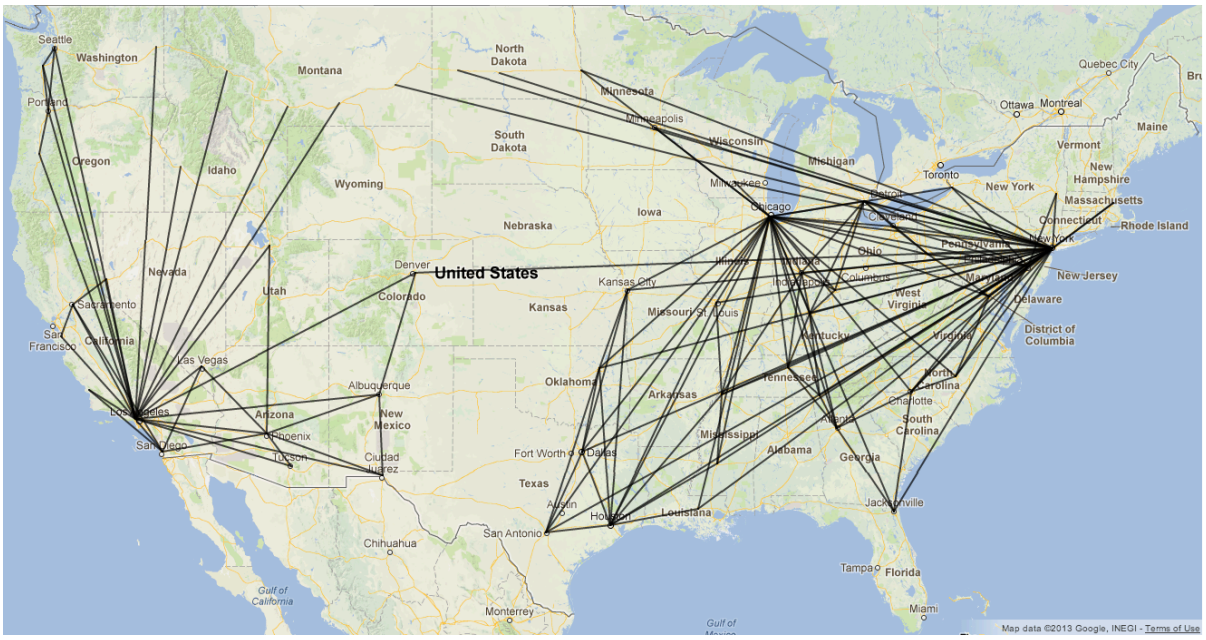

Figure 4.36: Visual representation of Internet2 geographical threshold topology

a distance threshold value that incurs minimal cost to obtain a connected graph is shown in Figure 4.35. In this case, while islands of nodes that are close to each other are richly connected, overall the graph is far from being biconnected. The geographical threshold

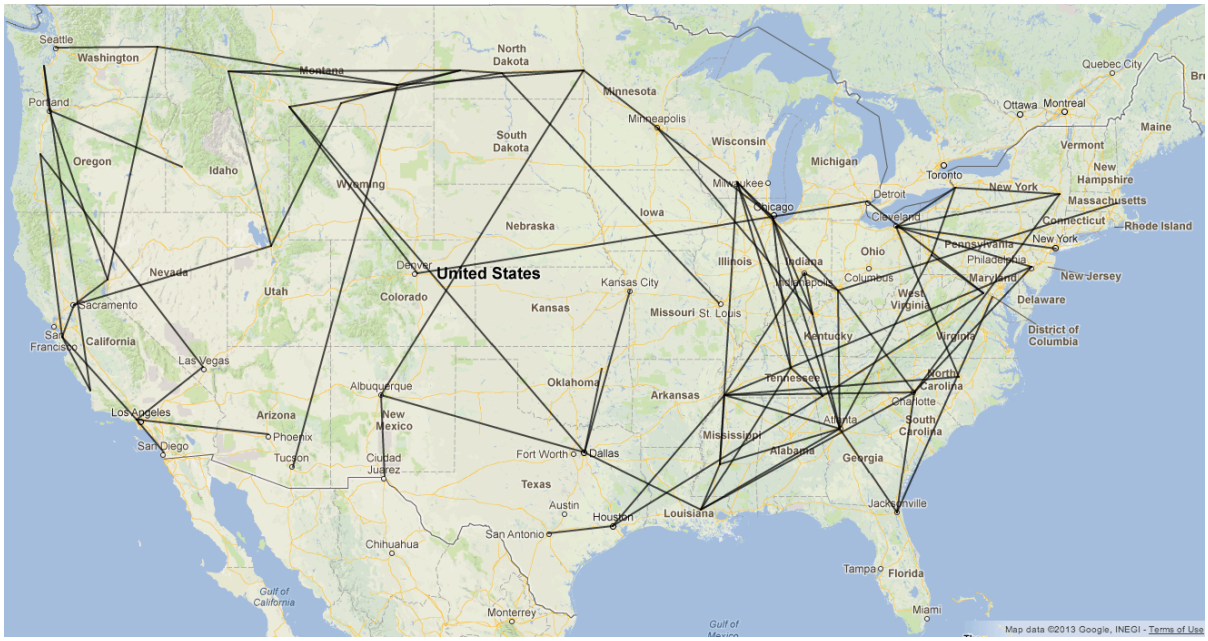

Figure 4.37: Visual representation of Internet2 Waxman topology

graph of the Internet2 topology using population of cities as node weights is shown in Figure 4.36. This synthetic graph resembles multiple star-like structures, because highly-populated cities become central nodes and connect to nodes that are far away. In this connected graph, there is only one link that connects east and west portions of the US. Finally, a location-constrained Waxman graph with $\beta = 0.6$ and $\alpha = 0.1$ values is shown in Figure 4.37. Because of the probabilistic nature of this graph model, the links between nodes are established randomly. In conclusion, Gabriel graphs are the closest to model physical level topologies with some caveats which we discuss in the next section.

4.4.5 On the Fitness of Synthetic Graph Models

In Section 4.4.4 we demonstrated that none of the synthetic geographical graph models we study capture the cost and structural properties perfectly. Based on our observations we present some ideas about how to develop a new geographic graph model that more closely captures the cost and structural behavior of physical topologies. First, we observe

that the presence of parameters within a graph model gives the user more control with regards to optimising the graph based on an objective function. Second, we note that while Gabriel graphs capture linear topologies that are horizontally aligned, they fall short in capturing star-like structures.

GTGs (geographical threshold graphs), on the other hand, generate star-like structures *aggressively* around heavily weighted nodes. For example, in Figures 4.38 and 4.39, we show the behavior of the Gabriel model and GTG model applied to a linear topology consisting of nodes horizontally aligned. In Figure 4.38, we see that the Gabriel model perfectly captures the linear topology, while in Figure 4.39 we see that the GTG model aggressively adds more links.

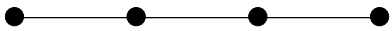

Figure 4.38: Gabriel graph model under linear geography

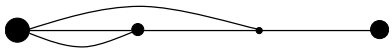

Figure 4.39: GTG graph model under linear geography

Next, consider a star-like graph as shown in Figure 4.40. While the Gabriel model aggressively changes it to a grid-like structure as depicted in Figure 4.41 showing the circles for determining the links, the GTG model can capture this star-like structure better than the Gabriel model depending on the node weight distribution, which we represent by giving different node sizes as shown in Figure 4.42. While each of these two models captures different structures better than the other, a better model would be able to select either the Gabriel model or the GTG model based on local structural criteria.

Finally, a detailed examination of Gabriel graphs show that they have two undesirable properties as compared to highly-engineered physical graphs. First, they add unnecessary ladder cross-connections between parallel linear segments in an attempt to increase the grid-like structure, and second they leave stub links that do not biconnect nodes on the

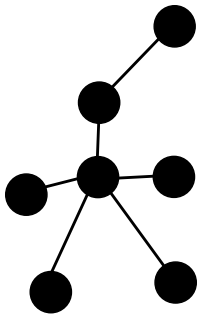

Figure 4.40: Actual graph model under star geography

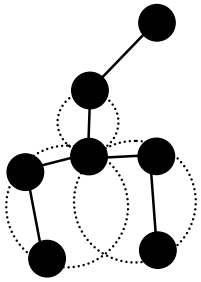

Figure 4.41: Gabriel graph model under star geography

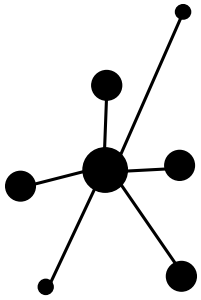

Figure 4.42: GTG graph model under star geography

edge into the rest of the graph. We will explore heuristics for a modified-Gabriel graph to address these issues in future work.

4.5 Summary

In this chapter, we investigate the graph properties of critical infrastructures using graph spectra and show that physical and logical level topologies have different structural characteristics (i.e. mesh vs. grid). Next, a framework is presented to model and analyse

multilevel and multiprovider networks and we show that single level analysis obscures realistic analysis of complex networks. Finally, we investigate existing synthetic graph generators and show that the Gabriel graph model best captures features of physical level topologies. However, to improve the connectivity of existing networks, new models and algorithms are required. In the next chapter, we present graph algorithms that optimise existing topologies for increased resilience in a cost-efficient manner.

Chapter 5

Network Design and Optimisation

.

Network connectivity can be improved by simply adding links to an existing graph. However, this might not be the most cost-efficient process. Moreover, networks cannot have unlimited resilience due to cost constraints because these two objectives fundamentally oppose one another. For example a complete (i.e. full mesh) graph has maximum resilience and the cost of building such a graph is also maximised. The design and optimisation of cost-efficient networks that are resilient against challenges and attacks has been studied by many researchers over the past few decades, but the resilient network design problem is NP-hard.

In this chapter, we approach resilient network design from a graph-theoretic perspective. We develop two *heuristic algorithms* that improve the connectivity of a graph by adding links. The graph properties we aim to improve are algebraic connectivity [72] and path diversity [71]. A secondary objective of our algorithms is to select the links that improve the resilience of the graph in the least costly fashion in which we capture the network cost as the total link length.

The work presented in this chapter has resulted in a publication [538] in which we develop

a graph algorithm that optimises the algebraic connectivity of a given graph. The rest of this chapter is organised as follows: In Section 5.1, the optimisation algorithm based on algebraic connectivity is presented. The optimisation algorithm based on path diversity is presented in Section 5.2. Results and analysis of our optimisation algorithms are presented in Section 5.3, and we conclude in Section 5.4.

5.1 Optimisation Based on Algebraic Connectivity

We develop a *heuristic algorithm* that improves the connectivity of a graph in terms of the *algebraic connectivity* metric by adding links. Algebraic connectivity $a(G)$ is defined as the second smallest eigenvalue of the Laplacian matrix [72] and it is widely used for topological optimisations [67, 68, 74]. A secondary objective of our algorithm is to select the links that improve the algebraic connectivity of the graph in the least costly fashion in which we capture the cost of network as the total link length. Furthermore, we parameterise our optimisation algorithm such that connectivity and cost are weighted depending on a cost-effect parameter named γ.

The heuristic to increase algebraic connectivity in a graph is based on adding links to the nodes that have least incident links (i.e. minimal degree nodes) [67, 68]. Our parameterised heuristic algorithm identifies and selects the links that increases the algebraic connectivity of a graph depending on the available budget. Moreover, the search of the optimal links is computationally less expensive in our algorithm compared to an exhaustive search. Our algorithm provides cost-efficient new links to improve a network's resilience measured by the algebraic connectivity metric. The assumptions, objective functions, and our heuristic algorithm is presented in Section 5.1.1. The evaluation of our algorithm on a sample graph is presented in Section 5.1.2.

5.1.1 Algebraic Connectivity Optimisation Algorithm

In this section, we describe our algorithm that optimises connectivity and cost of a topology. Our heuristic algorithm is implemented using Python. Furthermore, we assume that node geolocations are given for a particular graph to which the optimisation algorithm is applied, as would be the case for a deployed service provider.

Objectives

The objective of this algorithm is to identify the best links to be added to improve the connectivity of the graph. In this work, we use algebraic connectivity as a measure of connectivity, but we note that any graph connectivity property, such as average node degree, clustering coefficient, or betweenness can be used instead. For example, the clustering coefficient can be used to replace the algebraic connectivity or both the clustering coefficient and the algebraic connectivity can be used with a tuning parameter to control their effect in selecting the links.

Algorithm

The heuristic topology optimisation algorithm has three inputs: an input graph G_i, a number of required links L, and a cost-effect parameter γ. The input graph G_i has a number of nodes n_i with a number of links l_i. The number of required links L is the number of links that should be added to the graph. The cost-effect parameter γ is a tuning parameter between cost and algebraic connectivity. When $\gamma = 0$, the cost term of the rank function is neglected since it is zeroed. As a result, the algorithm selects the link that maximises the algebraic connectivity. On the other hand, when $\gamma = 1$, the algebraic connectivity is neglected and the least link cost is selected in each iteration. The algorithm adds links to the graph with L iterations. To keep track of the selected

links in each iteration, the algorithm adds these links to a list. In each iteration, the algorithm starts by adding the selected links from previous iterations to the graph. Then, the rank value is computed for each candidate link and the link with the maximum rank value is selected to be added. A ranking function is used to select the best candidate in each iteration. The rank value r is computed using:

$$r = (1 - \gamma)a(G) + \gamma(1 - C) \tag{5.1}$$

where C represents the length of the ranked link. This algorithm uses four functions: cost function $\mathsf{cost}(L)$, algebraic connectivity function $\mathsf{algConn}(G)$, link ranking function $\mathsf{maxLink}(D)$, and candidate link function $\mathsf{candidate}(G)$. The cost function $\mathsf{cost}(L)$ returns the cost of adding a link L. In this work, the cost is defined as the euclidean distance between the two ends of the link. The algebraic connectivity function $\mathsf{algConn}(G)$ takes a graph G and returns the second smallest eigenvalue of its Laplacian matrix. The $\mathsf{maxLink}(D)$ function returns the maximum ranked link. The $\mathsf{candidate}(G)$ takes a graph G as input and returns a set of candidate links to be added to the graph. The candidate links are a set of links that are examined every time a link is added to a graph. One option to use for the candidate links is the set of complement links of a graph is denoted as $\bar{G}$, which can be determined as the set of links in full mesh subtracted from the current links in a graph G. The number of complement links (cf. shown in column 4 Table 5.4) is computed as:

$$\frac{n_i(n_i - 1)}{2} - l_i \tag{5.2}$$

However, this number is computationally expensive as the number of nodes n_i gets larger, which results in a complexity of $O(Ln_i^2)$. In an attempt to decrease the number of

candidate links, we only examine the links connected to the lowest degree node in the graph. As a result, the algorithm complexity decreases to $O(Ln_i)$.

Both the algConn(G) and cost(L) functions are normalised to have a maximum value of one. Since the theoretical maximum value for the algebraic connectivity of a given graph is the number of its nodes, it is normalised by dividing it by the number of nodes. To normalise the cost function, it is divided by the maximum possible distance between any nodes in the graph. The pseudocode of our algorithm is shown in Algorithm 1.

Functions:
cost(L) := cost function
algConn(G) := algebraic connectivity function
candidate(G) := candidate links function
maxLink(D) := max value of a dictionary
Input:
G_i := input graph
L := number of required links
γ := cost effect parameter
Output:
an ordered list of the added links
begin
    selectedLinks = []; empty ordered list
    rank = {}; empty dictionary
    while $L > 0$ **do**
        G = G_i
        G.addlinks(AddedLinks)
        for link in candidate(G) **do**
            rank[link]=$(1-\gamma)$algConn(G) + γ(1−cost(link))
        end
        selectedLinks.add(maxLink(rank))
        $L = L - 1$
    end
    return selectedLinks
end

Algorithm 1: Algebraic connectivity optimisation algorithm

Options

In this work, we target two graph types: physical and logical level graphs [75]. For the logical level graph, the algorithm is applied with no additional conditions. However, for the physical level graph, we have added a filter that removes very long links from the candidate links set. This is because it is not practical to add very long links between cities such as a physical fibre link between Los Angeles and New York City. Therefore this raises the question of what the maximum length should be chosen. In our implementation, we provide it as a variable that can be set by the user. We choose the maximum length link in the input graph to be the threshold for long links in the dataset, which gives a good indicator for the maximum link length a provider can afford.

5.1.2 $a(G)$ Optimisation Algorithm Evaluation

In this section, we explain how our heuristic algorithm optimises a topology on a small-size graph. Figure 1 shows a sample graph with 8 nodes and 9 links as solid lines. The initial algebraic connectivity of this sample graph is 0.3432 and the initial cost (i.e. total link length in km) of the graph is 8,203. Our heuristic algorithm adds links to the least connected nodes, which in the example are nodes 0 and 7. The six candidate links for node 0 are shown as square dots, whereas five candidate links for node 7 are shown as long dashes and dots. Throughout this example, we describe how our algorithm operates if we are going to add *one* link $L = 1$ to the sample graph shown in Figure 5.1.

There can be a maximum of 28 links in this 8-node graph (maximum links can be calculated by $\frac{n(n-1)}{2}$). Since there are 9 links in the graph, if we were to do an exhaustive search, there would be $28 - 9 = 19$ candidate links (i.e. the complement links). In the sample graph shown in Figure 5.1, there are six candidate links that can be added to node 0 and there are five links for node 7 using our heuristic algorithm. Therefore, the

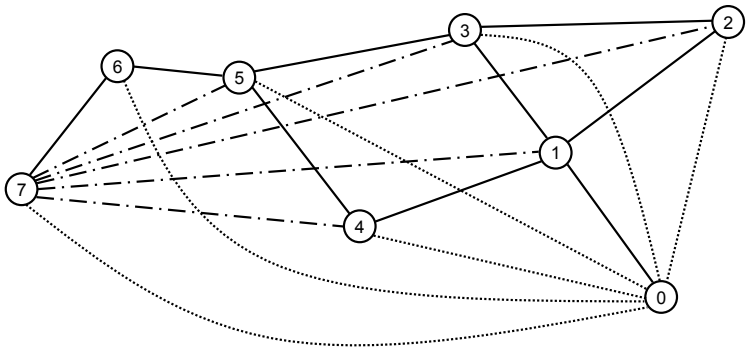

Figure 5.1: Graph example for algebraic connectivity based optimisation

candidate link set is reduced to 11, because our algorithm only considers candidate links from the least connected nodes. The algebraic connectivity and cost value of adding each link individually for $\gamma = 0$ and $\gamma = 1$ is shown in Table 5.1.

Table 5.1: $a(G)$ and cost values for the example graph

Link	$\boldsymbol{\gamma = 0}$		$\boldsymbol{\gamma = 1}$	
	$\boldsymbol{a(G)}$	$\boldsymbol{\Delta a(G)}$	**cost**	**Δ cost**
$0 \leftrightarrow 2$	0.3485	0.0053	**9,275**	**1,072**
$0 \leftrightarrow 3$	0.3588	0.0156	9,405	1,202
$0 \leftrightarrow 4$	0.3659	0.0227	9,848	1,645
$0 \leftrightarrow 5$	0.4079	0.0647	10,624	2,421
$0 \leftrightarrow 6$	0.5908	0.2476	11,228	3,025
$0 \leftrightarrow 7$	0.7713	0.4281	11,843	3,640
$7 \leftrightarrow 1$	**0.8345**	**0.4913**	11,302	3,099
$7 \leftrightarrow 2$	0.7071	0.3639	12,061	3,858
$7 \leftrightarrow 3$	0.6651	0.3219	10,915	2,712
$7 \leftrightarrow 4$	0.5918	0.2486	10,207	2,004
$7 \leftrightarrow 5$	0.5075	0.1643	9,463	1,260

When $\gamma = 0$, our algorithm ignores the cost associated with adding a link and selects the additional link that increases the algebraic connectivity of the graph the most. For $\gamma = 0$, the algorithm adds the link between node 1 and 7 in the example graph since it provides the highest algebraic connectivity among the 11 candidate links. When $\gamma = 1$, the cost is the dominant factor determining the addition of a link. Therefore, our heuristic

algorithm selects the link between node 0 and 2, since it incurs the lowest cost among the candidate set of links. The selection of links via our heuristic algorithm is highlighted bold in Table 5.1. Moreover, we perform an exhaustive search on the sample graph shown in Figure 5.1, and find that the link between node 1 and 7 has the highest algebraic connectivity among 19 possible links. The result of the exhaustive search for the least incurred cost link indicated that the link between node 3 and 4 is the best option, however, as mentioned above, our algorithm adds links to the minimal degree nodes. Therefore our algorithm selects the link between node 0 and 2 when $\gamma = 1$. We note that for physical-level networks $\gamma \to 1$ makes sense due to the significant cost of deploying fibre. On the other hand, $\gamma \to 0$ is more appropriate since the cost of virtual link deployment is negligible, whereas delay is a dominant factor in logical overlays. To conclude, our heuristic algorithm optimises graphs cost-efficiently while selecting the links that improves the algebraic connectivity the most based on the γ parameter value.

5.2 Optimisation Based on Path Diversity

In this section, first we develop an algorithm to calculate the TGD (described in Section 2.3.2) of a graph [71, 73]. We modify the TGD calculation algorithm such that instead of considering relatively more diverse paths [71, 73], we consider the effect of the diversity of all paths when calculating the TGD. Second, we introduce an algorithm for finding the optimal k-diverse paths considering both nodes and links using an exhaustive path search. Thirdly, we present an optimisation algorithm that improves the path diversity of a graph based on the TGD metric. Our graph optimisation algorithm considers adding links with the least cost among available choices.

The rest of the section is organised as follows: The algorithm to calculate the path diversity of a graph is explained in Section 5.2.1. The assumptions, objective functions,

and our heuristic algorithm is presented in Section 5.2.2. The evaluation of our algorithm on a sample graph is presented in Section 5.2.3.

5.2.1 Finding k-Diverse Paths

In this section, we present a new k-diverse paths algorithm that determines the k paths between a source s and destination d. We note that the EPD [71] is defined as the aggregation of path diversities for a selected set of paths between a given node-pair (s, d) where the value of k_{sd} captures the sum of the minimum path diversities among the selected paths. We redefine the value of k_{sd} as the sum of path diversities because it captures the difference between paths with the same minimum diversity. Our algorithm returns the most diverse paths considering both fully- and partially-disjoint nodes and links in a given path sorted by their length in case of equal diversity. This algorithm has four inputs: a source node s, a destination node d, a hop count threshold h, and a threshold for the number of returned diverse paths k. Moreover, this algorithm uses four functions: $\mathsf{all\_simple\_paths}(s, d, h)$, $\mathsf{sort}(L)$, $\mathsf{path2elements}(P)$, and $\mathsf{p\_div}(P)$. The $\mathsf{all\_simple\_paths}(s, d, h)$ function finds all possible loopless paths between a source s and destination d, with hop count threshold h for the path length. If h is not set, all possible paths are returned. The number of all possible simple paths can be as large as $n!$ where n is the number of nodes. This number can be infeasible to compute for large size graphs. Thus, the h parameter should be set based on the size of the graph. The $\mathsf{sort}(L)$ function sorts a list of tuples of three elements: link, diversity, and cost. The sorting is done in decreasing order of the diversity value and increasing order of the cost value for links with equal diversities. The $\mathsf{path2elements}(P)$ function converts a path P to a set of nodes and links elements as described in Section 2.3.2. The $\mathsf{p\_div}(P)$ function computes the diversity of the path with respect to the $\mathsf{selected\_elements}$ set. The pseudo code for k-diverse path selection is shown in Algorithm 2.

Functions:
all_simple_paths(s, d, h) := all simple paths between node s and node d with a threshold hop count h
sort(l) := sorting l function
path2elements(p) := convert path to link and node elements
p_div(P) := computes path diversity of path P
Input:
s := source node
d := destination node
h := hop count threshold value for examined paths
k := threshold value for returned diverse paths
Output:
an ordered list of the diverse paths
begin
    diverse_paths = []; empty ordered list
    selected_elements = {}; empty set
    for path in all_simple_paths(s, d, h) **do**
        diverse_paths.append(path,p_div(path),len(path))
        selected_elements.add(path2elements(path))
    end
    sort(diverse_paths)
    return diverse_paths[0:k]
end

Algorithm 2: k-diverse path selection algorithm

This algorithm has two phases: finding all simple paths and finding the k most diverse paths. In the first phase, all possible paths are determined between a source s and destination d with a hop count threshold h using the function all_simple_paths(s,d,h). For the second phase, the algorithm determines the most diverse paths among the returned paths via the all_simple_paths function. The shortest path P_0 is added to the selected paths in the first iteration and its elements (nodes and links) are added to the selected_elements set. Next, the algorithm iterates over the rest of the paths by computing the diversity of the path using the p_div(P) function and adding it along with the path length to the diverse_paths list while the path elements are added to the selected_elements. Finally, using the sort(L) function, all the tuples in the diverse_paths list are sorted in decreasing order of their diversity and in case there are multiple paths with the same diversity, these paths are sorted in increasing order of their hop count. The first k paths of the list are returned.

5.2.2 Path Diversity Optimisation Algorithm

In this section, we describe our algorithm that optimises the TGD of a given graph with its node locations by adding new cost-efficient links. Our heuristic algorithm is implemented using Python and uses the NetworkX library [512] for graph algorithms.

Objectives

The objective of this algorithm is to improve the TGD of a graph by adding a user-specified number of links. The algorithm adds one link that increases the TGD the most. If there are multiple links that give the same largest TGD value, the least cost link is selected. We measure the cost of a link in terms of the Euclidian distance of that link. The link addition process is repeated until the number of links requested by the user is added.

Algorithm

The topology optimisation algorithm has two inputs: an input graph G_i, a number of required links L. The input graph G_i has a number of nodes n_i with a number of links l_i. The number of required links L is the number of links that should be added to the graph. The algorithm adds links to the graph with L iterations. To keep track of the selected links in each iteration, the algorithm adds this link to the selectedLinks list. In each iteration, the algorithm starts by adding the selected links from previous iterations to the graph.

The candidate set contains the links that are connected to the lowest EPD pair(s) of the graph and not currently present in the graph. To find the best candidate link, each link in the candidate set is added to the graph and the EPD of the corresponding pair is computed and mapped to that link. Then, the link with the largest EPD is selected. In case there are multiple links with the same largest EPD, the least cost link is selected. This process is repeated until the user requested number of links are added.

This algorithm uses four functions: cost(L), epd(P), candidate(G), and bestLink(L). The cost function cost(L) returns the cost of adding a link L. In this work, the cost is defined as the euclidean distance between the two ends of the link. The effective path diversity function epd(P) computes the effective path diversity of the link L based on Equation 2.3. The bestLink(L) function returns the link with the highest EPD and lowest cost in case of multiple highest EPD values. The candidate(G) takes a graph G as input and returns a set of candidate tuples of two elements. The first element is a lowest EPD pair and the second element is a candidate link. The candidate links are the set of all links where one end is connected to a node in the lowest EPD pairs and the other end is connected to a node in the graph given that this link does not exist in the graph. For each pair and link in the candidate set, we add the link to the graph and compute the new EPD

value of that pair with its cost. Finally, the link with the highest EPD and the lowest cost is selected using bestLink(L) function and then added to the selectedLinks list. The algorithm repeats this process as many times as the user requested. The pseudo code for path diversity optimisation algorithm is shown in Algorithm 3.

Functions:
cost(L) := cost of link L
epd(P) := EPD value of pair P
candidate(G) := candidate links function
bestLink(L) := maximum EPD value of links list L
Input:
G_i := input graph
L := number of required links
Output:
an ordered list of the added links
begin
  selectedLinks = []; empty ordered list
  links_epd_list = []; empty ordered list
  while $L > 0$ **do**
    G = G_i
    G.addlinks(AddedLinks)
    for P, L in candidate(G) **do**
      links_epd_list.append((L,epd(P),cost(L)))
    end
    selectedLinks.add(bestLink(links_epd_list))
    $L = L - 1$
  end
  return selectedLinks
end

Algorithm 3: Path diversity optimisation algorithm

We only examine the links connected to the lowest EPD pair node in the graph since they contribute the most to lower the TGD of the graph. If there is more than one pair with the lowest EPD value, all the links connected to all the lowest EPD pairs are added to the candidate set. The other alternative is to examine all the links in complement graph of G_i, which has the size of $\frac{n_i(n_i-1)}{2} - l_i$. However, this number is computationally expensive as the number of nodes n_i gets larger, which results in a complexity of $O(Ln_i^2)$. Since our approach considers the links connected to the lowest EPD pairs, the complexity is

reduced to $O(Ln_i m_i)$, where m_i represent the number of lowest degree pairs at iteration i.

5.2.3 Path Diversity Optimisation Algorithm Evaluation

In this section, we explain how our heuristic algorithm optimises a topology on a small-size graph. Figure 5.2 shows a sample graph with 5 nodes and 5 links. In this example, the hop count threshold h is set to 10 and the number of diverse paths k is set to 4. The initial TGD value of this sample graph is 0.2023. Our heuristic algorithm examines the links connected to the least EPD pairs. The smallest EPD pairs are (1,2) and (3,4) with EPD of 0 since they have no alternative paths. Therefore, the candidate set consists of four possible links for each pair. To find the best candidate, we determine the resulting EPD of the corresponding pair after adding the candidate link and the cost incurred as shown in Table 5.2. Then, we find the link that gives the highest pair EPD. Among the eight candidates, there are four links that give a high pair EPD of 0.50. The next step is to find the lowest length link, which is the link (1,3). After adding this link, the new TGD of this graph is 0.4034, which is almost double the initial TGD.

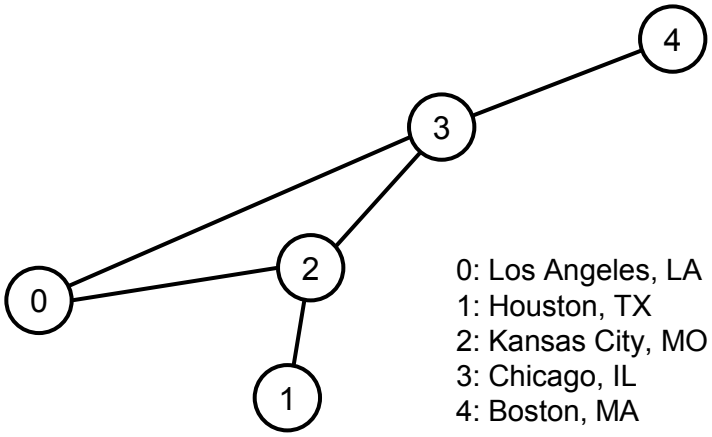

Figure 5.2: Graph example for path diversity based optimisation

Table 5.2: EPD and cost values for the candidate links in the example graph

Node Pair	Link	Pair EPD	Cost
(1, 2)	(1, 0)	0.50	2,177
(1, 2)	(1, 3)	**0.50**	**1,043**
(1, 2)	(1, 4)	0.46	2,311
(1, 2)	(2, 4)	0.00	1,988
(3, 4)	(3, 1)	0.00	1,043
(3, 4)	(4, 0)	0.50	4,058
(3, 4)	(4, 1)	0.46	2,311
(3, 4)	(4, 2)	0.50	1,988

5.3 Results and Analysis

In this section, we present results and analysis of topological optimisation based on algebraic connectivity in Section 5.3.1. This is followed by presentation of results for path diversity based optimisation of topologies in Section 5.3.2. We compare the two optimisation algorithms in Section 5.3.3. We note that full set of plots showing the analysis of our optimisation algorithm is presented in Appendix B.

5.3.1 Analysis of Optimisation Based on $a(G)$

We study physical- and logical-level topologies of three tier-1 service provider networks. We use Rocketfuel-inferred [33] logical-level topologies of AT&T, Level 3, and Sprint. Physical-level topologies of the three service providers were constructed using a third party map [505]. The details of generating physical-level topologies are presented in Chapter 4. The number of nodes, links, and complement links of these graphs are shown in Table 5.3.

Table 5.3: Topological dataset for algebraic connectivity optimisation

Network	Nodes	Links	Complement links
AT&T phy.	383	488	72,665
Level 3 phy.	99	132	4,719
Sprint phy.	264	313	34,403
AT&T log.	107	140	5,531
Level 3 log.	38	376	3,276
Sprint log.	28	76	302

Backbone Provider Network Analysis

Our algorithm is applied to three ISPs by adding 100 links. We show the graph algebraic connectivity and the cost incurred in terms of meters after adding each link. Moreover, we show the relation of cost and algebraic connectivity and the slope in these figures shows how the cost increases as the graph connectivity improves.

Selection of γ Values

γ parameter that ranges 0 to 1 controls the outcome of the algorithm as described in Section 5.1.1. In Equation 5.1, we have two terms: $(1-\gamma)a(G)$ and $\gamma(1-C)$. The $a(G)$ is the normalised algebraic connectivity value, which is low for sparse graphs and one for a full mesh graph. The value of C denotes the normalised cost of adding a link and it is low when the maximum possible link length in the input graph is larger than the average link length in the candidate set. Therefore, choosing the value of γ depends on the initial properties of the input graph. For the physical graphs, we choose for $\gamma = \{0, 10^{-9}, 10^{-7}, 10^{-5}, 1\}$ because the cost term is larger than the γ term by about six order of magnitude for physical level graphs. For the logical level graphs, we choose different values of $\gamma = \{0, 10^{-4}, 10^{-3}, 10^{-2}, 1\}$ because the cost term is larger than the algebraic connectivity term by about two order of magnitude.

Physical-level Topology Analysis

As explained in Section 5.1.1, an option is added in our heuristic algorithm to discard the links that are longer than the actual maximum link of the graph. Furthermore, physical level graphs have more nodes than the logical level graphs, which increases the number of shorter links for the candidate set. For these reasons, optimisation on physical level graphs results in selection of shorter links.

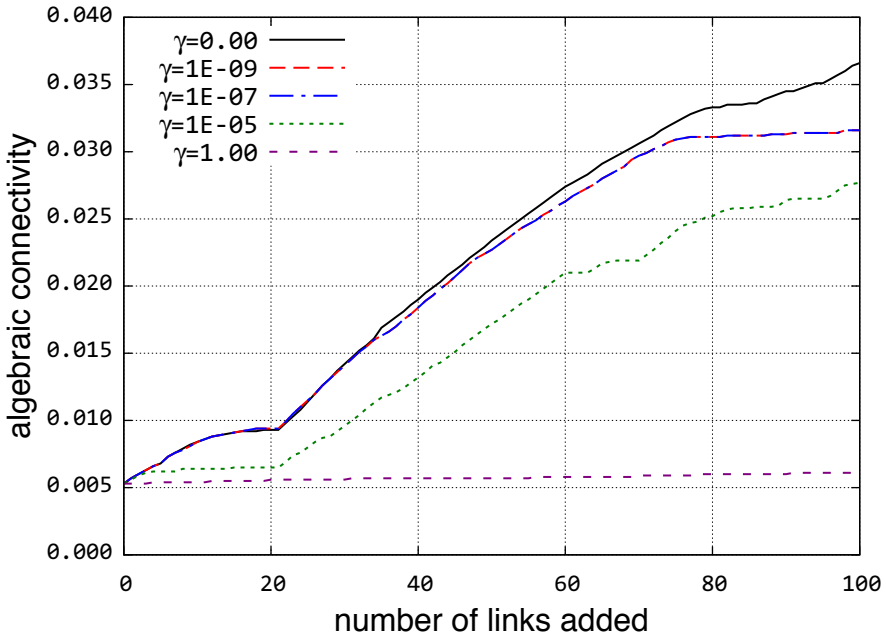

Figure 5.3: Connectivity improvement for Sprint physical topology

Algebraic connectivity improvement of the Sprint physical level topology after adding 100 links iteratively is depicted in Figure 5.3. The algebraic connectivity is higher for $\gamma = 0$ than the other values of γ, and for $\gamma = 1$ our algorithm considers minimising the cost, but not improving the algebraic connectivity. Moreover, we observe the occurrence of possible phase transition when $\gamma = 1$ for the physical-level graphs. For example, algebraic connectivity improvement of the Sprint physical topology starts with a moderate increase, and after about 20th link addition, the improvement (i.e. the slope of the curve) gets

steeper. The reasons for the occurrence of this phenomenon will be the subject of future work.

The cost incurred when adding 100 links iteratively to the Sprint physical level topology is shown in Figure 5.4. The cost in physical topology is the length of links to be laid between nodes, thus, short links are favorable in physical level topology optimisation for $\gamma = 1$.

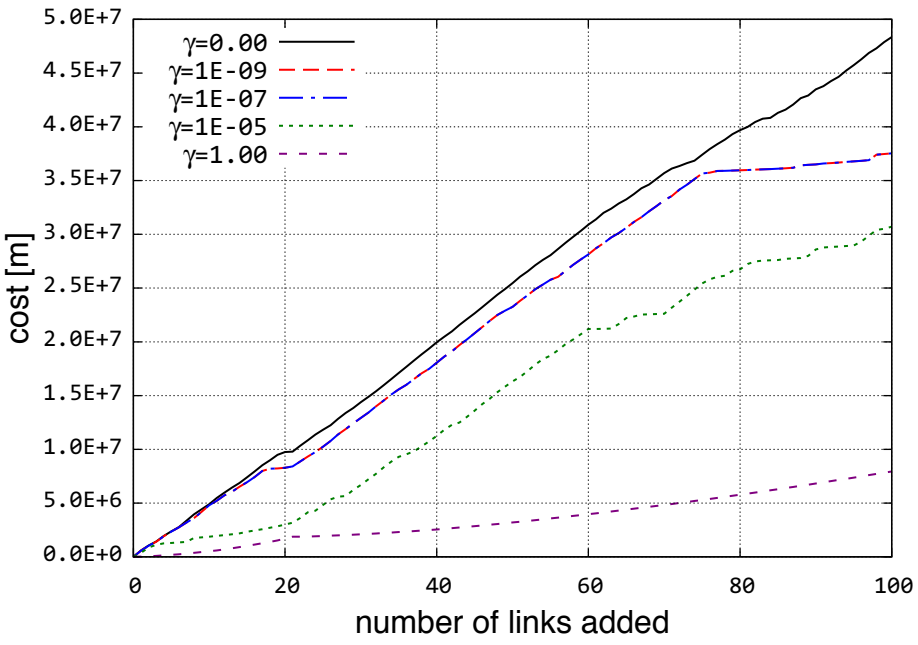

Figure 5.4: Cost incurred with adding links for Sprint physical topology

The relationship between connectivity and cost for the Sprint physical level topology is shown in Figure 5.5. For the Sprint example shown in Figure 5.5, if the cost is the constraint (i.e. $\gamma = 1$), the designer can improve the algebraic connectivity to 0.006 by adding 100 links. On the other hand, if there is available budget (i.e. $\gamma = 0$) the algebraic connectivity of the Sprint topology can be improved more than 0.035.

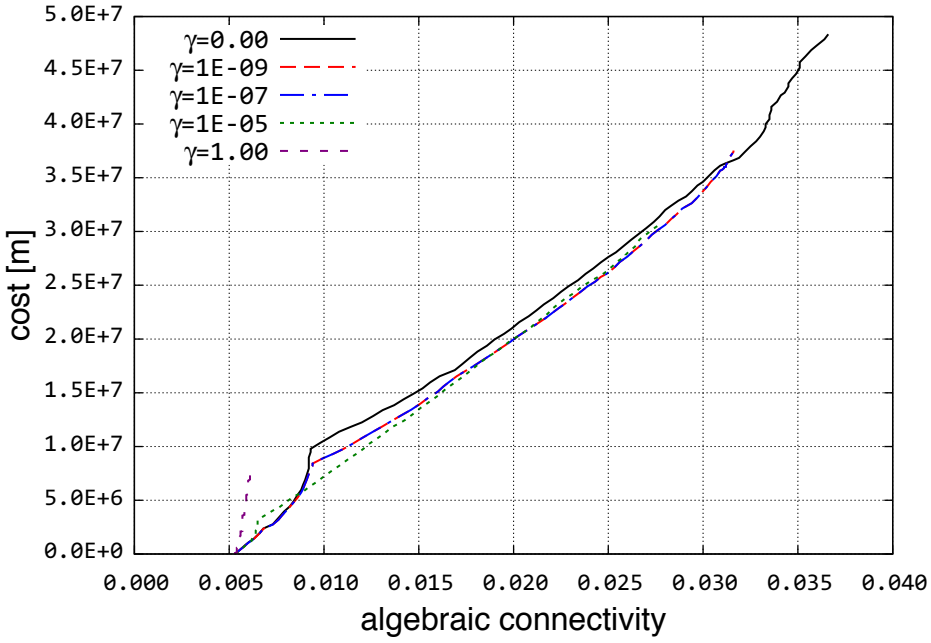

Figure 5.5: Connectivity and cost trade-offs for Sprint physical topology

Logical-level Topology Analysis

The optimisation of the Sprint logical level topology is discussed in this section. The algorithm has more candidate link options since it is not constrained by the maximum length of links in the input graph for logical level topologies. Therefore, the improvement of algebraic connectivity as links are added is higher than for physical level topologies. The algebraic connectivity improvement of up to two orders of magnitude, can be seen clearly for the Sprint logical level topology in Figure 5.6.

The cost incurred after adding 100 links for the Sprint logical level topology is shown in Figure 5.7. Similar to the physical level topologies, as the value of γ increases, the cost of building more connected graphs decreases.

The trade-offs between cost and connectivity for the Sprint logical level topology is shown in Figure 5.8. For example, to improve the algebraic connectivity of the Sprint logical

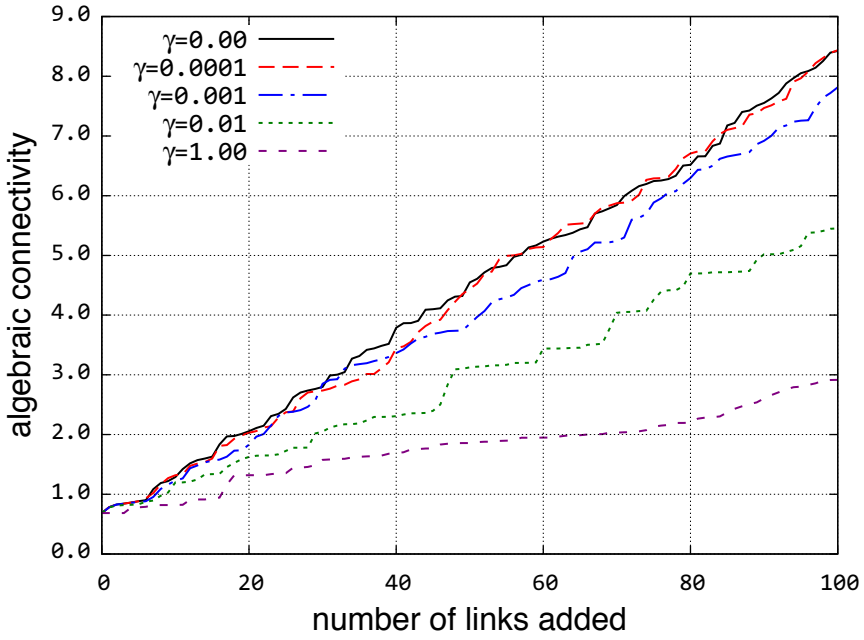

Figure 5.6: Connectivity improvement for Sprint logical topology

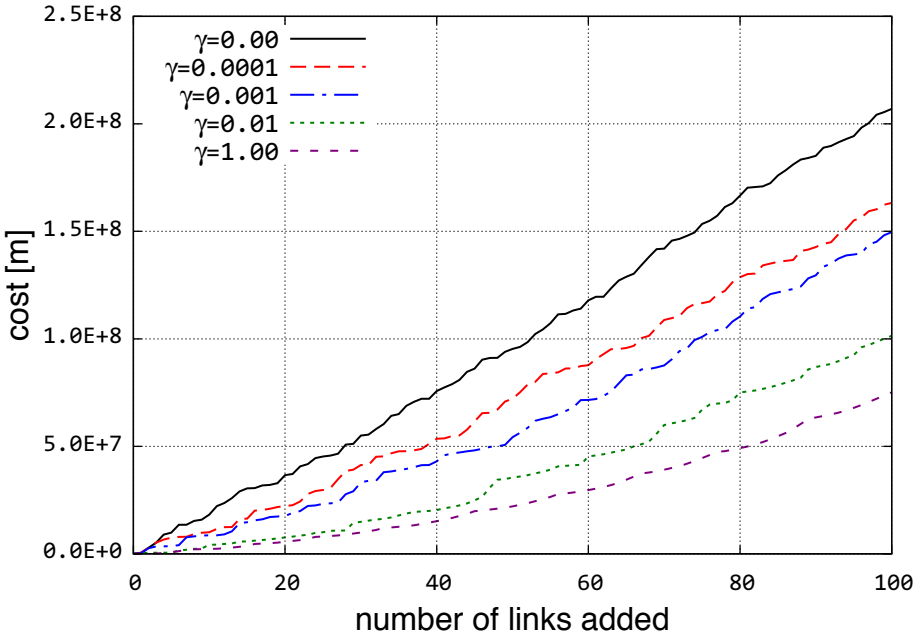

Figure 5.7: Cost incurred with adding links for Sprint logical topology

topology to a value of 10 in Figure 5.8, we should select the links returned from the algorithm when γ is 0.01 since it incurs the lowest cost.

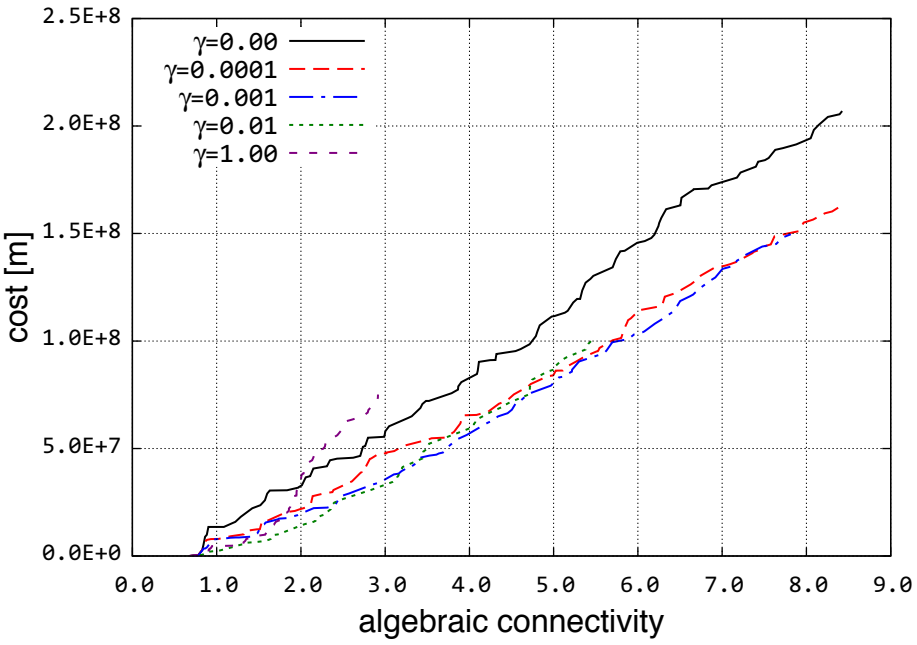

Figure 5.8: Connectivity and cost trade-offs for Sprint logical topology

Optimisation Comparison of Backbone Networks

Finally, we compare the optimisation output of the three backbone provider topologies using $\gamma = 0$ and $\gamma = 1$ as shown in Figures 5.9 through Figure 5.12, respectively. Even though Sprint and AT&T physical level topologies have a different number of nodes and links, the optimisation for $\gamma = 0$ results in about the same algebraic connectivity as when we add 100 links as shown in Figure 5.9. On the other hand, the Level 3 physical level topology starts from an even higher initial algebraic connectivity and significantly improves to a higher algebraic connectivity with larger cost than the others as shown in Figure 5.9. Similar conclusions can also be drawn when we compare our optimisation algorithm output for $\gamma = 1$ as shown in Figure 5.11. The main difference is that when there is no budget constraint (i.e. $\gamma = 0$), the algebraic connectivity and cost is an order of magnitude higher. Moreover, the slope in these figures shows how the cost increases as we improve the connectivity of the graph. Small slope value of a curve implies high

gain in the graph connectivity for a given budget, which is favorable. On the other hand, large slope value means that the cost is high while the improvement is low.

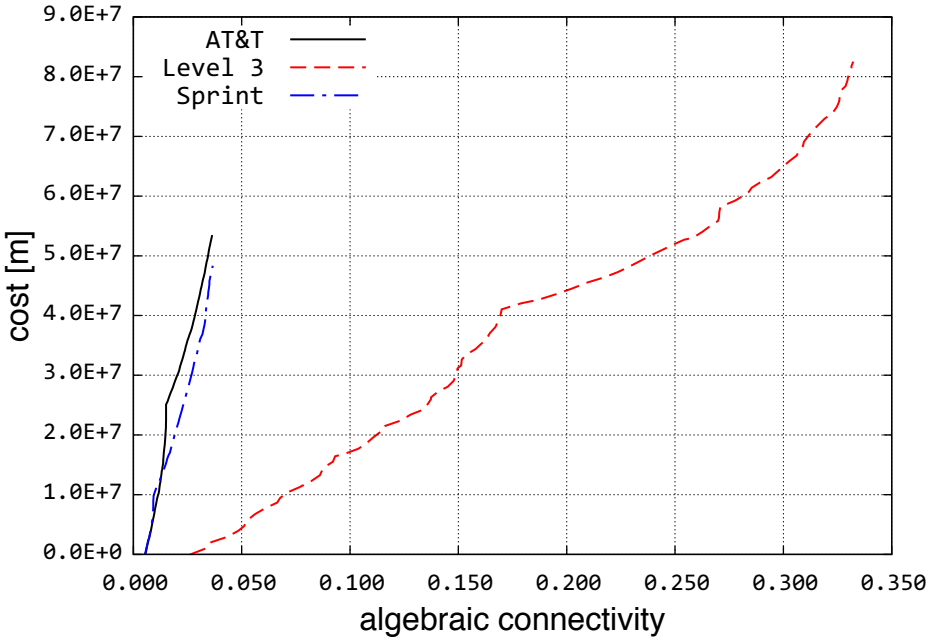

Figure 5.9: Algebraic connectivity and cost effect for $\gamma = 0$ for physical-level topologies

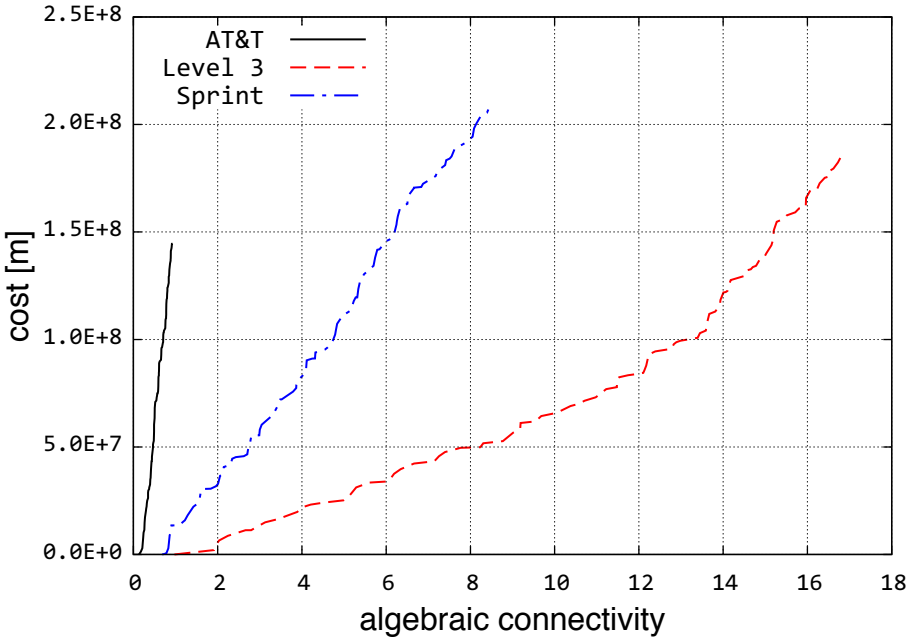

Figure 5.10: Algebraic connectivity and cost effect for $\gamma = 0$ for logical-level topologies

The tradeoffs between cost and algebraic connectivity for the logical topology graphs when $\gamma = 0$ is shown in Figure 5.10. For AT&T, we see that algebraic connectivity does not improve for adding the first links, with total cost around 1.5×10^8. On the other hand, for Level 3 and Sprint, the algebraic connectivity increases significantly as links are added as shown in Figure 5.10. When $\gamma = 1$, which means the least cost links are selected, the algebraic connectivity does not improve much as links are added to the graph compared to $\gamma = 0$ as shown. Another interesting result is that the Sprint and Level 3 topologies incur about the same cost after adding 100 links, however the algebraic improvement for Level 3 is twice or more for $\gamma = 0$ and $\gamma = 1$ values. Finally, physical level topologies have lower gain in terms of the algebraic connectivity since long links are removed from the candidate set and these links can be the highest contributors to the algebraic connectivity.

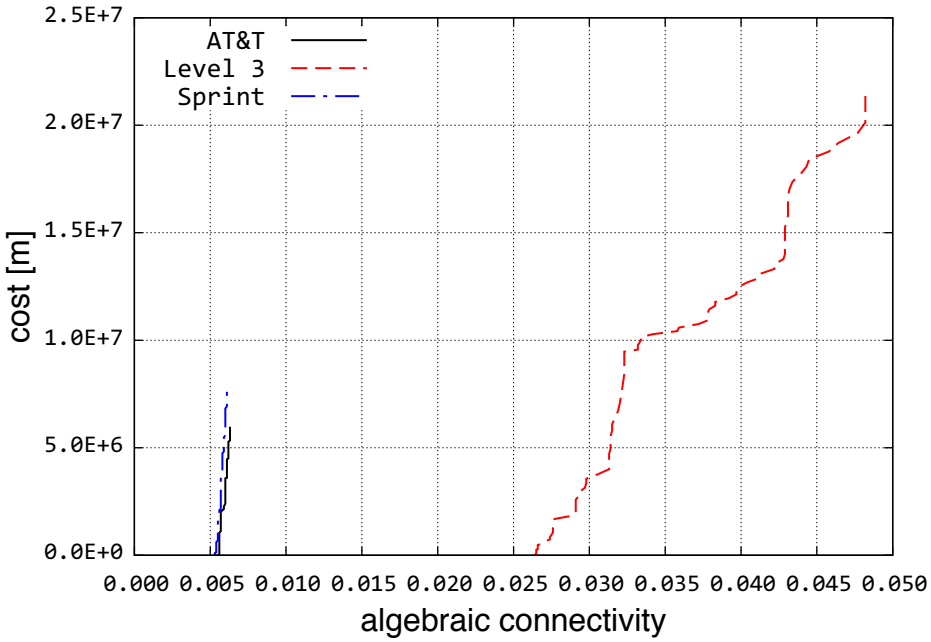

Figure 5.11: Algebraic connectivity and cost effect for $\gamma = 1$ for physical-level topologies

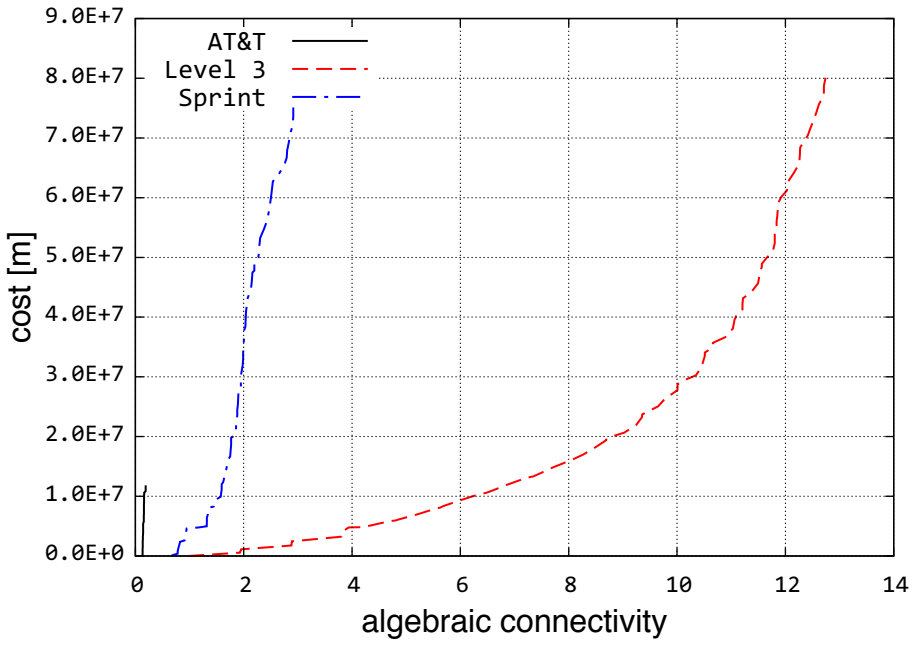

Figure 5.12: Algebraic connectivity and cost effect for $\gamma = 1$ for logical-level topologies

5.3.2 Analysis of Optimisation Based on Path Diversity

In this section, we use CORONET [508,509], Internet2 [507], and Level 3 [504] fibre-level topologies for evaluating our path diversity optimisation algorithm. In Table 5.4 we list a number of relevant quantities for each of the provider networks. A detailed analysis of graph metrics for the given physical networks was presented in Chapter 4. Next, we apply the optimisation algorithm on three realistic backbone networks and study the results. Then, we apply three centrality-based attacks to the resulting optimised and non-optimised graphs and show how the robustness changes for each graph.

Table 5.4: Topological dataset for path diversity optimisation

Network	**Nodes**	**Links**	**Avg. Degree**	**Diameter**	**Avg. Hopcount**
CORONET	75	99	2.64	17	6.45
Internet2	57	65	2.28	14	6.69
Level 3	99	132	2.67	19	7.65

Optimisation Analysis

In this section, we apply the optimisation algorithm on three realistic backbone service provider graphs and study the TGD improvement and the cost incurred for each graph as we add 20 links. We vary the number of value of k and h while λ is set to 0.5.

Varying the Hop Count Threshold h

The hop count threshold h is a parameter that controls the length of the shortest path returned by the k diverse algorithm introduced in Section 5.2.1. Therefore, to get the optimal diverse paths, the value of h should be larger or equal to the diameter of the graph in order to examine all of the possible paths in the graph. However, for large graphs, large values of h may take an impractical time to calculate. Here, we apply the algorithm with several values of hop count thresholds $h = \{5, 10, 15\}$ while the value of k is set to 12. These values show how varying the parameter h affects the value of TGD.

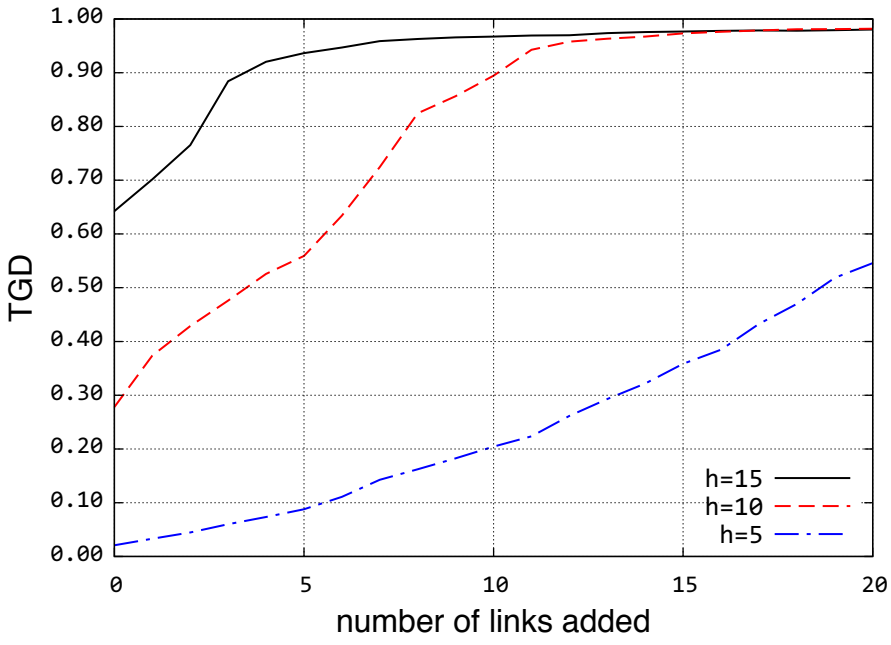

Figure 5.13: Internet2 TGD improvement

Figure 5.13 depicts the results of each hop count threshold, which shows the TGD improvement as 20 links are added to the Internet2 topology. As the hop count threshold increases, the size of the candidate set also increases, which in turn increases the probability to have a higher EPD value. As a result, a 5 hop count threshold has the lowest TGD while 10 and 15 have the median and the highest TGD, respectively as shown in Figure 5.13.

The cost does not follow a pattern as the hop count threshold increases since the cost for the highest EPD link for 10 hop count threshold could be less than the cost of the highest EPD for 15 hop threshold and vice versa as shown in Figure 5.14. Thus, the cost incurred depends on the initial topological properties such as the number of nodes and links, average degree, and nodes' locations.

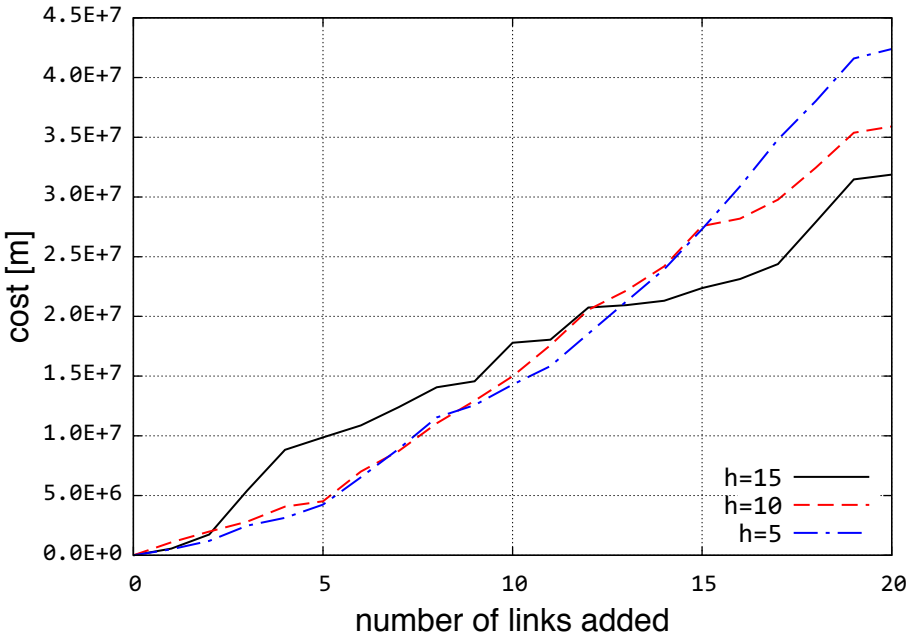

Figure 5.14: Internet2 cost incurred

The cost needed to achieve a certain TGD for Internet2 topology is shown in Figure 5.15, which shows that as the hop count threshold increases, the cost to achieve a certain TGD

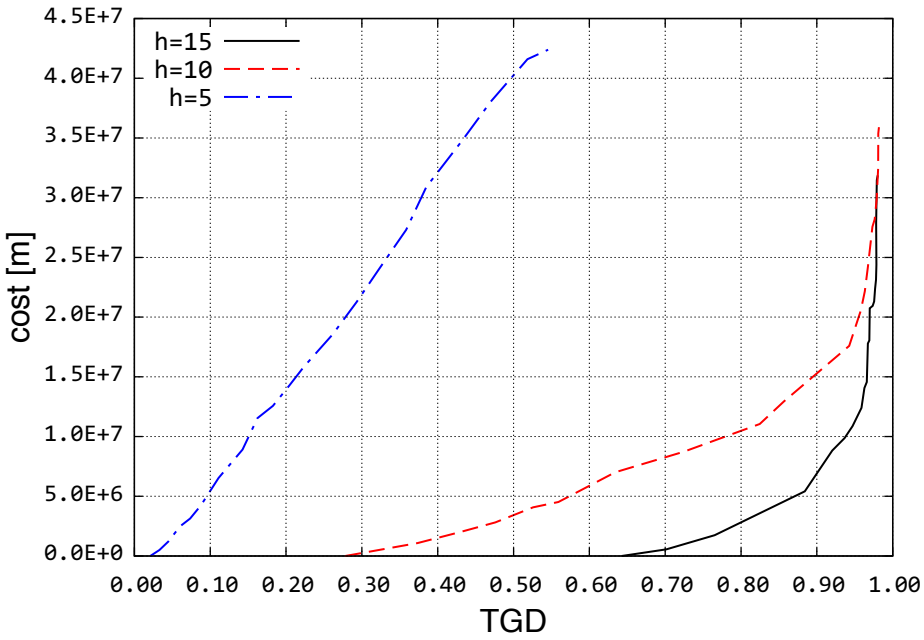

Figure 5.15: Internet2 cost and TGD

decreases in general. We expect this outcome to mostly occur because higher hop count threshold starts from a higher TGD for the same reason mentioned earlier.

Varying the Number of Diverse Paths k

The number of diverse path k is a parameter that controls the number of the returned most diverse paths by the k diverse algorithm introduced in Section 5.2.1. The value of k depends on the application of the graph. For example, if the provider uses a multipath routing protocol with a threshold for the number of multipaths used, k can be chosen to match that parameter for accurate path diversity. Choosing a high value of k does not have a processing complexity penalty similar to choosing a higher value of h.

We apply the algorithm with several values of the number of diverse path threshold $k = \{4, 8, 12\}$ while the value of h is set to 15. As the value of k increases, the length of diverse path set also increases, which in turn increases effective path diversity for a given

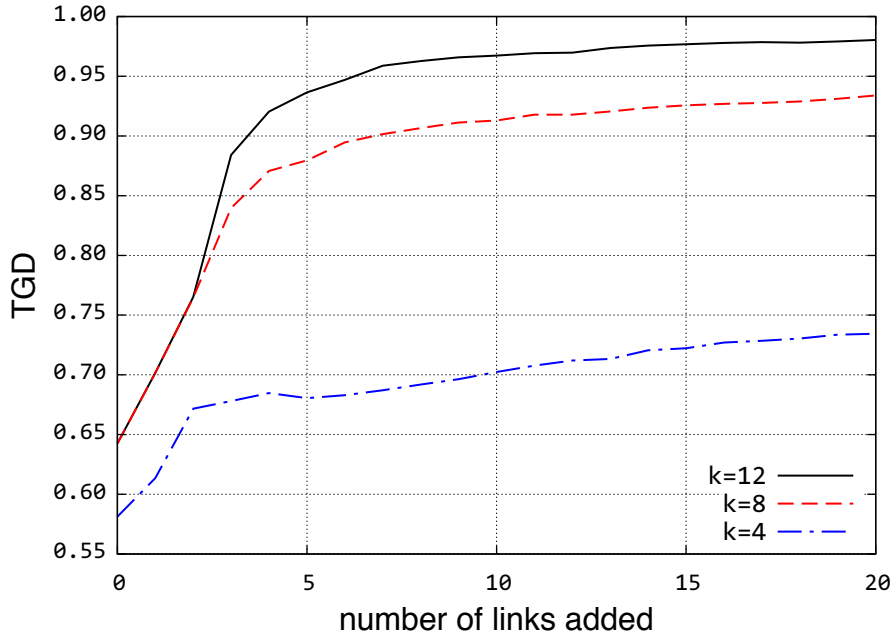

Figure 5.16: Internet2 TGD improvement

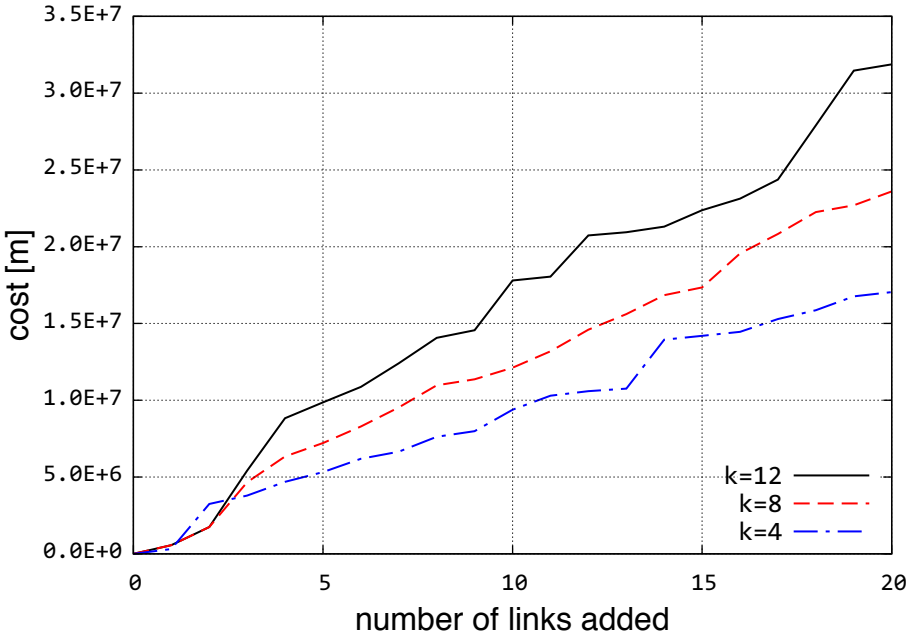

Figure 5.17: Internet2 cost incurred

pair of nodes. Consequently, as the value of k increases, the corresponding TGD increases as shown in Figure 5.16 for Internet2 topology. However, the length of diverse paths set

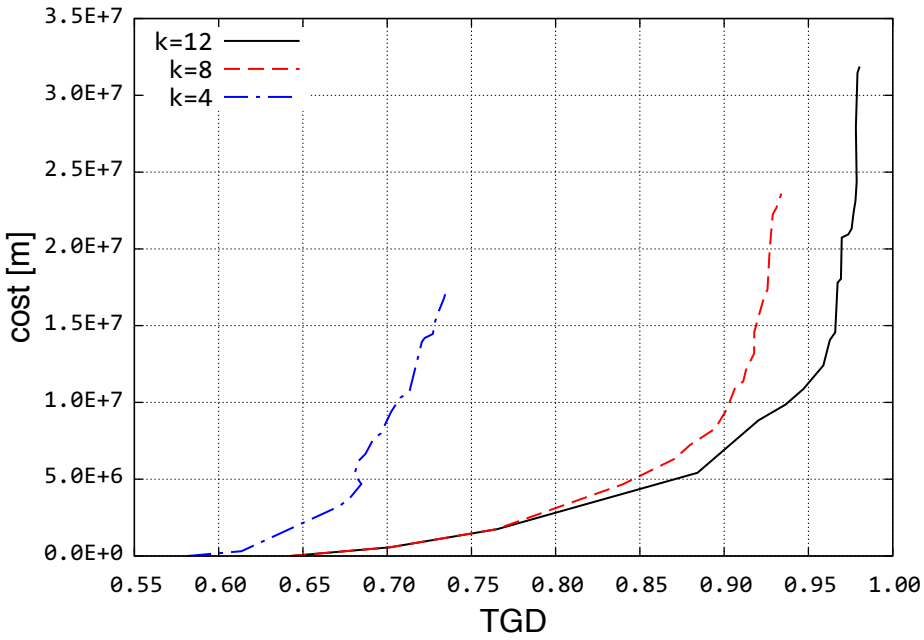

Figure 5.18: Internet2 cost and TGD

is actually m, which does not increase as the maximum diversity of the remaining paths is zero as mentioned in Section 2.3.2. For this reason, the two k values 8 and 12 have similar outcomes as depicted in Figure 5.16. The cost does not follow a pattern as the value increases for the same reason as h increases. Moreover, the cost incurred depends on the initial topological properties as shown in Figure 5.17. The cost needed to achieve a certain TGD for the Internet2 graph is shown in Figure 5.18, which show that as the value of k increases, the cost to achieve a certain TGD decreases in general as long as increasing the value of k actually increases the value of m.

Robustness Evaluation

In this section, we present the set of attacks used to evaluate the flow robustness (cf. Section 4.2.2) of the resulting optimised and non-optimised topologies. Then, we apply these attacks and show the results.

Graph centrality attacks

We use a graph theoretic model to attack a given graph and show how its flow robustness changes after each node removal. In this work, we use three centrality metrics: node betweenness, node closeness, and node degree. Thus, we have three attack models, in which the node with the highest centrality is removed. The node betweenness attack targets the node through which the highest number of shortest paths pass. The node closeness attack targets the closest node to all the other nodes in terms of hop count. The highest degree node attack targets the node with the highest number of neighbours. The list of removed nodes is determined adaptively for each attack model. This means the node centrality values are calculated after each node is removed and the highest is selected to be the next node to be removed. This is done repeatedly until all nodes are selected. The adaptive removal of nodes gives a more correct selection for the highest centrality than the non-adaptive removal, in which the highest targeted number of nodes are selected based on a single evaluation.

Lowest degree optimisation

For comparison purposes, we introduce an intuitive optimisation algorithm to improve the connectivity of a given graph via adding links to the smallest degree nodes. This algorithm adds one link repeatedly until a number of links request by the user is added. On each iteration, one end of the link is connected to the least degree node and the other end is connected to the next least degree node. If there are multiple least degree candidate links, the least cost link is selected to be added.

Robustness evaluation results

In this section, we show the results of applying the graph centrality attacks to path diversity optimised (PD-optimised), lowest degree optimised (LD-optimised), and non-optimised topologies. For the set of PD-optimised graphs, we choose the set generated using the hop count threshold $h = 15$ and the number of diverse path threshold $k = 12$ because both have more diverse and accurate results. For each graph, we apply the attack by removing half of its original number of nodes. The flow robustness is calculated after each node removal. The node betweenness attack has the highest negative impact on flow robustness because it targets the most vital nodes in the Internet2 graph as shown in Figures 5.19 through 5.21. The second highest negative impact on flow robustness is done by the highest closeness node attack since the target node has the highest closeness to all the other nodes in terms of hop count. The least negative impact on flow robustness comes from the highest degree node since it has a higher number of neighbors but is not necessarily used by most shortest paths.

Both PD-optimised and LD-optimised graphs are more resilient than non-optimised graphs because they have 20 additional links. For example, the total flow robustness of the PD-optimised Level 3 graph under the betweenness attack is 10.1 while it is 6.5 for the LD-optimised and 5.7 for the non-optimised graphs. Among the three provider graph analyses, the PD-optimised graphs are more resilient than the LD-optimised and non-optimised graphs for betweenness and closeness attacks. For degree-based centrality attack, LD-optimised graphs have higher flow robustness since links are added to the lowest degree nodes, which are targeted *the least* as shown in Figure 5.21. Therefore, the links connected to the lowest degree nodes using LD-optimisation contribute more to flow robustness than links added using PD-optimisation during the degree-based attack.

PD-optimised graphs have higher flow robustness in most physical-level graphs because

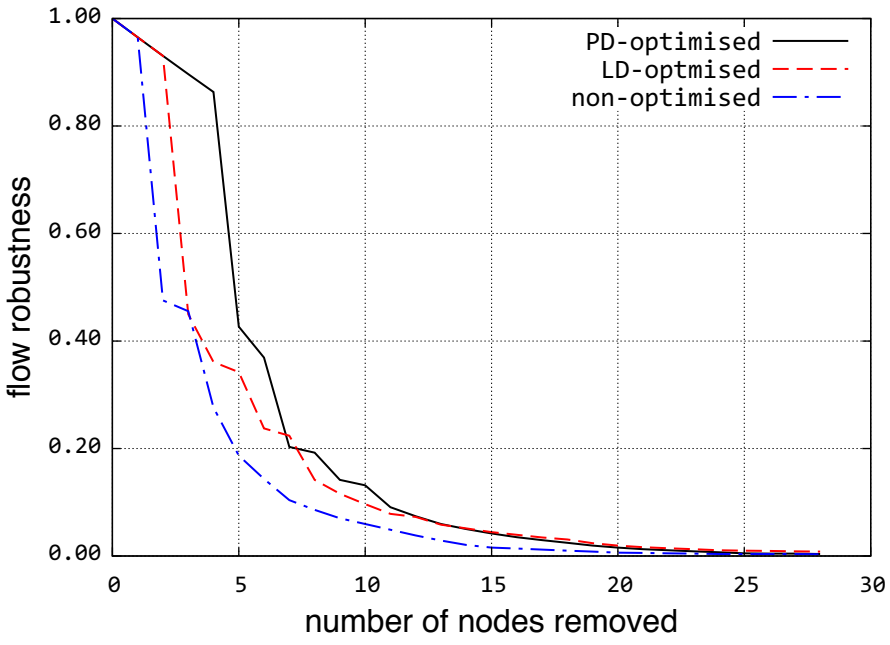

Figure 5.19: Robustness of Internet2 against betweenness-based attack

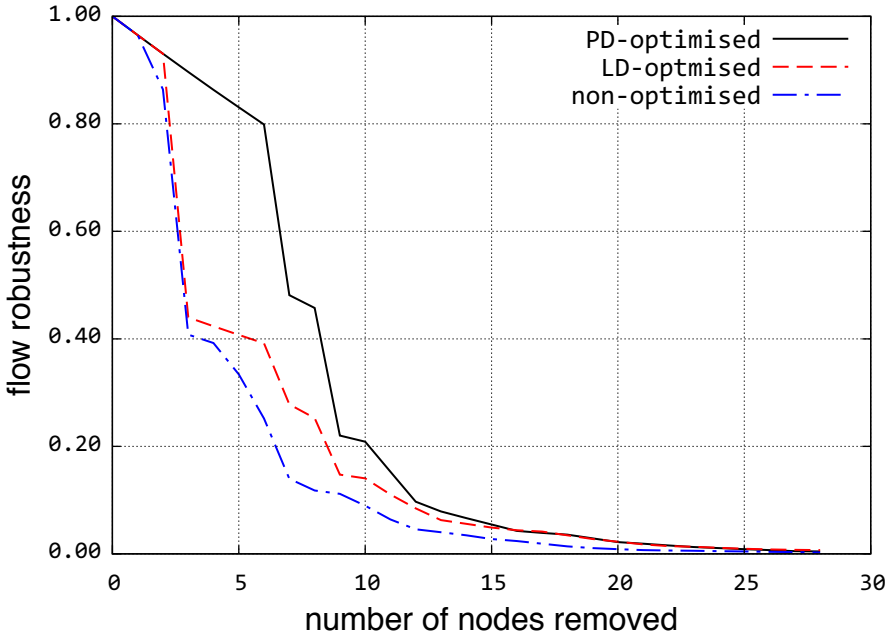

Figure 5.20: Robustness of Internet2 against closeness-based attack

in PD-optimised graphs, links are added to increase the number of diverse paths the most among all communicating nodes in the graph. Thus, when a node is removed from a PD-

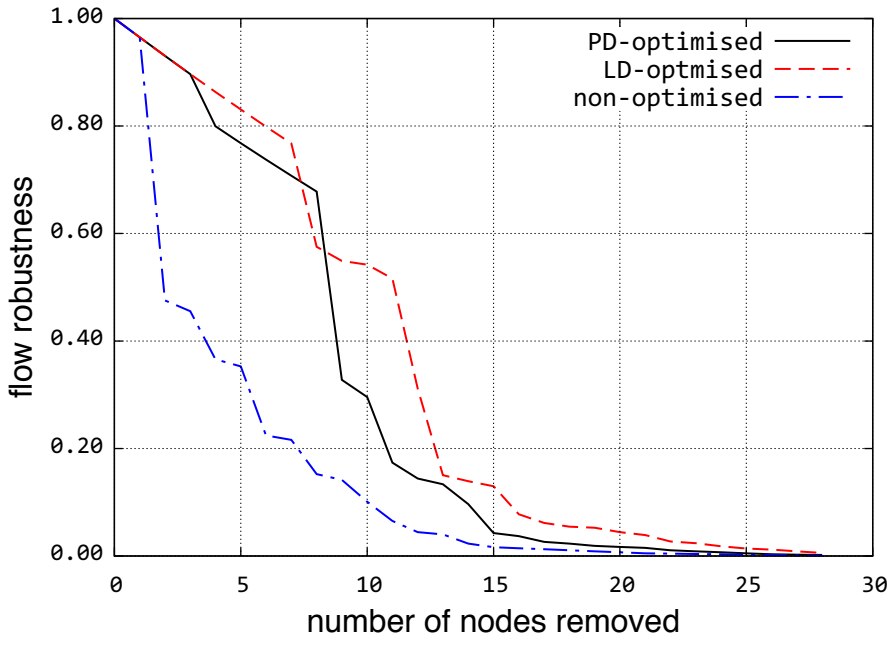

Figure 5.21: Robustness of Internet2 against degree-based attack

optimised graph, it slightly affects the other communicating nodes since they have more alternative paths to reroute their communication through. In contrast, when a node is removed from an LD-optimised graph, the other communicating nodes are more affected because they have fewer alternative paths among the communicating nodes. This is because LD-optimisation adds links based on the objective of increasing the connectivity of the lowest degree node rather than increasing the number of diverse paths.

5.3.3 Comparison of Optimisation Algorithms

In this section, the cost incurred by adding links and flow robustness against centrality-based attacks of the optimised topologies are compared. The CORONET, Internet2, and Level 3 geographic physical-level topologies are used, because the similar size and order of these graphs make them fit for a comparative evaluation while their small size allow a reduced time of calculation of graph properties. The topological characteristics

of these graphs are given in Table 5.4. 20 links are added in each physical-level topology and 50 nodes are removed to compare the flow robustness of optimised topologies. For algebraic connectivity $a(G)$ optimisation results we choose $\gamma = 0$ in which the cost weight of adding a link is ignored (i.e. maximising the $a(G)$ of the graph). We choose the hop count threshold $h = 15$ and the number of diverse paths threshold $k = 12$ for path diversity PD optimisation results, since these parameter values result in more diverse paths. For comparison of the optimisation algorithms we remove the constraint that limits the length of additional links.

Cost Comparison of Optimisation Algorithms

The cost incurred for adding 20 links for which algebraic connectivity and path diversity optimisation algorithms are applied to CORONET, Internet2, and Level 3 topologies are shown in Figures 5.22, 5.23, and 5.24, respectively.

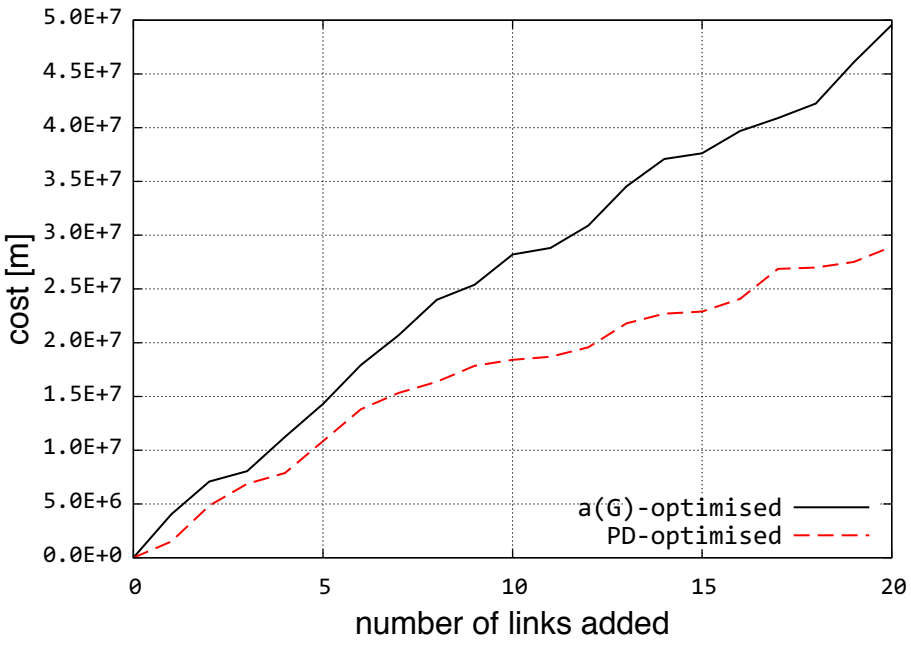

Figure 5.22: Cost comparison of graph optimisation algorithms for CORONET

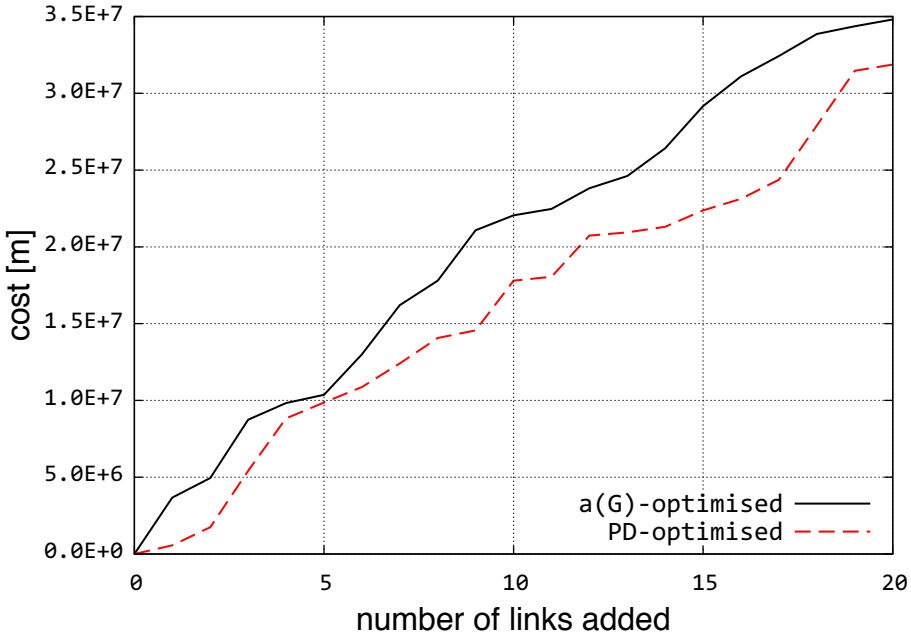

Figure 5.23: Cost comparison of graph optimisation algorithms for Internet2

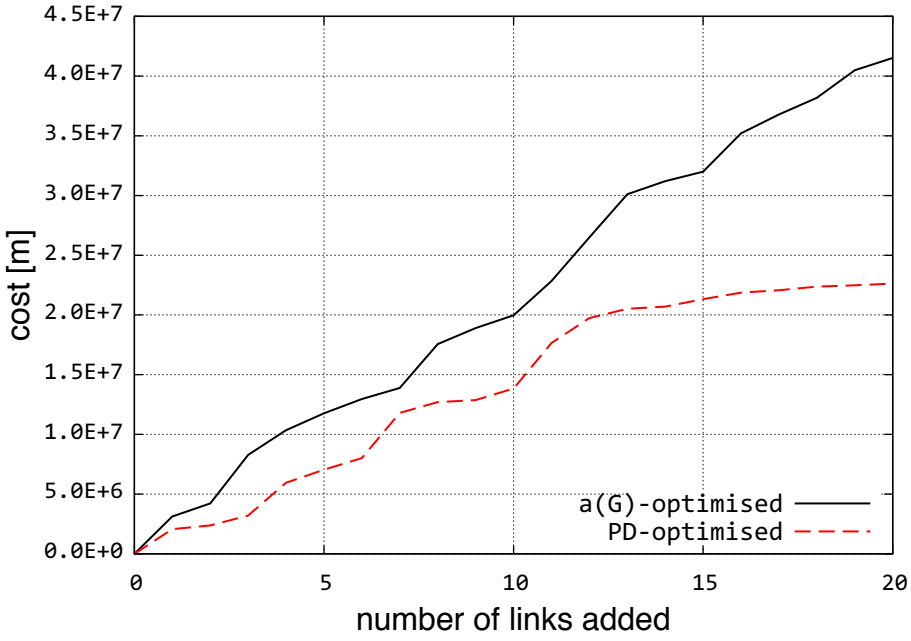

Figure 5.24: Cost comparison of graph optimisation algorithms for Level 3

For the CORONET topology, the $a(G)$ optimisation of 20 link additions incurs a cost of 5,000 km and PD optimisation of 20 link additions incurs a cost of 2,800 km. For

the Internet2 topology, the $a(G)$ optimisation of 20 link additions incurs a cost of 3,500 km and PD optimisation of 20 link additions incurs a cost of 3,200 km. For the Level 3 topology, the $a(G)$ optimisation of 20 link additions incurs a cost of 4,100 km and PD optimisation of 20 link additions incurs a cost of 2,300 km. By observing these three optimised topologies for the parameters used, the $a(G)$ optimisation requires more cost than the PD optimisation algorithm for adding 20 links. In other words, with the selected parameters, the $a(G)$ optimisation adds links that are longer compared to the links added via the PD optimisation.

Flow Robustness Comparison of Optimisation Algorithms

The flow robustness (explained in Section 4.2.2) of $a(G)$- and PD-optimised topologies of CORONET, Internet2, and Level 3 when subjected to betweenness-, closeness-, and degree-based attacks are shown in Figures 5.25, 5.26, and 5.27, respectively.

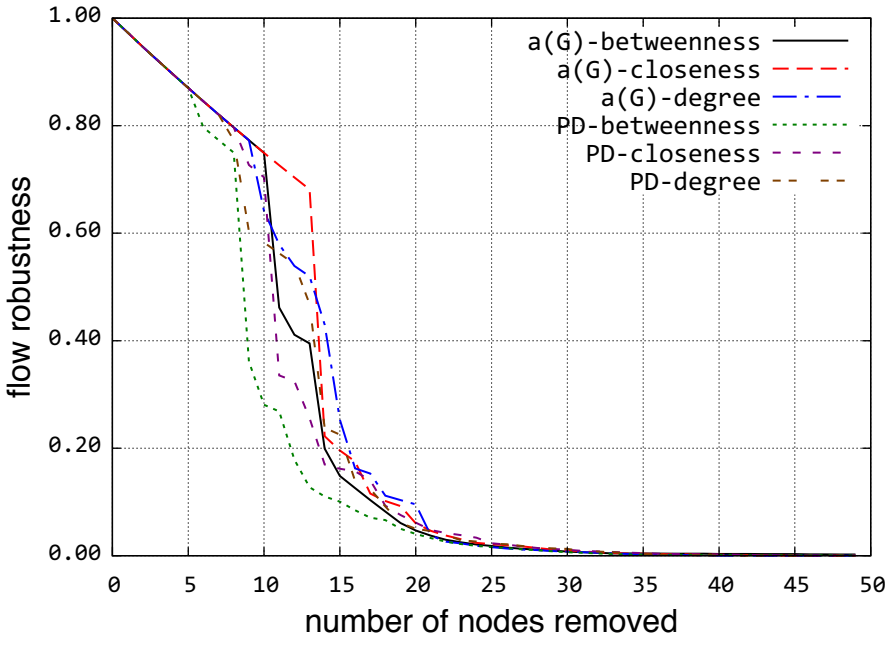

Figure 5.25: Robustness comparison of graph optimisation algorithms for CORONET

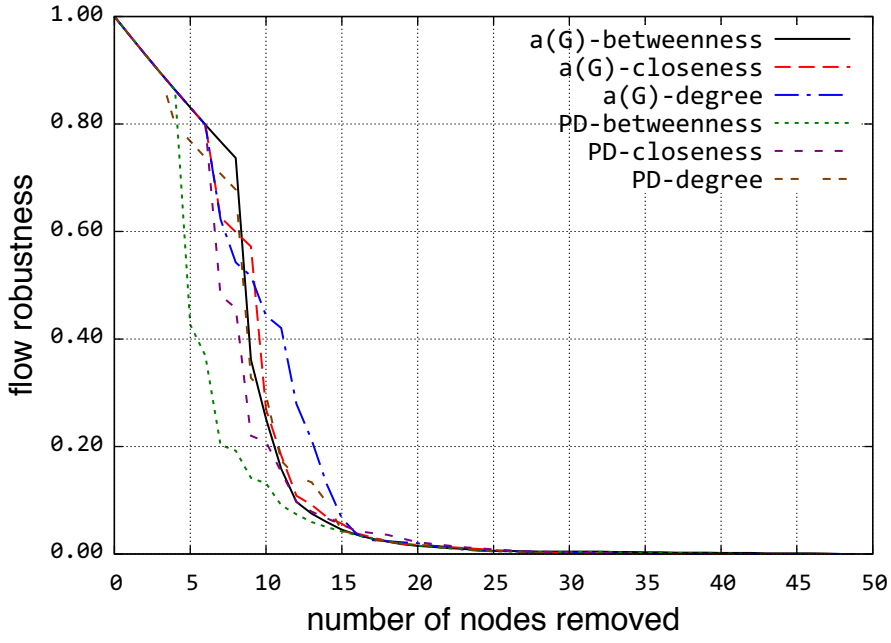

Figure 5.26: Robustness comparison of graph optimisation algorithms for Internet2

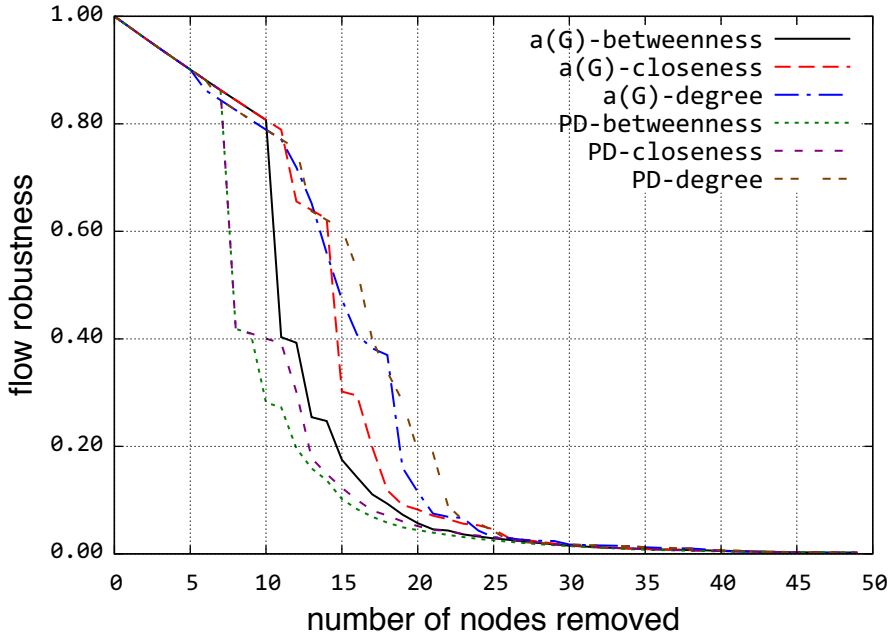

Figure 5.27: Robustness comparison of graph optimisation algorithms for Level 3

Generally, the $a(G)$-optimised topologies have a higher value of flow robustness compared to the PD-optimised topologies. From these flow robustness comparison plots, we can also

infer that betweenness-based attacks result in the worst damage. Moreover, closeness-based attacks inflict more damage to the optimised graphs than the degree-based attacks. We speculate that the longer links that are added by the $a(G)$ optimisation algorithm make these graphs more resilient compared to the PD-optimised graphs when subjected to centrality-based attacks.

5.4 Summary

In this chapter, we present two heuristic algorithms that add links that most greatly increase the connectivity of a graph, especially in terms of algebraic connectivity and path diversity metrics. The parameterised algorithms enable a designer to fine-tune the connectivity of a network based on the available budget. Furthermore, the comparison of the two algorithms shows the trade-offs between cost and resilience when designing the networks against challenges. In the next chapter, we present network resilience evaluation methodology we have developed and the evaluation of algorithms on the GpENI Future Internet testbed via experimentation.

Chapter 6

Network Resilience Evaluation

Understanding network challenges and their impact can help us to optimise existing networks and improve the design of future networks; therefore, it is imperative to have a framework and methodology to study them. We cannot thoroughly study the effects of challenges in live networks without impacting users. Testbeds are useful, but do not provide the scope and scale necessary to understand the resilience of large, complex networks, although progress is being made in this direction [48, 49]. Simulations arguably provide the best compromise between tractability and realism to study challenges; however, this is nontrivial [50].

Resilience evaluation of networks is necessary to improve the existing networks and to design better ones. In Section 2.4 we presented background information on analytical models, simulation tools, and experimentation testbeds used for resilience evaluation. In this chapter, we present the KU-CSM (KU Challenge Simulation Module) framework to evaluate network dependability and performability in the face of various challenges. The initial version of the KU-CSM framework [18] is substantially improved to evaluate the resilience of challenged networks. We use a simulation-based approach to analyse the effects of perturbations to the normal operation of networks. This framework can simulate challenges on logical or physical topologies with realistic node coördinates using

the ns-3 discrete event simulator. The framework models challenges that can be static or dynamic and can evolve temporally and spatially. Moreover, we present experiments validating the optimisation algorithms presented in Chapter 5. The experiments are run on the GpENI PlanetLab nodes using the tinc tunneling software.

The work presented in this chapter has resulted in several publications. We developed the ns-3 simulation models of attacks, random failures, and correlated failures [20]. We showed the differences between modelling the logical router-level and physical topologies [23]. The type of experiments GpENI supports were presented [49] and experiments validating optimisation algorithm results were shown [546]. The rest of this chapter is organised as follows: We present the simulation framework of a variety of challenges in Section 6.1. GpENI testbed overview and validation of graph opimisation algorithms is presented in Section 6.2. We conclude this chapter in Section 6.3.

6.1 Simulation Framework

In this section, we present our simulation framework to evaluate the resilience of network topologies when subject to a variety of challenges. The challenge simulation models are developed in the ns-3 [547] network simulator. Network configuration and challenge specification files are fed to our preprocessor that is the input to an ns-3 simulation.

6.1.1 Methodology Overview

Simulation via abstraction is one of the techniques to analyse networks in a cost-effective manner. We have chosen ns-3 [547] since it is open source, flexible, provides mixed wired and wireless capability (unlike ns-2 [548]), and the models can be extended. Unfortunately, the simulation model space increases multiplicatively with the different number of challenges and network topologies being simulated. Hence, for n different topologies

subjected to c different challenges, $n \times c$ models must be generated and simulated. Our framework decouples the challenge generation from topologies by providing a comprehensive challenge specification framework, thereby reducing the simulation model space to n network + c challenge models. We have created an automated simulation model generator that will combine any recognised challenge specifications with any provided topology, thus increasing the efficiency of simulation generation. Our simulation framework consists of four distinct steps as shown in Figure 6.1.

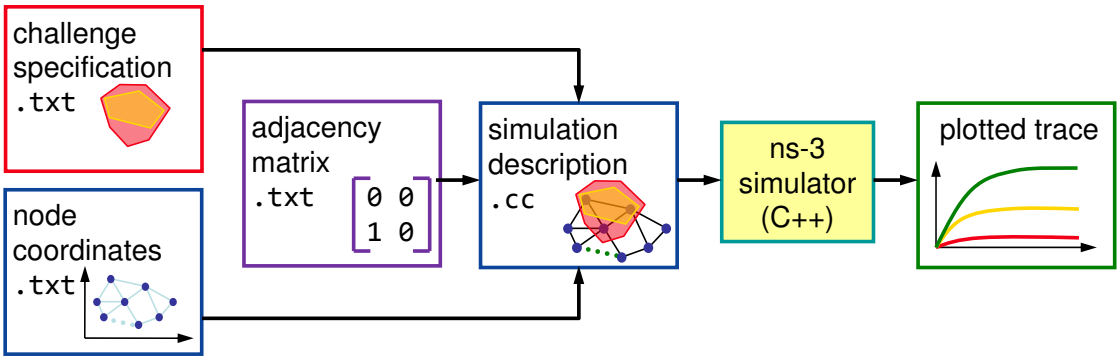

Figure 6.1: KU-CSM framework flow diagram

The first step is to provide a challenge specification that includes the type of the challenge and configuration of the challenge scenario. The second step is to provide a description of the network topology, consisting of node geographical or logical coördinates and an adjacency matrix. The third step is the automated generation of ns-3 simulation C++ code based on the topology and challenge descriptor. Finally, we run the simulations and analyse the network performance throughout the challenge scenario. Additionally, the simulation framework can also be enabled to generate ns-3 network animator (NetAnim) traces for visualisation purposes. A NetAnim screenshot of the Rocketfuel [33] based Sprint backbone network topology of 27 nodes and 68 links is shown in Figure 6.2. We have provided partial simulation code to ns-3 community that automates generation of topologies based on an adjacency matrix and node coördinates [549], which has been

incorporated to the ns-3.10 standard release.

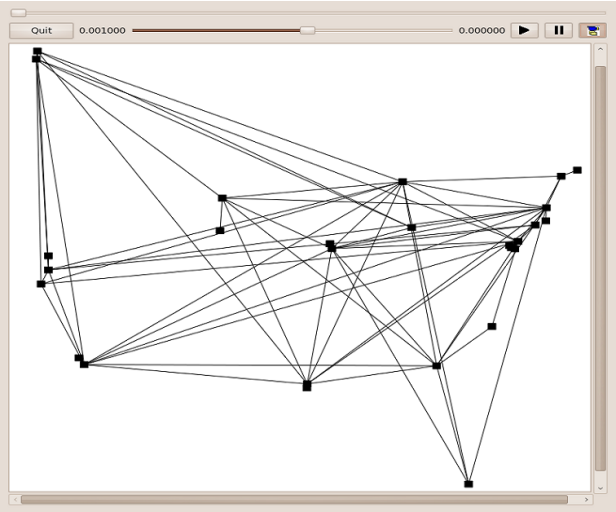

Figure 6.2: NetAnim screen shot of inferred Sprint topology

6.1.2 Implementation of Challenge Models

In the following subsections we present the details of implementation of challenge models (cf. Chapter 3) in the ns-3 discrete event simulator.

Non-malicious Challenges

In the case of wired domain challenges in this category, the number of nodes or links k subject to random failure during a challenge period (t_1, t_2) is listed in the challenge specification file. Nodes or links are shut down for the duration of the challenge if the probability of failure of that node or link is greater than the probabilistic failure rate threshold p_{r} provided as a parameter. This type of challenge models random node and link failures that are uncorrelated with respect to topology and geography.

Malicious Attacks

Malicious attacks result from the exploitation of structural knowledge of the network by an attacker who wishes to inflict maximum damage with limited resources. We use topological properties of the graph in order to determine the *critical* elements in the network, based on measures such as the degree of connectivity of nodes, and betweenness of nodes and links (betweenness is the number of shortest paths through a particular element [513, 515]). The critical nodes or links are shut down for the duration of the challenge period (t_1, t_2).

Large-scale Disasters

The challenge specification for area-based challenges resulting from large-scale disasters is an n-sided polygon with vertices located at a particular set of geographic coördinates (x_i, y_i) or a circle centered at specified coördinates (x_c, y_c) with radius r. The simulation framework then determines the nodes and links that are encompassed by the polygon or circle, and disables them during the challenge interval. We use the Computational Geometry Algorithms Library (CGAL) [550], which is an open source library with efficient geometric algorithms implemented in C++. We also implement dynamic area-based challenges, in which the challenge area can evolve in shape over time: scale (expand or contract), rotate, and move on a trajectory during the simulation. Large-scale regional failure scenarios previously only have been modelled as a static circle [138] for evaluating the performance of path restoration after a failure. Examples of the need to simulate arbitrary polygons are to model large-scale power blackouts, EMP weapons [215], coronal mass ejections [211], and large-scale natural disasters such as hurricanes and tsunami.

6.1.3 Network Challenge Simulations

In this section we present our results that demonstrates the utility of our approach. We first evaluate the perturbations to logical topologies. Next, we present evaluation of physical topologies when faced by correlated challenges, which is a *necessary* condition for evaluation of networks. We use ns-3.7.1 release and the simulation parameters are as follows: The network is composed of bidirectional wired links with 10 Mb/s bandwidth and 2 ms transmission delay. Routing is accomplished using the Dijkstra shortest-path-first algorithm, recalculated at each time step, with reconvergence delay as a simulation parameter. The traffic is constant bit rate (CBR) at 40 kb/s between every node pair, with 1000 Byte packets. These parameters are chosen such that there is no congestion under normal operation, but the network is not significantly over-provisioned so that we will see the effect of node and link failures. We measure the network's aggregate performance under challenges in terms of aggregate packet delivery ratio (PDR).

Non-malicious and Malicious Challenges

First, we evaluate the performance of three separate topologies under the presence of malicious and non-malicious challenges. The topologies we choose are the Sprint logical topology based on the Rocketfuel map [33] (Figure 6.3) and two synthetic topologies (Figure 6.4 and 6.5). The synthetic topologies are generated using the KU-LoCGen topology generation tool [5, 536, 551]. KU-LoCGen generates topologies with geographic constraints and places links between nodes using various models; in this case the modified Waxman [545] model. The resulting synthetic topologies have the same number of nodes at the same geographic locations as the inferred Sprint topology, however the number of links and connectivity of the nodes differ. The two synthetic graphs chosen for this work consist of a richly connected and poorly connected topology to demonstrate the range

of robustness results from this simulation framework. The graph characteristics of three topologies[1] are presented in Table 6.1.

Figure 6.3: Sprint inferred topology

Figure 6.4: Synthetic topology 1

We evaluate the performance of the sample topologies under the presence of malicious and non-malicious challenges with the PDR of the network shown in Figures 6.6, 6.7, and 6.8

[1]The topological data used in this chapter is slightly outdated compared to the data presented in Table 4.2, as the work presented in this chapter was performed earlier. However, this does not change the methodology presented. The topological data presented earlier contains minor modification to reflect correct topology.

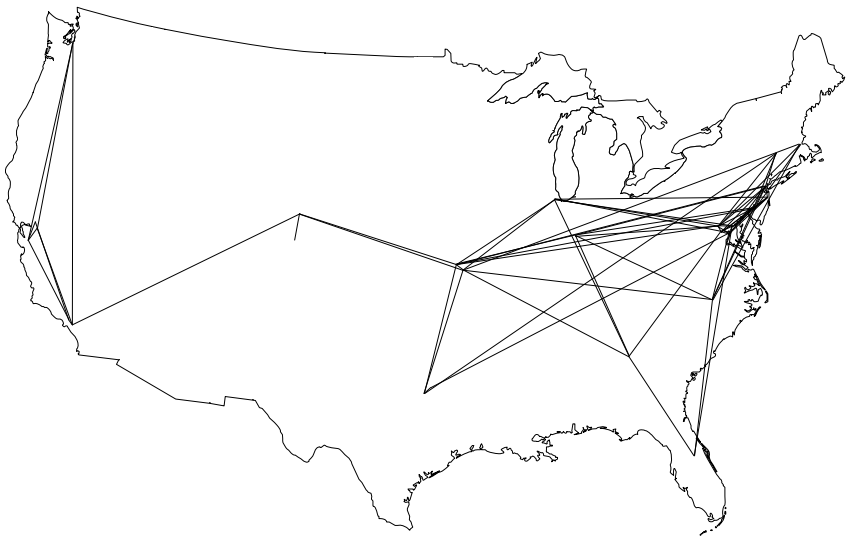

Figure 6.5: Synthetic topology 2

Table 6.1: Topological characteristics of sample networks

Network Topology	**Sprint**	**Synthetic 1**	**Synthetic 2**
number of nodes	27	27	27
number of edges	68	74	68
maximum degree	12	9	10
average degree	5.04	5.5	5.04
clustering coeff.	0.43	0.29	0.38
network diameter	6	4	6
average hopcount	2.4	2.2	2.9
max. node betweenness	144	76	302
max. link betweenness	72	31	140

for link failures and in Figures 6.9, 6.10, and 6.11 for node failures with up to 10 links or nodes down. We measure the instantaneous PDR at the steady-state condition during the challenges for each point. We also note that for random failures, we average the results over 100 runs. For malicious challenges (betweenness or degree of connectivity), first we calculate the betweenness (or degree of connectivity) for each network element in the topology, and provide the challenge file as the list of the elements to be brought down in order as a function of the x-axis.

Figures 6.6, 6.7, and 6.8 shows the PDR during the link perturbations to Sprint inferred

(Figure 6.3), synthetic 1 (Figure 6.4), and synthetic 2 (Figure 6.5) topologies respectively. We evaluate the PDR during link failures for two cases: 10 random link failures and an attack using the 10 highest-ranked links based on link betweenness values. Except for the synthetic topology 1, intelligent link attacks have a more degrading impact than the random failures. The PDR of 100% for both random and attack cases for the synthetic 1 topology (Figure 6.7) can be attributed to this topology's lower average hop count, network diameter, clustering coefficient, and higher average degree. The synthetic topology 1 also has six more links compared to the other two topologies: 74 vs. 68. On the other hand, the link attack on highest betweenness link for synthetic topology 2 results in a PDR drop to 60%. Visual inspection of synthetic topology 2 (Figure 6.5) clearly identifies the link failure between the central and west US is the cause of this since the network partitions. We can also infer the same conclusion by examining the link betweenness of synthetic topology 2 in Table 6.1, in which this link has 140 shortest paths.

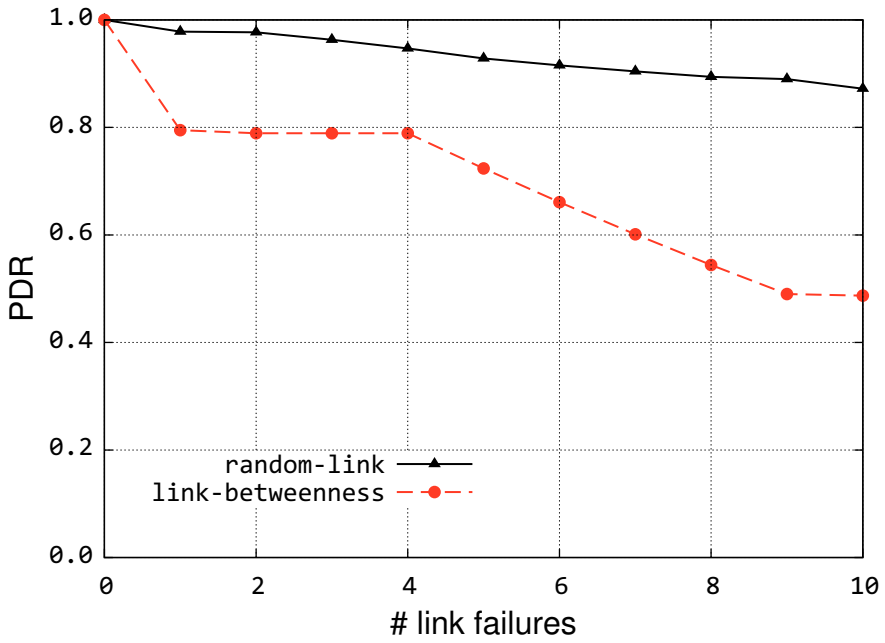

Figure 6.6: PDR during link perturbations for Sprint inferred topology

The performance of sample topologies against malicious and non-malicious node pertur-

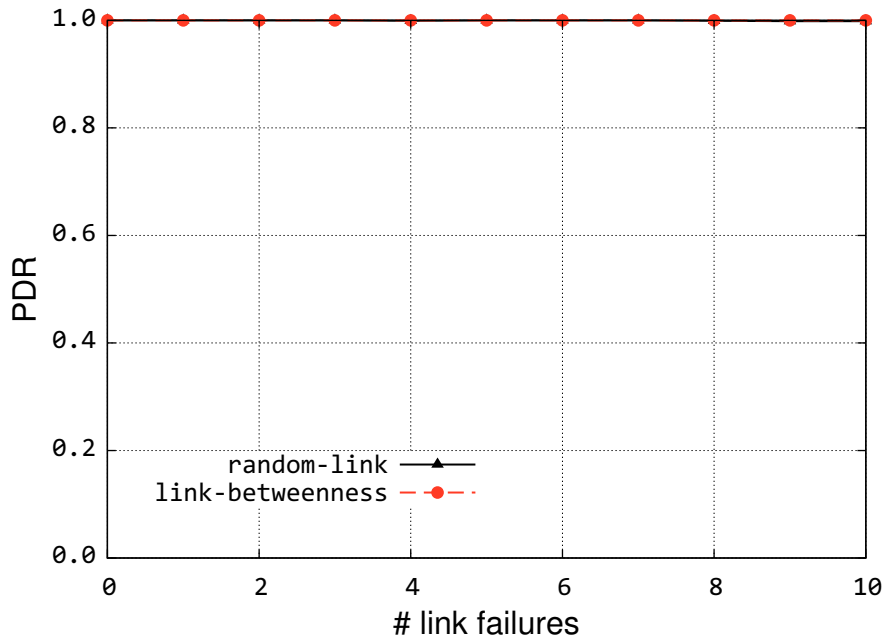

Figure 6.7: PDR during link perturbations for synthetic topology 1

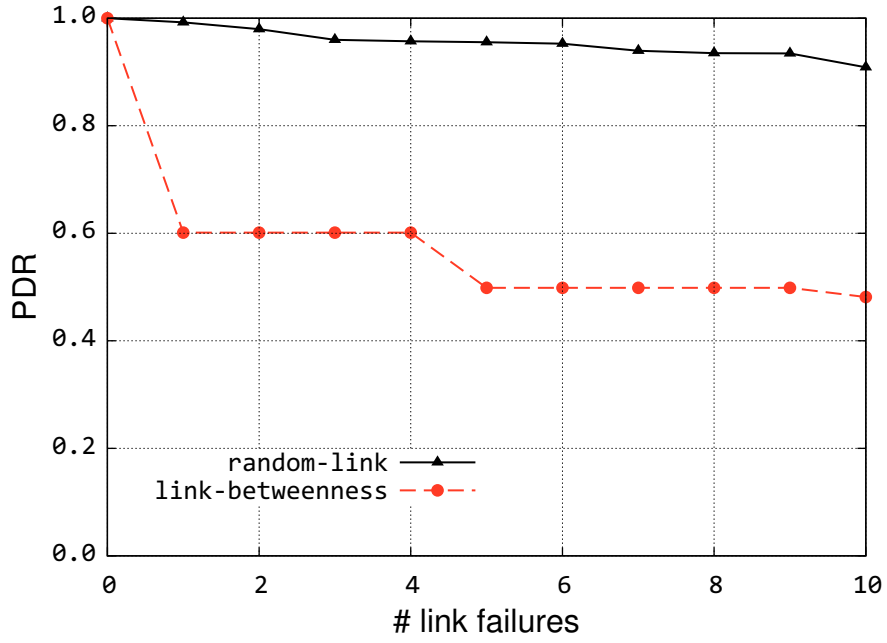

Figure 6.8: PDR during link perturbations for synthetic topology 2

bations is shown in Figures 6.9, 6.10, and 6.11. We evaluate the PDR during node failures for three cases: 10 random node failures, attack of the 10 highest ranked nodes based on

betweenness, and attack of the 10 highest ranked nodes based on degree of connectivity. Figures 6.9, 6.10, and 6.11 show that node failures are worse than link attacks or failures (compared to Figures 6.6, 6.7, and 6.8), since each node failure is the equivalent of the failure of all links incident to that node. Our results indicate that attacks launched with knowledge of the network topology can cause the most severe degradation. We can also infer the tradeoff between robustness and the cost of building topologies using our framework. It should be noted that KU-LoCGen performs topology generation under cost constraints of a fixed and variable cost of each link, and thus we can compare the resilience of various cost points, with increasing cost providing increasing resilience due to better network connectivity when there are more links.

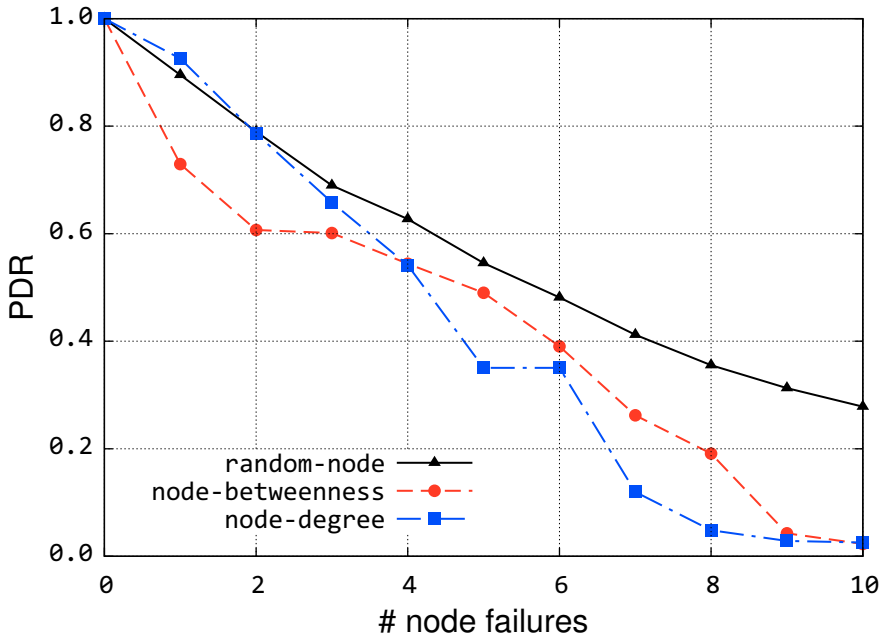

Figure 6.9: PDR during node perturbations for Sprint inferred topology

The performance evaluation of sample networks with varying failure probabilities is shown in Figure 6.12 and Figure 6.13. We averaged 100 simulation runs for the probabilistic failure scenarios. The seed to the random number generator is generated via the system clock, therefore we used a different seed for each run. The random variables used in the

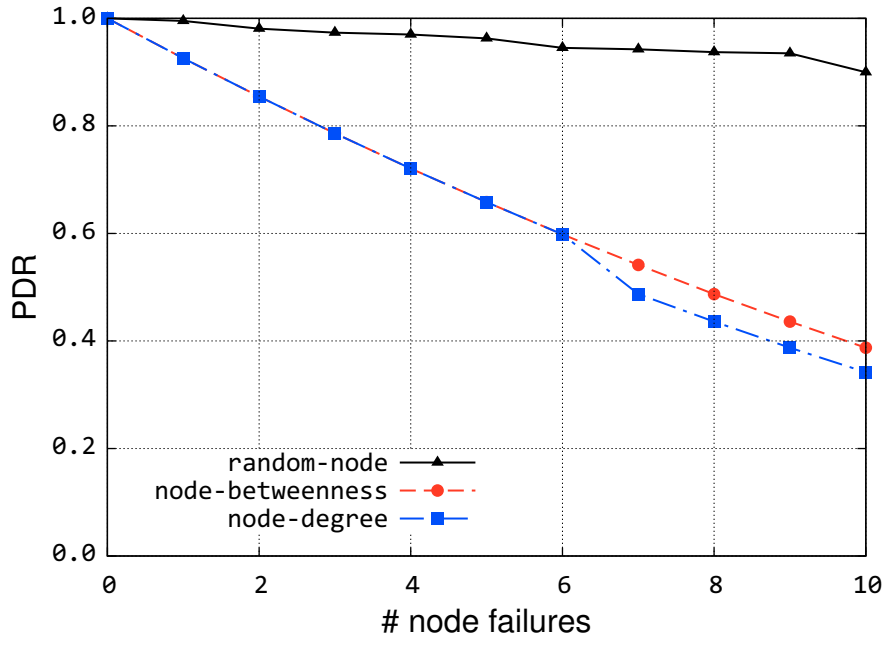

Figure 6.10: PDR during node perturbations for synthetic topology 1

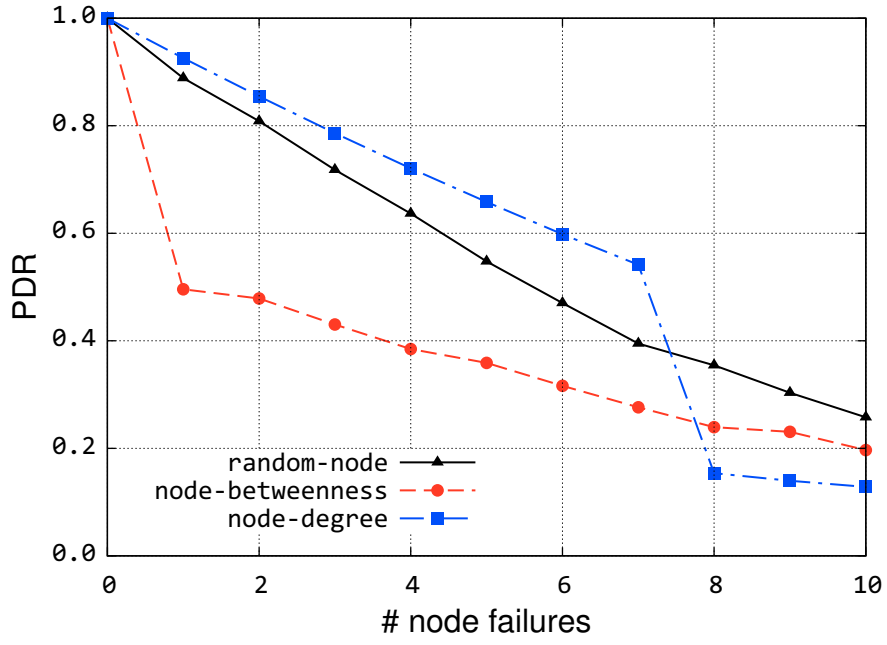

Figure 6.11: PDR during node perturbations for synthetic topology 2

simulations were uniformly distributed. In these scenarios, PDR is calculated when the network elements are in the down state. The state transition for each element occurs

if the probability of failure of a network element is greater than the specified value p_r provided in the challenge specification file.

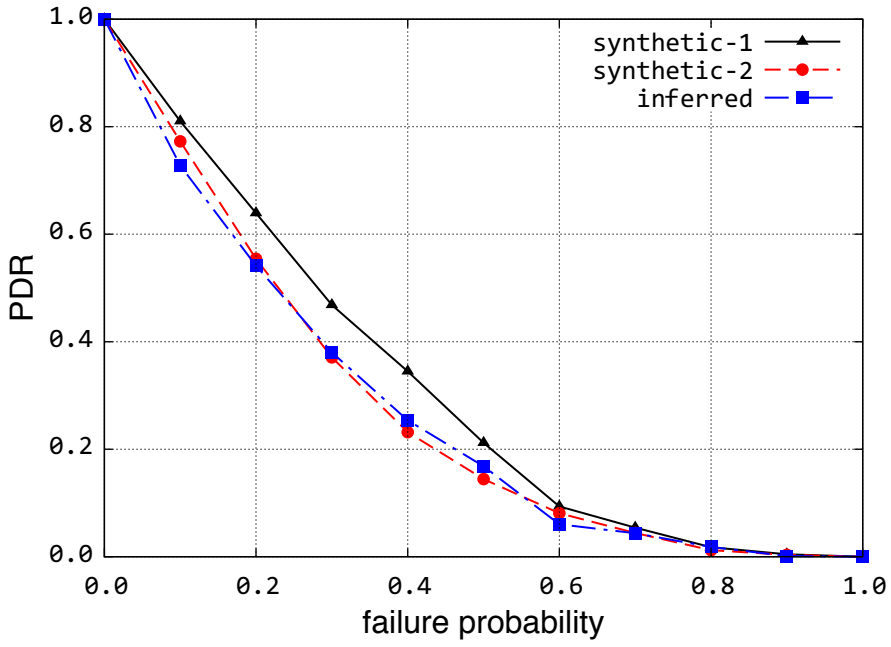

Figure 6.12: PDR during statistical node failures

The performance of the sample networks with increasing probabilistic node failure is shown in Figure 6.12. The PDR value varies between 100% and 0% as the node failure probability increases from 0% to 100%. The curves are close to each other since each sample topology has the same number of nodes, and failure probabilities are uniformly distributed. In particular, the synthetic 2 topology and Sprint inferred topologies show similar characteristics since the average degree values of those topologies are the same as listed in Table 6.1.

Figure 6.13 shows the PDR during the probabilistic link failures for synthetic 1, synthetic 2, and Sprint inferred topologies respectively. While the performance of the Sprint inferred and synthetic 2 topologies are close to each other, synthetic 1 topology has better performance for the probabilistic link failure scenario since it has more links compared to

the other two topologies. Compared to the probabilistic node failures, probabilistic link failures do not impact the networks as much, since the impact of a node failure includes one or more links being brought down.

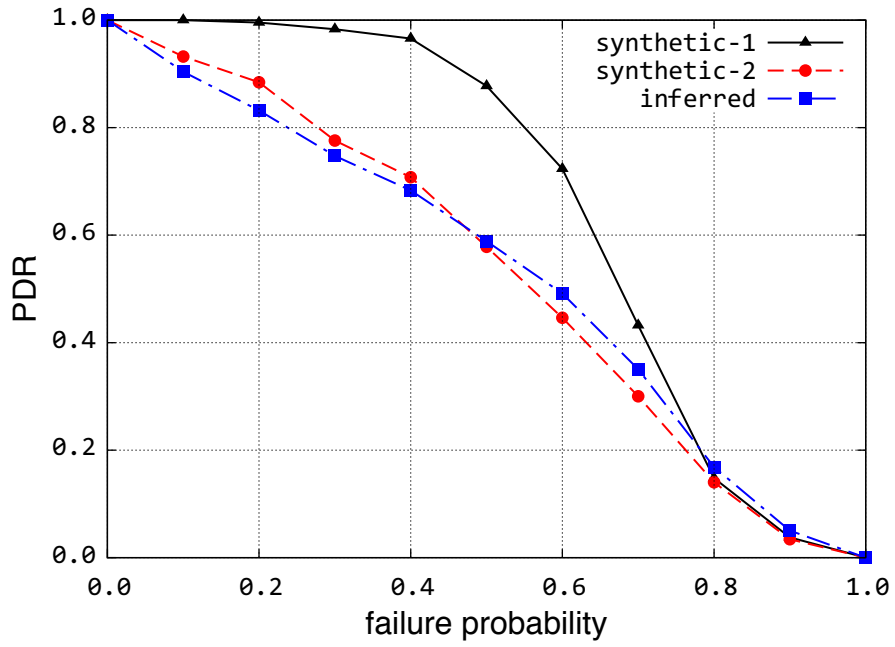

Figure 6.13: PDR during statistical link failures

Area-based Challenges

As previously discussed, our framework uses circles or polygons to model geographically correlated failures representative of large-scale disasters needed to evaluate network survivability [52, 53]. Area-based challenges in our model can be stationary or evolving in time. Next, we present the results of three scenarios that demonstrate area-based challenges that evolve spatially and temporally. In all scenarios, we use the Rocketfuel-based Sprint logical topology as shown in Figure 6.3. Application traffic is generated from 2 to 29 s. and challenge scenarios were applied from 10 until 22 s. for the performance plots.

Scaling circle:

To demonstrate a scaling circle area-based challenge scenario, we simulate a circle centered at (−74.00°, 40.71°), in New York City (NYC) as shown in Figure 6.14(a), with a radius of 1° (approximately 111 km). We choose the scenario to be representative of an electromagnetic pulse (EMP) attack [215]. The PDR is shown in Figure 6.14(b). We choose the simulation parameters such that the radius doubles in every 4 s. As can be seen, the PDR reduces as the circular area doubles. The PDR drop depends on how many nodes and links resides in the circle for each step.

Moving circle:

Next, we demonstrate an area-based scenario that can evolve spatially and temporally, such as to model a hurricane. We simulate a moving circle in a trajectory from Orlando (−81.37°, 28.53°) to NYC (−74.00°, 40.71°). Three snapshots of the evolving challenge are shown in Figure 6.15(a). The radius of the circle is kept at 2° (approximately 222 km). We choose the simulation parameters for illustration such that the circle reaches NYC in seven seconds (sped up to constrain simulation time), with route recomputation every 3 s.

As shown in Figure 6.15(b), PDR reduces to 93% as the challenge starts only covering the node in Orlando at 10 s. As the challenge moves towards NYC in its trajectory, the PDR reaches one at 13 s. In this case, the challenge area includes only the link between Orlando and NYC, but since there are multiple paths, a single link failure does not affect the PDR, showing that *diversity for survivability* is crucial [2,73]. As the challenge moves into the northeast US region at 16 s., the PDR drops to 66% as the challenge covers several nodes and links. The simulation shows that as the circle moves out of the more crowded region of the network, the PDR improves, until the challenge ends at 22 s.

Scaling polygon:

(a) Scaling circle

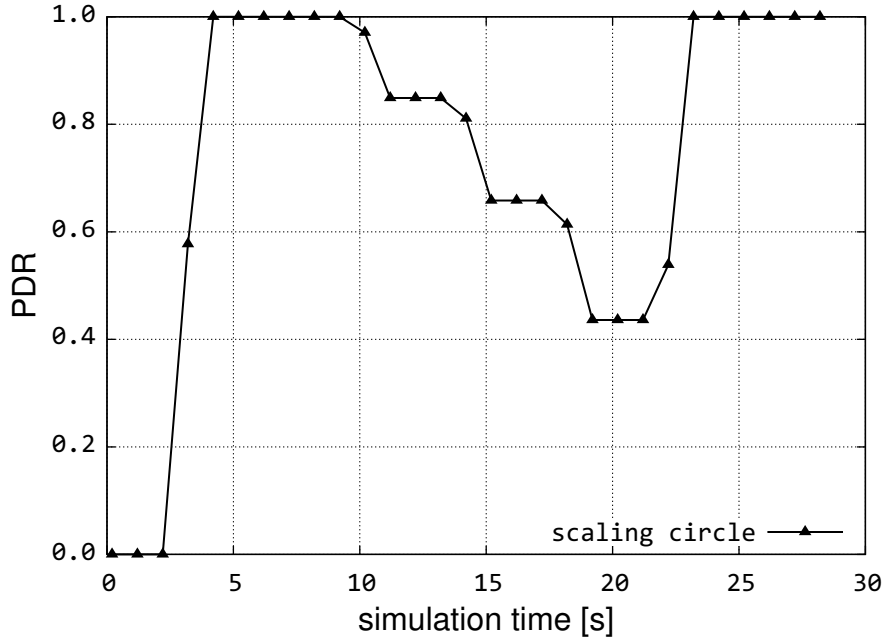

(b) Scaling circle PDR

Figure 6.14: Scaling circle challenge scenario and PDR for Sprint logical topology

Polygons are useful to model specific geographic challenges such as power failures that can cause large-scale network disruption as in the 2003 Northeast US blackout [255]. For a scaling polygon example, we show a 6-sided irregular polygon in the Midwest region of

(a) Moving circle

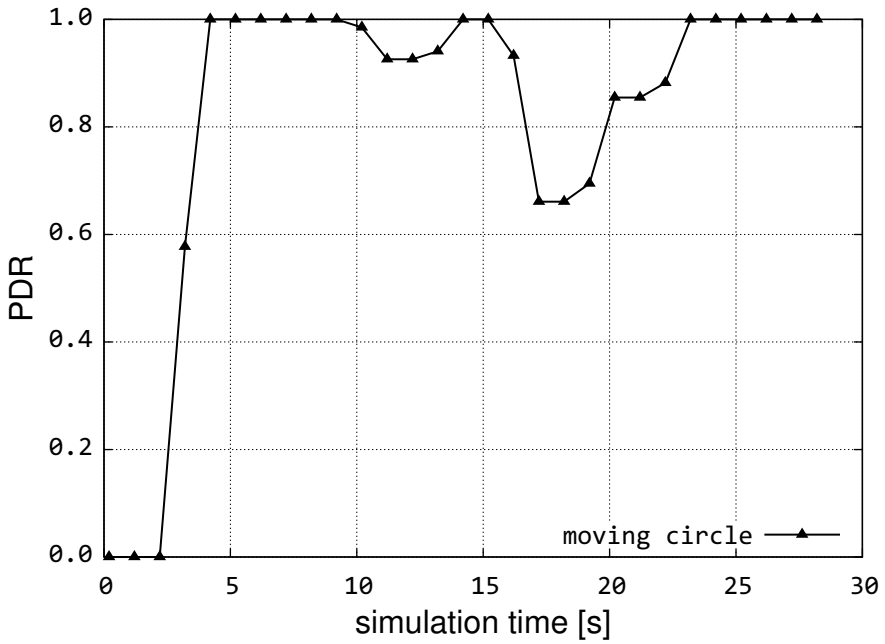

(b) Moving circle PDR

Figure 6.15: Moving circle challenge scenario and PDR for Sprint logical topology

the US, roughly representative of the North American Electric Reliability Corporation (NERC) Midwest region [215], with vertices at: [(−87.91 °, 43.04 °), (−89.09 °, 42.27 °), (−89.64 °, 39.8 °), (−88.54 °, 39.12 °), (−88.24 °, 40.12 °), (−87.65 °, 41.85 °)] as shown in

Figure 6.16(a).

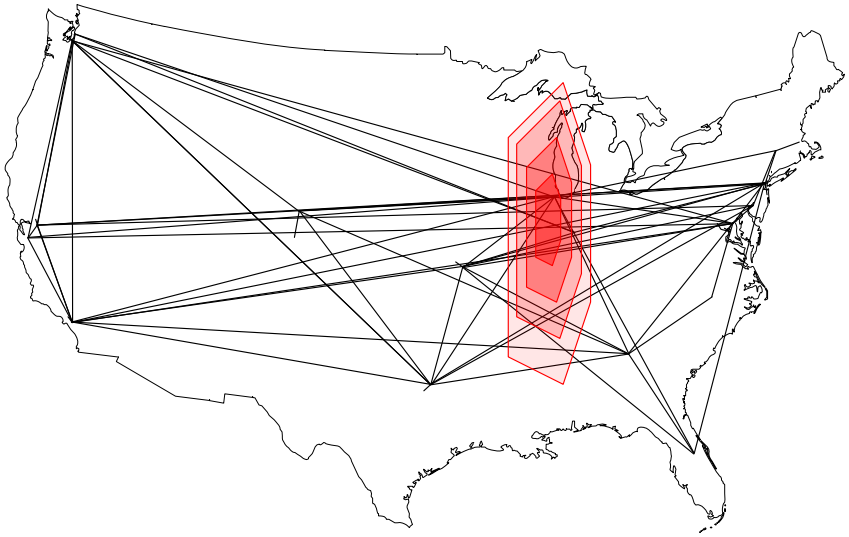

(a) Scaling polygon

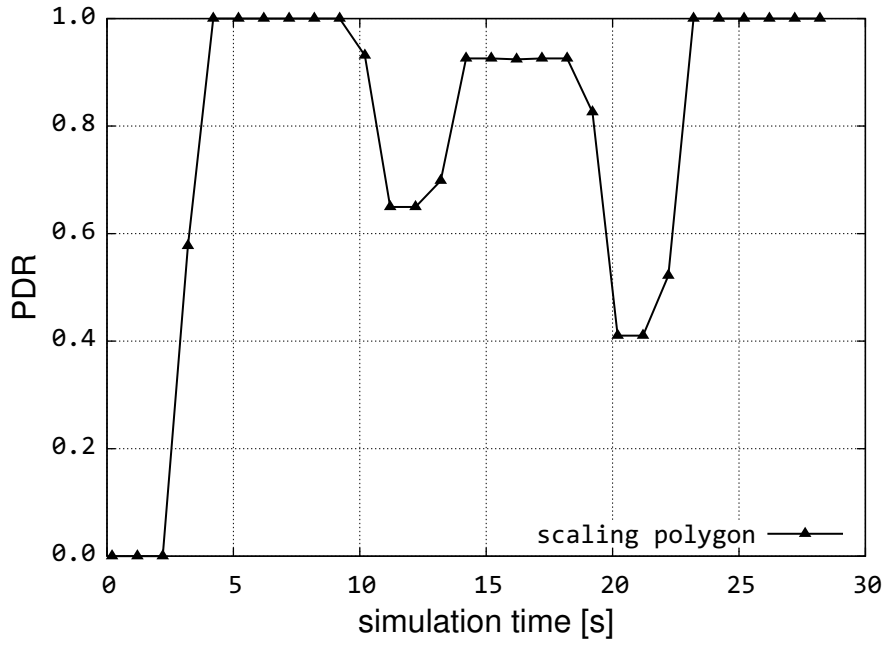

(b) Scaling polygon PDR

Figure 6.16: Scaling polygon challenge scenario and PDR for Sprint logical topology

The PDR throughout the simulation is shown in Figure 6.16(b). In this scenario, the edges of the irregular polygon increase 1.8 times every 3 s. At 10 s. the challenge affects

16 links, which causes the PDR to drop to 65%. The PDR then increases to 93%, even though more links and nodes are affected at 13 s. because of route reconvergence. As the polygon increases in size, the PDR drops to as low as 41%, because the challenge area partitions the network at 21 s. This type of scenario can be used either to understand the relationship between the area of a challenge and network performability, or to model a temporally evolving challenge, such as a cascading power failure that increases in scope over time.

Impact of Challenges on Physical Topologies

Network performance analysis under a variety of challenges is possible with this framework. We showed results of how this framework can be used on a layer 3 logical topology in Section 6.1.3 and 6.1.3. In this section we investigate the network performance of Sprint's geographical physical layer topology.

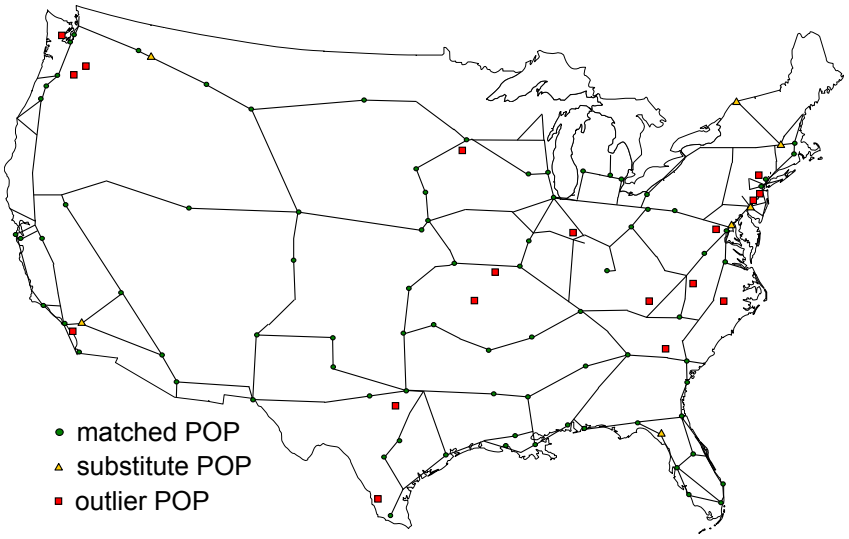

Figure 6.17: Sprint MPLS PoP locations

A fibre-optic topology does not necessarily use all nodes to be traffic sources and sinks. There can be signal regenerators, cross-connects, and ADMs. Therefore, to realistically place source and sink points we utilised the Sprint global MPLS map [552], with a total

of 115 MPLS PoPs in the US. Among these 115 PoPs, 83 exactly match to the physical topology we constructed. 7 more PoP locations closely match to a city on the physical topology. For example the fibre-optic route map has a point in Coeur d′Alene, Idaho, while the MPLS map has a PoP located in Post Falls, Idaho, which are very close to each other, so we consider the Coeur d′Alene node on the physical topology adjacency matrix a traffic source and sink point. 25 PoP locations did not match to the physical topology well. For example the MPLS PoP in Springfield, Missouri does not lie on any Sprint fibre routes. These MPLS PoPs are backhauled over other service providers, and are thus excluded from our Sprint traffic matrix. The resulting traffic matrix has 90 source/sink pairs. The Sprint fibre-optic routes with Sprint MPLS PoP locations are shown in Figure 6.17.

Challenge Simulations on Physical Topologies:

The physical topology has 245 cities of which 90 MPLS PoP locations match to the cities on the physical topology. Since not all cities are traffic source or sinks, the statistical failure scenarios would not be useful determining the performability of the network. Hence, we focus on simulating area-based challenges against the physical topologies representing large-scale disasters. We run the same area-based scenarios on the physical topology that we ran on the Sprint logical topology (Section 6.1.3) as depicted in Figures 6.18(a), 6.19(a), and 6.20(a).

The performance of the physical topology is shown Figures 6.18(b), 6.19(b), and 6.20(b). The characteristics of the performance curves closely match between physical and logical topologies for the same area-based challenge scenarios. The difference is the PDR values. This is expected since the number of traffic sources and the sinks differ between each topology. We also increase the link bandwidth from 10 Mb/s to 100 Mb/s to prevent artificial drop of packets in the physical topology scenarios, since the maximum link

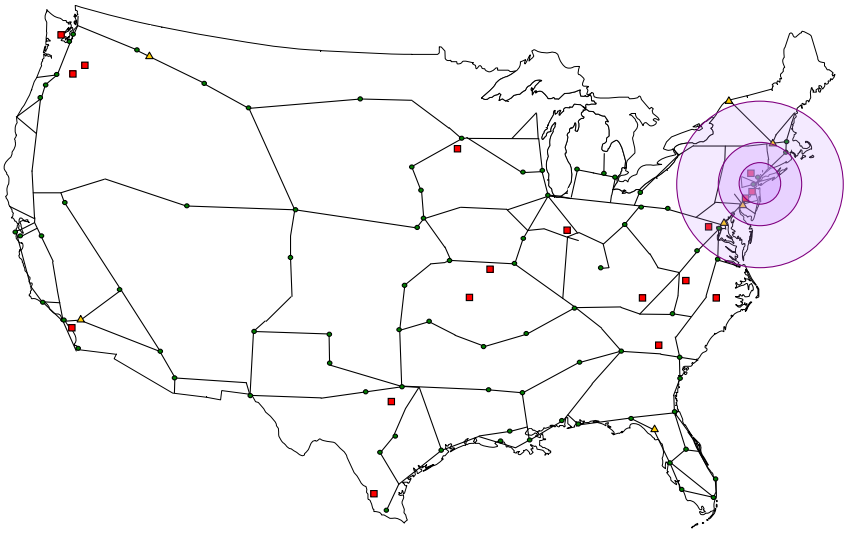

(a) Scaling circle

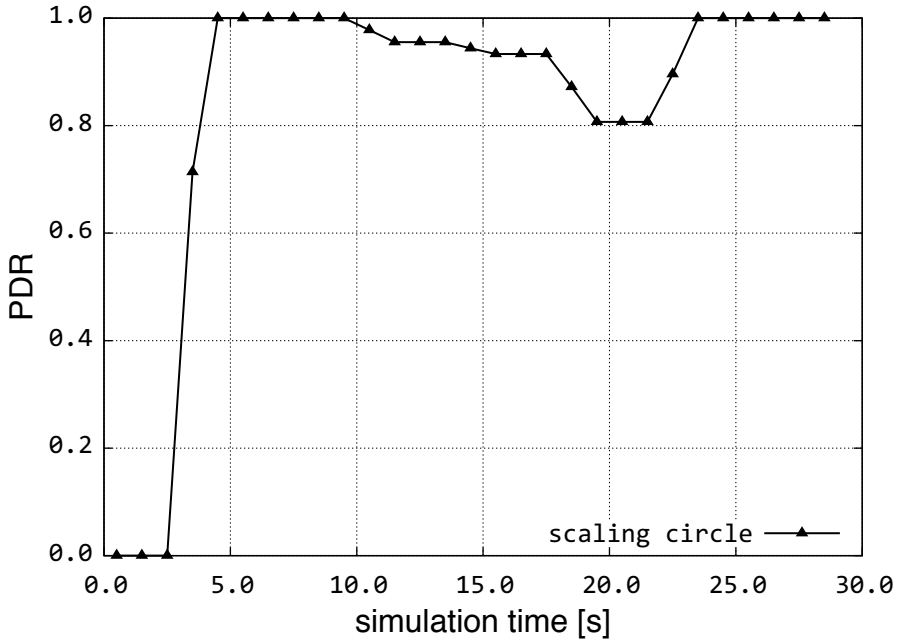

(b) Scaling circle PDR

Figure 6.18: Scaling circle challenge scenario and PDR for Sprint physical topology

betweenness in the physical topology is 8,012 and the maximum link betweenness on the Sprint logical topology is 72.

Next, we demonstrate an area-based scenario representative of a hurricane hitting south

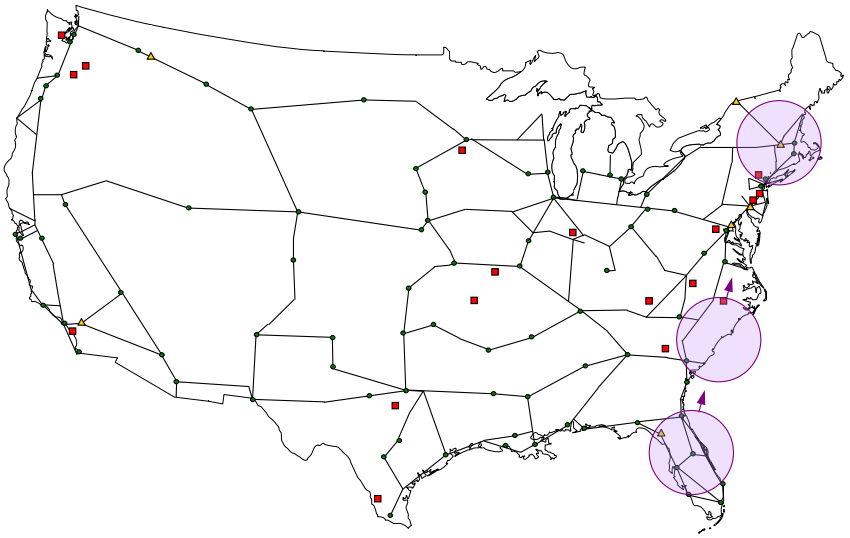

(a) Moving circle

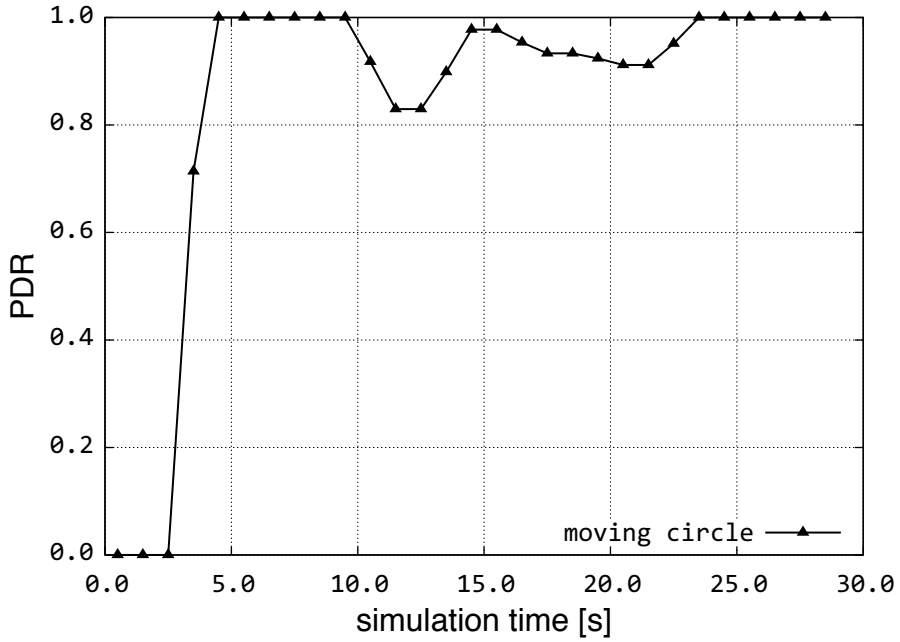

(b) Moving circle PDR

Figure 6.19: Moving circle challenge scenario and PDR for Sprint physical topology

central US as shown in Figure 6.21. In the smallest area are the nodes in New Orleans and Biloxi of which only the New Orleans node is a MPLS PoP node. In the second circular area challenge, the nodes are: New Orleans, Baton Rouge, Lafayette, Biloxi, and

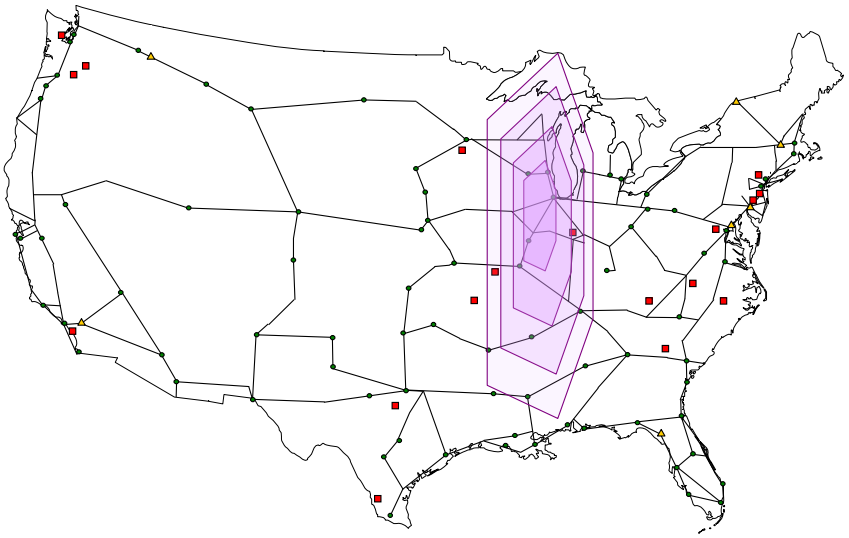

(a) Scaling polygon

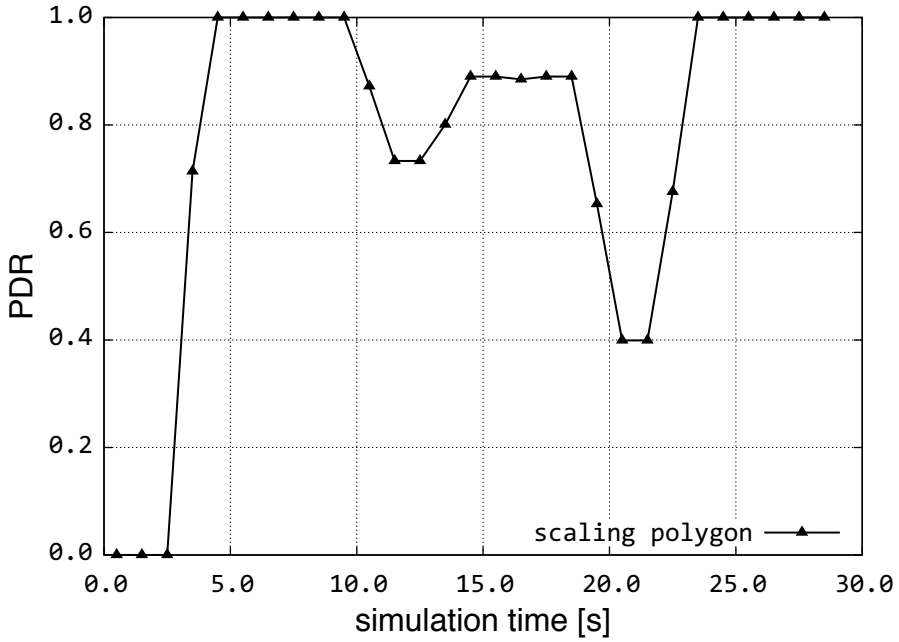

(b) Scaling polygon PDR

Figure 6.20: Scaling polygon challenge scenario and PDR for Sprint physical topology

Mobile, in which 4 out of the 5 affected nodes are PoP nodes. In the largest affected area there are a total of 10 nodes, 6 of which are the PoP nodes. However, none of the three circular challenge areas cover any logical links or nodes on the map in Figure 6.3,

permitting us to investigate the differences between logical and physical topologies.

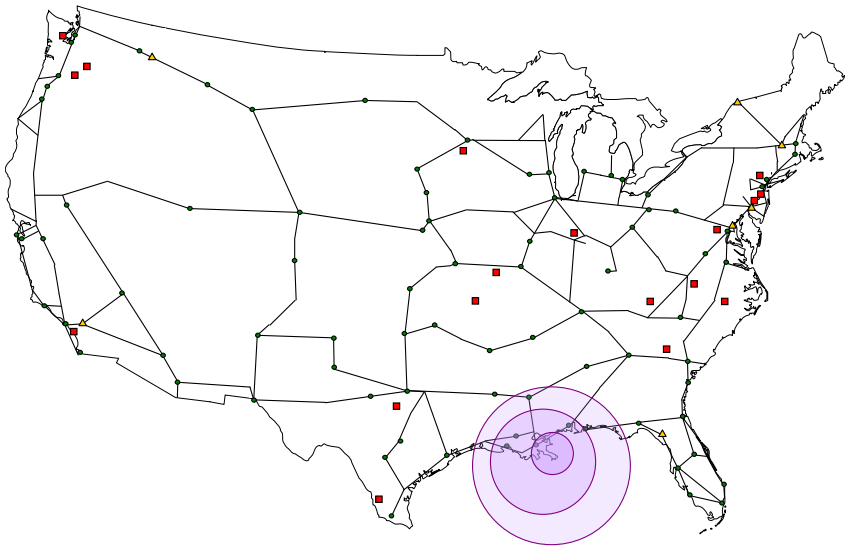

Figure 6.21: South central area-based challenge scenario

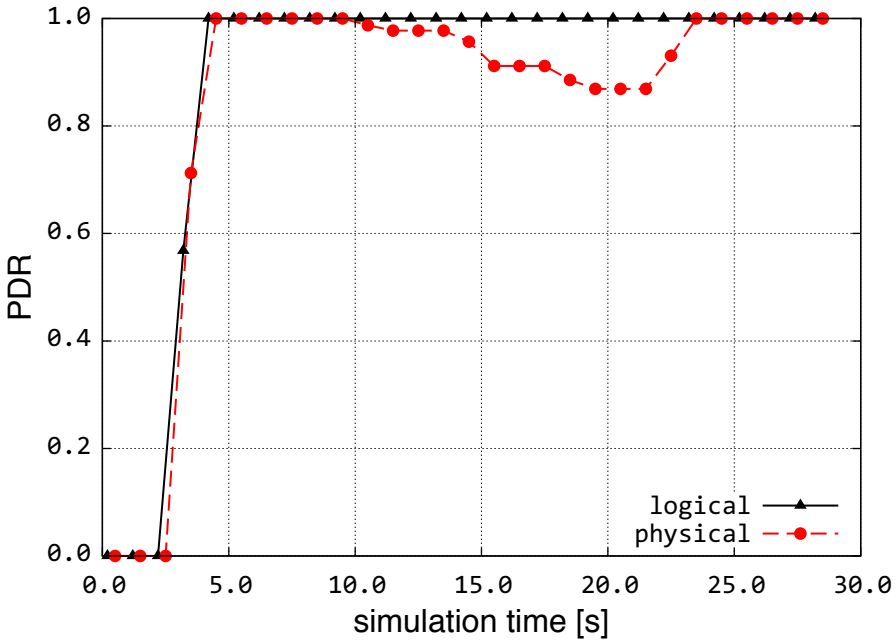

Figure 6.22: PDR during south central US challenge scenario

The network performance of physical and logical topologies when the south central US region is challenged is shown in Figure 6.22. Since there are no nodes or links in the logical topology impacted, the PDR appears to be 100%. On the other hand, the PDR

of the physical topology drops to 98%, 91%, and 86%, respectively, as the challenge area covers more nodes and links. This demonstrates that it is imperative to study the impact of area-based challenges on the *physical* topologies. Traditional layer-3 logical topologies are insufficient to understand the impact of physical challenges against the network infrastructure.

To conclude, networks face challenges that are inherent in the environment and the consequences of these challenges can be costly. The resulting impact of these challenges is related to the probability of occurrence, the magnitude of a challenge, and the duration of a challenge. In this work we study the temporal and spatial characteristics of network challenges. Even if a challenge is an act of nature, statistical recording of such challenges can provide for the allocation of resources at the right time to reduce the impact of the challenges. Another factor contributing the probability of occurrence is social behavior: malicious attacks [262, 553] affect the severity of the consequences. The magnitude of a challenge can be characterised by the following:

- challenge area
- number of impacted nodes and links in challenge area
- significance of network elements in the challenge area
- traffic carried through the affected area

As seen in our simulation results, magnitude of a challenge impacts the result of challenges. Finally, the duration of a challenge is critical factor impacting the network performance. This is related to the ATIS/ANSI (unservability, duration, extent) triple [4] (cf. Section 2.4, Figure 2.5). Natural or human-made disasters causing power outages increase the network downtime [369, 554]. Our framework could provide valuable insight

for probable consequences of network failures during the varying challenge duration. Understanding the characteristics of challenges can provide insight into the mitigation strategies to cope with various challenges.

6.2 Experimental Evaluation on GpENI Testbed

Experimentation is another technique to evaluate the resilience of networks. We presented some of the recent experimentation efforts to evaluate network resiliency in Section 2.4. In this section we first present an overview of the GpENI Future Internet testbed. Next, we present the experimentations performed to evaluate the heuristic algorithm that improves the algebraic connectivity of a graph described in Section 5.1.

6.2.1 GpENI Testbed Overview

The Great Plains Environment for Network Innovation – GpENI [48] is an international programmable network testbed centered on a regional network between The University of Kansas (KU), University of Missouri – Kansas City (UMKC), Kansas State University (KSU), University of Nebraska – Lincoln (UNL), supported with Brocade OpenFlow switches and Ciena CoreDirectors, in collaboration with the Kansas Research and Education Network (KanREN). GpENI is funded in part by National Science Foundation GENI (Global Environment for Network Innovations) program and the EU FIRE (Future Internet Research and Experimentation) programme. International topology is anchored on Lancaster University in the UK, ETH Zürich and Uni-Bern in Switzerland, G-Lab at Kaiserslautern in Germany, and NorNet at Simula in Norway.

Objectives

The aim of the GpENI project is to build a collaborative research infrastructure in Kansas, the Great Plains region, and internationally. It provides a programmable network infrastructure enabling GpENI member institutions to conduct experiments in Future Internet architecture. The flexible GpENI infrastructure supports the GENI program, mesoscale OpenFlow deployment, and GeniRack access. The GpENI testbed provides an open environment on which the networking research community can develop and conduct network experiments with emphasis on network resilience.

GpENI Topology and Network Infrastructure

GpENI is built upon tinc-meshed VLAN tunneled and OpenFlow interconnected between the principal GpENI institutions, with direct connectivity across GPN, Internet2, GANT2, and JANET backbones. Administration of Midwest GpENI infrastructure is assisted by KanREN. Each university has a GpENI node cluster interconnected to one-another and the rest of GENI by Ethernet VLAN, and within KanREN by OpenFlow. GpENI is part of GENI control framework Cluster B. GpENI is one of two network infrastructure projects (along with Mid-Atlantic Crossroads) that runs the PlanetLab control framework and interfaces with other Cluster B participants, running GUSH experiment control and Raven code deployment.

38 node clusters are coming up in 17 nations, shown in Figure 6.23. Each GpENI node cluster consists of several components, physically interconnected by a Gigabit Ethernet switch to allow arbitrary and flexible experiments. Each cluster consists of the following components: GpENI management and control processor, PlanetLab programmable nodes managed by MyPLC with GENIwrapper SFA sub-aggregate manager, VINI-based pro-

grammable routers, a managed Gigabit Ethernet switch, and site specific experimental nodes.

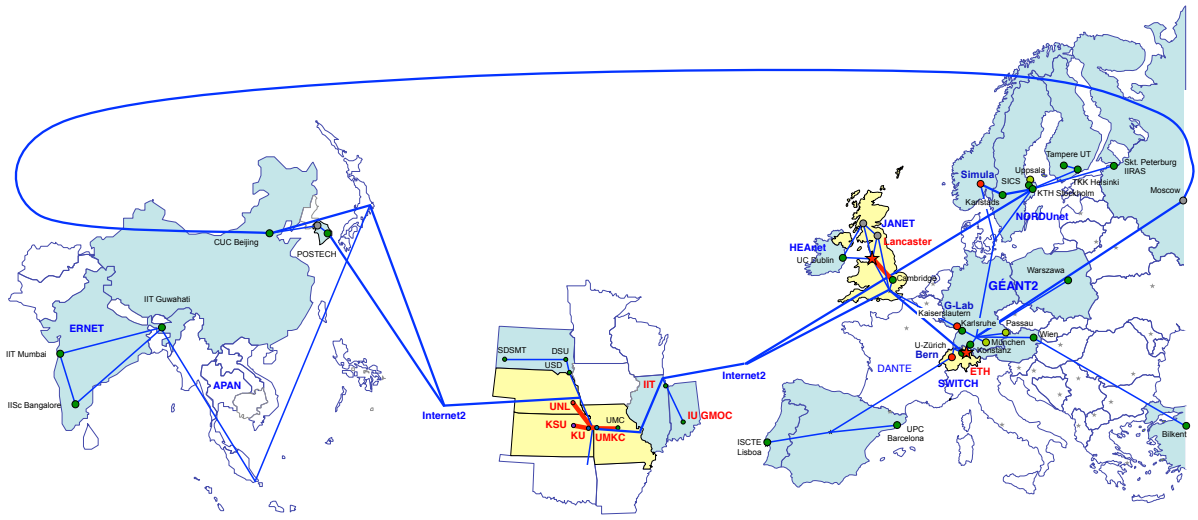

Figure 6.23: GpENI international connectivity

Experimentation

GpENI is undergoing significant regional and international expansion, with institutions providing node clusters tunneled (L2TPv3 or IP) into KU. Moreover, we deploy tinc to perform arbitrary L2 meshing to experiment with network algorithms on GpENI testbed. Next, we explain the experiments we run on the GpENI testbed to validate the graph algorithm presented in Section 5.1.

6.2.2 Graph Algorithm Evaluation on GpENI

We develop a heuristic algorithm that improves the connectivity of a graph in terms of the algebraic connectivity metric by adding links [538]. Algebraic connectivity is defined as the second smallest eigenvalue of the Laplacian matrix and it is widely used for topological optimisations as described in Section 2.3.1. A secondary objective of our algorithm is to select the links that improve the algebraic connectivity of the graph in

the least costly fashion in which we capture the cost of network as the total link length. The heuristic to increase algebraic connectivity in a graph is based on adding links to the nodes that have the fewest incident links (i.e. minimal degree nodes).

Large scale resilience experiments are run over interconnected PlanetLab clusters using tinc VPN tunneling software [555]. The tinc project allows creation of arbitrary topologies while preventing broadcast storms. We create sample topologies consisting of five GpENI PlanetLab nodes (i.e. KSU, KU, Cambridge, KIT, Bern) as shown in Figures 6.24 and 6.25 [556]. The sample binary-tree topology as shown in Figure 6.24 has the root node in Cambridge. The KU node is the highest-degree node in the partial-mesh topology shown in Figure 6.25.

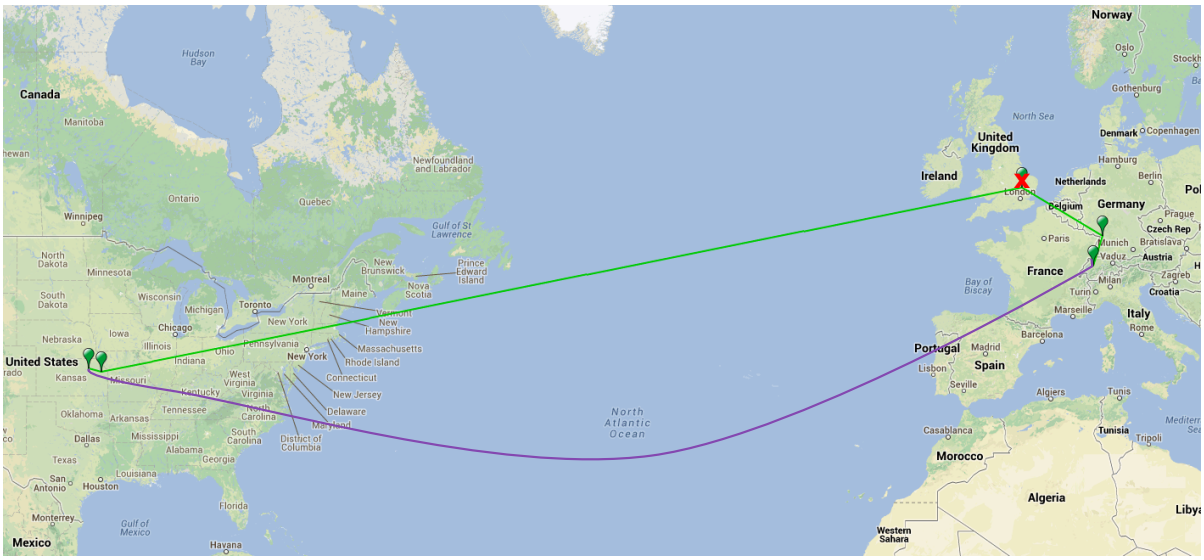

Figure 6.24: Example binary-tree topology

We measure the network performance in terms of flow robustness (described in Section 4.2.2), which quantifies resilience as the fraction of node pairs that remain connected in a network. Simultaneous ping traffic between every pair of node in each topology is generated. We pause tinc processes to emulate challenges against critical nodes in each scenario topology. Flow robustness is measured on the sample topologies with and with-

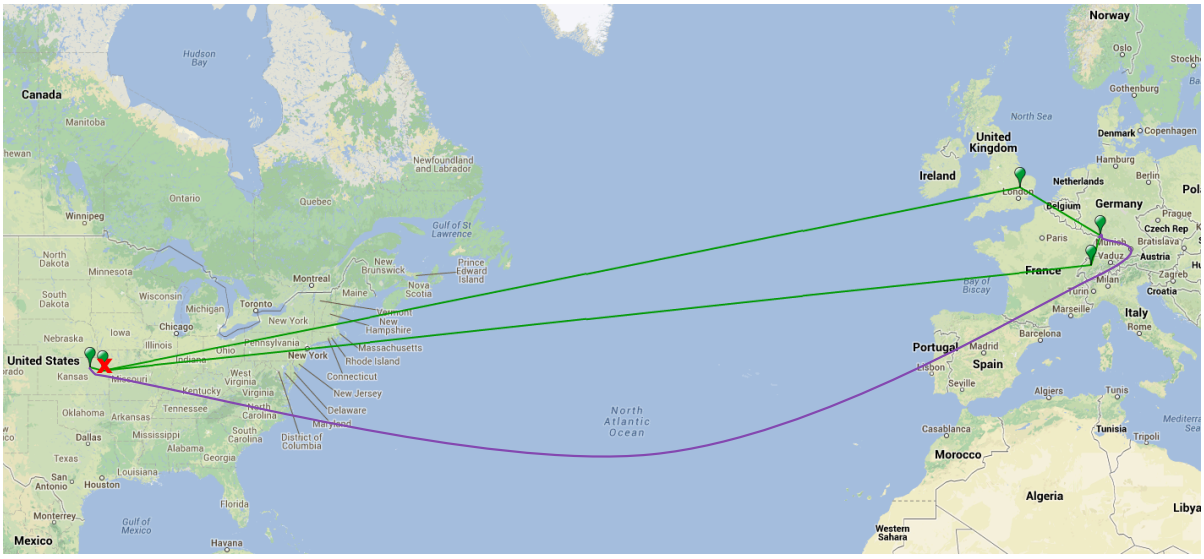

Figure 6.25: Example partial-mesh topology

out our optimisation algorithm being applied as shown in Figures 6.26 and 6.27 .

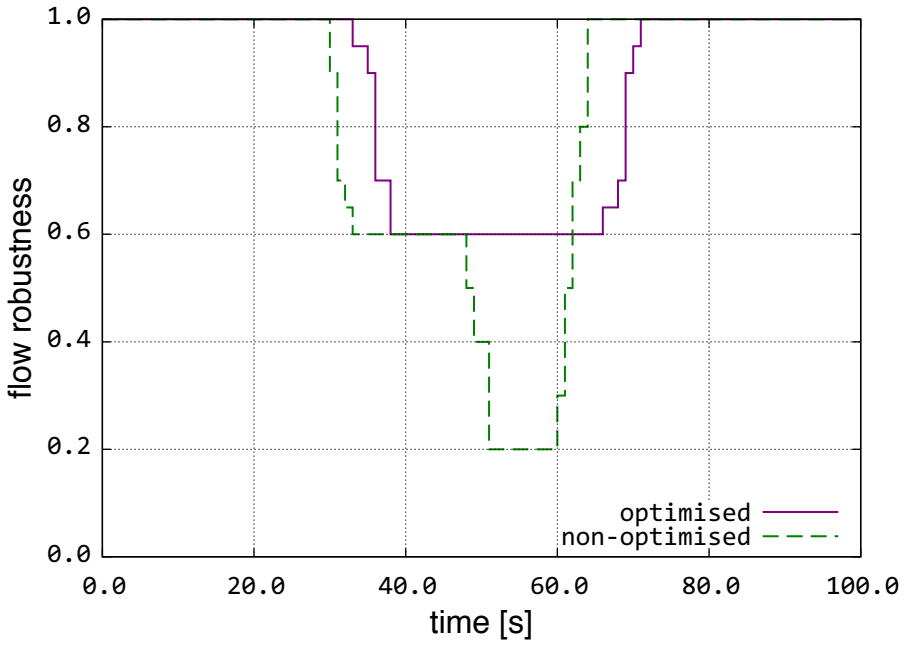

Figure 6.26: Robustness of optimised and non-optimised binary-tree topologies

We plot the flow robustness of the binary-tree scenario as shown in Figure 6.26. The scenario represents an attack against the highest betweenness node (Cambridge) in this

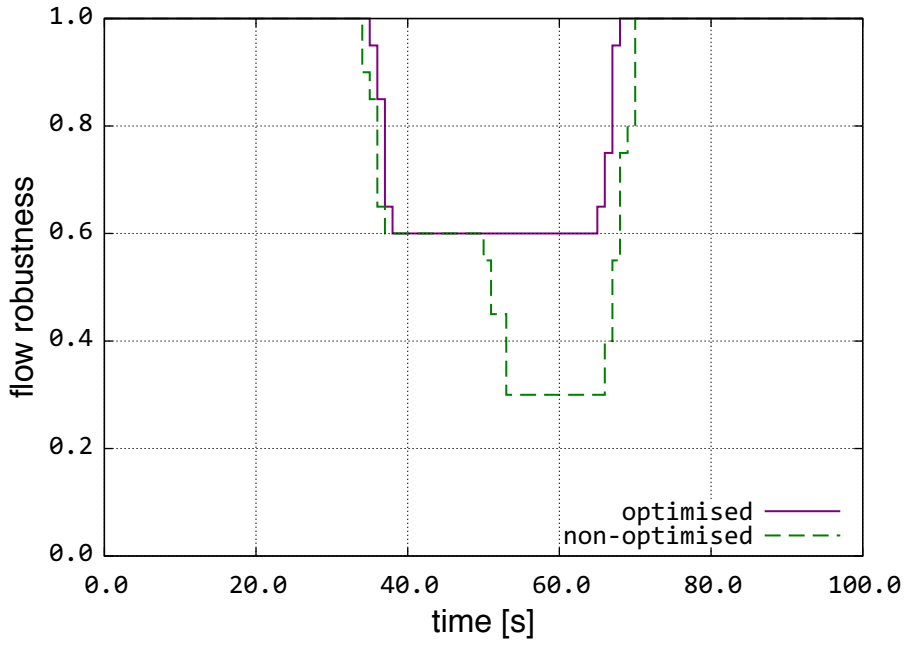

Figure 6.27: Robustness of optimised and non-optimised partial-mesh topologies

tree topology as shown in Figure 6.24. The optimised topology performs better since additional link (between KSU and Bern) provide alternate path between node pairs. Flow robustness of the partial-mesh scenario is shown in Figure 6.27. In this scenario the highest degree node (KU) is attacked in a partial-mesh topology as shown in Figure 6.25. The optimised topology (with additional link between KSU and KIT) has a flow robustness of 0.6, where as non-optimised topology has a flow robustness of 0.3. The flow robustness of non-optimised partial-mesh topology is better than the non-optimised binary-tree topology when critical nodes are attacked because nodes are connected more connected in the partial-mesh topology. This resilience experiment demonstrates creation of arbitrary topologies and application of our heuristic algorithm that improves algebraic connectivity metric on the large-scale GpENI testbed.

We have mainly measured the resilience as flow robustness and PDR (packet delivery ratio) when evaluating the networks throughout this work. We consider the area under

these curves to evaluate network resilience. While flow robustness captures how well the components of a graph are connected, aggregate PDR captures the ratio of packets delivered to the total number of packets in the network. In essence, flow robustness and PDR captures similar measures for analysing network resilience from a graph-theoretical perspective. We note that, correct parameter settings are essential so that the nodes and links are not overloaded and packets are not dropped due to congestion in simulations. Considering the analytical, simulation, and experimentation results, they all give a fairly good assessment of the network resilience.

6.3 Summary

The KU-CSM evaluates the impact of a variety of challenges including large-scale correlated failures on the networks, which allows for analysis of network performance in the ns-3 network simulator. The geographic failure scenarios consist of a regular circle centered at a point with constant radius R as well as n-sided polygons for irregular-shaped challenge models. We note the importance of modelling irregular shapes is to reflect reality, since not all challenges are circular shaped. Moreover, KU-CSM can model area-based challenges that evolve temporally and spatially. Next, we evaluate the heuristic algorithm that increases the algebraic connectivity of a graph on the GpENI Future Internet testbed. The results shown in this work cross-validate the analytical, simulation, and experimentation evaluation of network resilience.

Chapter 7

Conclusions and Future Work

This dissertation presents network models, design and optimisation of networks, and a methodology to evaluate networks under challenges. We find that there are a variety of challenges and that no single mechanism can address the full resilience requirement of a network service. The cost and resilience trade-offs should be considered according to the service requirements of the network. Moreover, the resilience evaluation methodologies cross-validate results obtained by simulations and experimentations. This chapter presents conclusions drawn from the major contributions of the dissertation and directions for future work.

7.1 Conclusions

In Chapter 3, we presented the known and potential communication network challenges. We tried to answer the question *Why networks fail?* While doing so, we systematically considered challenges within each of the relevant resilience disciplines. Based on the identified challenges, we provided a taxonomy of challenges that can assist a designer when considering essential factors among several ones for cost-efficient resilient network design. The aim of this survey is to assist network designers in avoiding the mistakes

of the past and also to aid in developing Future Internet architectures. We strived to have a complete and comprehensive taxonomy; however, it will require refinement as new challenges arise. We expect that such a taxonomy will be beneficial for network designers and foster coöperation among researchers.

The documentation and sharing of challenges and their impact is important. From a policy perspective, lack of failure data and improper documentation for resilience analysis of networks have been reported by several researchers [30, 274, 557, 558]. This can be attributed to service providers' unwillingness to share outage information due to security and competitive reasons. An *open* database similar to the US FCC (Federal Communications Commission) NORS (Network Outage Reporting System) [559] for sharing outage information would benefit the research community toward increasing the resilience of the Internet.

In Chapter 4, we have shown that realistically modelling of the Internet requires a collective and systemic analysis of all of its structural properties. Intuitively, fibre-level topologies are laid along right-of-way of the freeways, since it is less costly than line-of-sight installation. We analytically show structural similarities between these physical infrastructures by using the normalised Laplacian spectra. While freeways and physical fibre route graphs share similar grid-like structural characteristics, logical overlays clearly differ from the physical underlays in terms of well-studied graph metrics, spectral properties, and flow robustness values. Physical topologies have higher distance metric values, whereas logical topologies have higher centrality metric values.

Existing models of the Internet often employ a single level perspective. We have developed a multilevel and multiprovider framework to model and evaluate the multilevel nature of the Internet. Using the flow robustness metric we evaluated multilevel graphs and analysed combined communication and transport networks with our multilevel frame-

work. We confirmed that dynamic routing helps alleviate the impact of perturbations and that adaptive challenges degrade multilevel network performance more than non-adaptive challenges. Moreover, as we demonstrated, multilevel graphs yield different performance measures than single level graphs under network perturbations.

Physical level topologies are *necessary* to study the resilience of networks more realistically. We discuss the *fitness* of four geographical graph models applied to graphs with node locations given by those of six actual networks. We evaluate the cost of these synthetically generated graphs based on a cost model, and we find that among the synthetic graph models we studied, the Gabriel model yields topologies with the smallest cost. Furthermore, the cost incurred using synthetic models depends on the number of nodes and the geographic distribution of these nodes. We analyse the topologies generated by the synthetic geographic graph models, and visual inspection of these topologies shows that the Gabriel graphs best capture the grid-like structure of physical level topologies. We then show that geographical physical level graphs are dominated by degree-2 nodes, and removal of them provides more accurate structural metrics, particularly for degree distribution.

In Chapter 5, we introduce two heuristic graph algorithms that optimises the connectivity of a given graph with node locations and is computationally less costly than an exhaustive optimisation. First, we use algebraic connectivity as a measure to improve the connectivity of the graph. This algorithm minimises the cost of adding new links by selecting shorter links with high algebraic connectivity. We introduce a tuning parameter γ to control the effect of the cost function while selecting new links. Furthermore, the candidate links that are being added to improve the connectivity of the graph can be constrained by a length limit in our algorithm. We apply this algorithm to physical- and logical-level topologies of three backbone providers. The results show trade-offs between improving algebraic connectivity and minimising cost, from which a cost-efficient set of

link addition can be chosen based on the value of γ. We show that the algebraic connectivity improvement for the physical level graphs are less than the logical level graphs because of the differences in their structural characteristics, as well as the link length limitation imposed on the candidate links that are added to the physical level topologies.

Next, we introduce a k-diverse path algorithm that considers both the diversity of the nodes and links in the returned paths. We present a new heuristic algorithm that optimises the total path diversity of a given graph with node locations. This algorithm improves the TGD of a graph by adding the cost-efficient link that increases the lowest EPD pair the most. We apply the optimisation algorithm to three realistic physical level topologies. Using the flow robustness graph metric, the path diversity optimised graphs are compared to both lowest degree optimised and non-optimised graphs as they are attacked using node removal based on highest node centrality graph metrics. The path diversity optimised graphs show better resilience to these attacks compared to the lowest degree optimised and non-optimised graphs. Finally, we compare the two heuristic algorithms and show the trade-offs between resilience and cost when designing and optimising networks.

In Chapter 6, we introduce the KU-CSM framework to evaluate network performance when faced by realistic stationary or evolving challenges. This framework separates network topology from challenge specification, which increases tractability and flexibility. We demonstrated that while logical topologies are appropriate for statistical challenge scenarios or analysing network-level attacks, physical topologies are necessary to realistically study geographically correlated failures. Our results indicate that network performance varies depending on the type and severity of the challenge applied. Next, we present experiments conducted on the GpENI testbed for resilience evaluation. The experiments involve evaluating the graph optimisation algorithm on the large-scale GpENI infrastructure. We conclude that analytical, simulation, and experimental evaluations of

network resilience are consistent. While experimentations are costly to build, simulations provide tractable options to evaluate network resilience with the right settings.

7.2 Future Work

In this work we focus on challenges, which are adverse events that result in service failures and we provide a comprehensive survey of consequences of these challenges. A future research direction is to systematically classify *failures*. Such classification may prove to be useful for further analysis of temporal and spatial characteristics of failures. Combined with risk-cost-complexity analysis [26], it can help efficient design of the system with limited resources for maximal resilience gain. Another question in regards to different challenges is how do we quantify one challenge compared to another? In other words, considering two challenges with characteristics differing from one another, how do we assess which one is worse? Investigating the level of services provided by the network and the cost and time to recover and remediate can help this second question. However, further research is needed to evaluate these resilience aspects.

The critical infrastructures increasingly depend on each other. A question for which we don't have an answer is: What is the level of dependency between critical infrastructures? Recently, terrorist attacks in Turkey on gas pipelines impacted the Internet in Northern Iraq [560]. A pipeline explosion near Charleston, WV in the US impacted the fibre cables in which resulted telephone disruptions in several states [561]. Moreover, motorway traffic was rerouted on the Interstate Highway 77 since it was damaged near the explosion. As critical infrastructure, such as water distribution and transportation, are essential for the society, how the failure of one critical infrastructure impacts communication networks is still unknown.

We presented two heuristic graph algorithms that increase the resilience of a graph in a cost-efficient manner. For our future work, we would like to run our heuristic algorithm using graph properties such as clustering coefficient and graph spectra. For example, we can add clustering coefficient to replace the algebraic connectivity or we can add both with a parameter to weight their effect in ranking the links that need to be added to improve connectivity of the graph. It is worthwhile to investigate a variety of graph metrics and analyse which one helps capturing the resilience properties the most. We will also modify our algorithm to achieve a specified graph metric value with a constrained budget. Finally, we plan to design an algorithm that improves the connectivity the most based on node additions while minimising the cost.

Recently, there has been interest from the research community in modelling correlated failures. However, these challenge models typically have been oversimplified in modelling a tornado as a line or an EMP weapon as a static circle and we have shown realistic area-based challenge models [23]. Moreover, a major missing piece among these is a lack of models that study social, economical, and policy challenges. For example, how would a nationwide Internet outage impact our overall communication? What is the consequence of depeering ISP A from ISP B? Another research direction in the form of experimentation is how to evaluate graph resilience within the limited resources of the experimentation testbed. This question is fundamentally about the scaling of an experimentation scenario on the limited resources of the experimentation testbed.

Bibliography

[1] James P.G. Sterbenz and David Hutchison. ResiliNets: Multilevel Resilient and Survivable Networking Initiative Wiki. `http://wiki.ittc.ku.edu/resilinets`, April 2006.

[2] James P.G. Sterbenz, David Hutchison, Egemen K. Çetinkaya, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller, and Paul Smith. Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Computer Networks*, 54(8):1245–1265, 2010.

[3] James P.G. Sterbenz, David Hutchison, Egemen K. Çetinkaya, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller, and Paul Smith. Redundancy, Diversity, and Connectivity to Achieve Multilevel Network Resilience, Survivability, and Disruption Tolerance (invited paper). *Springer Telecommunication Systems*, 2012. (accepted April 2012).

[4] T1A1.2 Working Group. Network Survivability Performance. Technical Report T1A1.2/93-001R3, Alliance for Telecommunications Industry Solutions (ATIS), November 1993.

[5] James P.G. Sterbenz, Egemen K. Çetinkaya, Mahmood A. Hameed, Abdul Jabbar, Qian Shi, and Justin P. Rohrer. Evaluation of Network Resilience, Survivability,

and Disruption Tolerance: Analysis, Topology Generation, Simulation, and Experimentation (invited paper). *Telecommunication Systems*, 52(2):705–736, 2013.

[6] Nigel Edwards and Owen Rees. A Model for Failures in Dependable Systems. Technical report APM.1143.01, ANSA, March 1994.

[7] Yun Liu and Kishor S. Trivedi. Survivability Quantification: The Analytical Modeling Approach. *International Journal of Performability Engineering*, 2(1):29–44, 2006.

[8] Steven M. Rinaldi, James P. Peerenboom, and Terrence K. Kelly. Identifying, Understanding, and Analyzing Critical Infrastructure Interdependencies. *IEEE Control Systems Magazine*, 21(6):11–25, 2001.

[9] Jeff Papows. *Glitch: The Hidden Impact of Faulty Software*. Prentice Hall Press, Upper Saddle River, NJ, USA, 1st edition, 2010.

[10] Richard Power and Dario Forte. Ten years in the wilderness – a retrospective Part I: Nine false notions and nine steps to success. *Computer Fraud & Security*, 2006(1):8–13, 2006.

[11] Allen Householder, Kevin Houle, and Chad Dougherty. Computer Attack Trends Challenge Internet Security. *IEEE Computer*, 35(4):5–7, 2002.

[12] John C. McDonald, Paul Baran, Floyd Becker, Cullen M. Crain, Howard Frank, Lewis E. Franks, Paul E. Green, Erik K. Grimmelmann, E. Fletcher Haselton, Amos E. Joel, Donald Kuyper, Richard B. Marsten, David L. Mills, Lee M. Paschall, and Casimir S. Skrzypczak. *Growing Vulnerability of the Public Switched Networks: Implications for National Security Emergency Preparedness*. The National Academy Press, Washington, D.C., 1989.

[13] Algirdas Avižienis, Jean-Claude Laprie, Brian Randell, and Carl Landwehr. Basic Concepts and Taxonomy of Dependable and Secure Computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1):11–33, 2004.

[14] Jean-Claude Laprie, Algirdas Avižienis, and H. Kopetz, editors. *Dependability: Basic Concepts and Terminology*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1992.

[15] Algirdas Avižienis and Jean-Claude Laprie. Dependable Computing: From Concepts to Design Diversity. *Proceedings of the IEEE*, 74(5):629–638, 1986.

[16] Algirdas Avižienis, Jean-Claude Laprie, and Brian Randell. Dependability and Its Threats: A Taxonomy. In Renè Jacquart, editor, *Building the Information Society*, volume 156 of *IFIP International Federation for Information Processing*, pages 91–120. Springer Boston, 2004.

[17] Linlin Xie, Paul Smith, Mark Banfield, Helmut Leopold, James P.G. Sterbenz, and David Hutchison. Towards Resilient Networks Using Programmable Networking Technologies. In David Hutchison, Spyros Denazis, Laurent Lefevre, and Gary Minden, editors, *Active and Programmable Networks*, volume 4388 of *Lecture Notes in Computer Science*, pages 83–95. Springer Berlin / Heidelberg, 2009.

[18] Rabat A. Mahmood. Simulating Challenges to Communication Networks for Evaluation of Resilience. Master's thesis, The University of Kansas, Lawrence, KS, August 2009.

[19] Michael Fry, Mathias Fischer, Merkouris Karaliopoulos, Paul Smith, and David Hutchison. Challenge Identification for Network Resilience. In *Proceedings of the 6th IEEE/EURO-NF Conference on Next Generation Internet (NGI)*, pages 1–8, Paris, June 2010.

[20] Egemen K. Çetinkaya, Dan Broyles, Amit Dandekar, Sripriya Srinivasan, and James P.G. Sterbenz. A Comprehensive Framework to Simulate Network Attacks and Challenges. In *Proceedings of the 2nd IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, pages 538–544, Moscow, October 2010.

[21] James P.G. Sterbenz, Egemen K. Çetinkaya, Mahmood A. Hameed, Abdul Jabbar, and Justin P. Rohrer. Modelling and Analysis of Network Resilience (invited paper). In *Proceedings of the 3rd IEEE/ACM International Conference on Communication Systems and Networks (COMSNETS)*, pages 1–10, Bangalore, January 2011.

[22] ENISA Virtual Working Group on Network Providers Resilience Measures. Network resilience and security: Challenges and measures. Technical Report WP 2009 – WPK 1.2 VWG 1, ENISA – European Network and Information Security Agency, December 2009.

[23] Egemen K. Çetinkaya, Dan Broyles, Amit Dandekar, Sripriya Srinivasan, and James P.G. Sterbenz. Modelling Communication Network Challenges for Future Internet Resilience, Survivability, and Disruption Tolerance: A Simulation-Based Approach. *Telecommunication Systems*, 52(2):751–766, 2013.

[24] Paul Smith, David Hutchison, James P.G. Sterbenz, Marcus Schöller, Ali Fessi, Merkouris Karaliopoulos, Chidung Lac, and Bernhard Plattner. Network Resilience: A Systematic Approach. *IEEE Communications Magazine*, 49(7):88–97, 2011.

[25] Yue Yu, Michael Fry, Alberto Schaeffer-Filho, Paul Smith, and David Hutchison. An Adaptive Approach to Network Resilience: Evolving Challenge Detection and Mitigation. In *Proceedings of the 8th IEEE International Workshop on the Design*

of Reliable Communication Networks (DRCN), pages 172–179, Krakow, October 2011.

[26] Marcus Schöller, Paul Smith, and David Hutchison. Assessing Risk for Network Resilience. In *Proceedings of the 3rd IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, pages 1–7, Budapest, October 2011.

[27] Alberto Schaeffer-Filho, Paul Smith, and Andreas Mauthe. Policy-driven Network Simulation: a Resilience Case Study. In *Proceedings of the ACM Symposium on Applied Computing (SAC)*, pages 492–497, TaiChung, March 2011.

[28] Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, and Stephen Wolff. A Brief History of the Internet. *ACM Computer Communication Review*, 39(5):22–31, 2009.

[29] Chris W. Johnson and Rhys Williams. Computational Support for Identifying Safety and Security Related Dependencies between National Critical Infrastructures. In *Proceedings of the 3rd IET International Conference on System Safety*, pages 1–6, Birmingham, October 2008.

[30] Bianca Schroeder and Garth A. Gibson. A Large-Scale Study of Failures in High-Performance Computing Systems. *IEEE Transactions on Dependable and Secure Computing*, 7(4):337–351, 2010.

[31] Benoit Donnet and Timur Friedman. Internet topology discovery: a survey. *IEEE Communications Surveys & Tutorials*, 9(4):56–69, 2007.

[32] Chris Metz. Interconnecting ISP Networks. *IEEE Internet Computing*, 5(2):74–80, 2001.

[33] Neil Spring, Ratul Mahajan, David Wetherall, and Thomas Anderson. Measuring ISP topologies with Rocketfuel. *IEEE/ACM Transactions on Networking*, 12(1):2–16, 2004.

[34] Ramakrishnan Durairajan, Subhadip Ghosh, Xin Tang, Paul Barford, and Brian Eriksson. Internet Atlas: A Geographic Database of the Internet. In *Proceedings of the 5th ACM HotPlanet Workshop*, pages 15–20, Hong Kong, August 2013.

[35] Hamed Haddadi, Miguel Rio, Gianluca Iannaccone, Andrew Moore, and Richard Mortier. Network Topologies: Inference, Modeling, and Generation. *IEEE Communications Surveys & Tutorials*, 10(2):48–69, 2008.

[36] Dmitri Krioukov, kc claffy, Marina Fomenkov, Fan Chung, Alessandro Vespignani, and Walter Willinger. The Workshop on Internet Topology (WIT) Report. *ACM Comput. Commun. Rev.*, 37(1):69–73, 2007.

[37] Maciej Kurant and Patrick Thiran. Layered complex networks. *Phys. Rev. Lett.*, 96:138701, April 2006.

[38] Maciej Kurant, Patrick Thiran, and Patric Hagmann. Error and attack tolerance of layered complex networks. *Phys. Rev. E*, 76(2):026103, August 2007.

[39] Deep Medhi and David Tipper. Multi-Layered Network Survivability – Models, Analysis, Architecture, Framework and Implementation: An Overview. In *Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX)*, volume 1, pages 173–186, Hilton Head Island, SC, January 2000.

[40] Howard Frank and Wushow Chou. Topological optimization of computer networks. *Proceedings of the IEEE*, 60(11):1385–1397, 1972.

[41] Robert S. Wilkov. Analysis and design of reliable computer networks. *IEEE Transactions on Communications*, 20(3):660–678, 1972.

[42] J. M. McQuillan. Graph theory applied to optimal connectivity in computer networks. *ACM SIGCOMM Comput. Commun. Rev.*, 7(2):13–41, April 1977.

[43] Mario Gerla and Leonard Kleinrock. On the topological design of distributed computer networks. *IEEE Transactions on Communications*, 25(1):48–60, January 1977.

[44] Robert R. Boorstyn and Howard Frank. Large-scale network topological optimization. *IEEE Transactions on Communications*, 25(1):29–47, January 1977.

[45] John G. Klincewicz. Hub location in backbone/tributary network design: a review. *Location Science*, 6(1–4):307–335, December 1998.

[46] Michael O. Ball. Complexity of network reliability computations. *Networks*, 10(2):153–165, 1980.

[47] Michael O. Ball. Computational complexity of network reliability analysis: An overview. *IEEE Transactions on Reliability*, 35(3):230–239, August 1986.

[48] James P.G. Sterbenz, Deep Medhi, Byrav Ramamurthy, Caterina Scoglio, David Hutchison, Bernhard Plattner, Tricha Anjali, Andrew Scott, Cort Buffington, Gregory E. Monaco, Don Gruenbacher, Rick McMullen, Justin P. Rohrer, John Sherrell, Pragatheeswaran Angu, Ramkumar Cherukuri, Haiyang Qian, and Nidhi Tare. The Great Plains Environment for Network Innovation (GpENI): A Programmable Testbed for Future Internet Architecture Research. In Thomas Magedanz, Anastasius Gavras, Nguyen Huu Thanh, and Jeffry S. Chase, editors, *Testbeds and Research Infrastructures. Development of Networks and Communities*, volume 46

of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 428–441. Springer Berlin Heidelberg, 2011.

[49] Justin P. Rohrer, Egemen K. Çetinkaya, and James P.G. Sterbenz. Progress and Challenges in Large-Scale Future Internet Experimentation using the GpENI Programmable Testbed. In *Proceedings of the 6th ACM International Conference on Future Internet Technologies (CFI)*, pages 46–49, Seoul, June 2011.

[50] David M. Nicol, William H. Sanders, and Kishor S. Trivedi. Model-Based Evaluation: From Dependability to Security. *IEEE Transactions on Dependable and Secure Computing*, 1(1):48–65, 2004.

[51] ResiliNets Topology Map Viewer. `http://www.ittc.ku.edu/resilinets/maps/`, January 2011.

[52] R.J. Ellison, D. Fisher, R.C. Linger, H.F. Lipson, T. Longstaff, and N.R. Mead. Survivable network systems: An emerging discipline. Technical Report CMU/SEI-97-TR-013, Software Engineering Institute, Carnegie Mellon University, 1997.

[53] James P.G. Sterbenz, Rajesh Krishnan, Regina Rosales Hain, Alden W. Jackson, David Levin, Ram Ramanathan, and John Zao. Survivable Mobile Wireless Networks: Issues, Challenges, and Research Directions. In *Proceedings of the 1st ACM Workshop on Wireless Security (WiSe)*, pages 31–40, Atlanta, GA, September 2002.

[54] Larry Niven. *The Flight of the Horse*. Ballantine Books, New York, 1973.

[55] Linlin Xie, Paul Smith, David Hutchison, Mark Banfield, Helmut Leopold, Abdul Jabbar, and James P.G. Sterbenz. From Detection to Remediation: A Self-Organized System for Addressing Flash Crowd Problems. In *Proceedings of the IEEE International Conference on Communications (ICC)*, pages 5809–5814, Beijing, May 2008.

[56] Abhishek Chandra and Prashant Shenoy. Effectiveness of Dynamic Resource Allocation for Handling Internet Flash Crowds. Technical Report TR03-37, University of Massachusetts Amherst, Amherst, MA, November 2003.

[57] Richard A. Kemmerer and Giovanni Vigna. Intrusion detection: a brief history and overview. *IEEE Security & Privacy*, 35(4):27–30, 2002.

[58] Animesh Patcha and Jung-Min Park. An overview of anomaly detection techniques: Existing solutions and latest technological trends. *Computer Networks*, 51(12):3448–3470, 2007.

[59] Juan M. Estevez-Tapiador, Pedro Garcia-Teodoro, and Jesus E. Diaz-Verdejo. Anomaly detection methods in wired networks: a survey and taxonomy. *Computer Communications*, 27(16):1569–1584, 2004.

[60] E. Mannie and D. Papadimitriou. Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS). RFC 4427 (Informational), March 2006.

[61] Małgorzata Steinder and Adarshpal S. Sethi. A Survey of Fault Localization Techniques in Computer Networks. *Science of Computer Programming*, 53(2):165–194, 2004.

[62] Bezalel Gavish. Topological design of computer communication networks – the overall design problem. *European Journal of Operational Research*, 58(2):149–172, 1992.

[63] Abdullah Konak and Alice E. Smith. Network reliability optimization. In Mauricio G.C. Resende and Panos M. Pardalos, editors, *Handbook of Optimization in Telecommunications*, pages 735–760. Springer US, 2006.

[64] Hakki C. Cankaya, Ana Lardies, and Gary W. Ester. Network Design Pptimization from an Availability Perspective. In *Proceedings of the 11th International Telecommunications Network Strategy and Planning Symposium (NETWORKS)*, pages 359–364, Vienna, June 2004.

[65] Steven Chamberland, Marc St-Hilaire, and Samuel Pierre. On the point-of-presence optimization problem in IP networks. *Canadian Journal of Electrical and Computer Engineering*, 30(3):137–143, 2005.

[66] D. Fay, H. Haddadi, A. Thomason, A.W. Moore, R. Mortier, A. Jamakovic, S. Uhlig, and M. Rio. Weighted Spectral Distribution for Internet Topology Analysis: Theory and Applications. *IEEE/ACM Transactions on Networking*, 18(1):164–176, 2010.

[67] Ali Sydney, Caterina Scoglio, and Don Gruenbacher. Optimizing algebraic connectivity by edge rewiring. *Applied Mathematics and Computation*, 219(10):5465–5479, 2013.

[68] Huijuan Wang and Piet Van Mieghem. Algebraic connectivity optimization via link addition. In *Proceedings of the 3rd ICST International Conference on Bio-Inspired Models of Network, Information and Computing Sytems (BIONETICS)*, pages 22:1–22:8, Hyogo, Japan, November 2008.

[69] L. da F. Costa, F. A. Rodrigues, G. Travieso, and P. R. Villas Boas. Characterization of complex networks: A survey of measurements. *Advances in Physics*, 56(1):167–242, 2007.

[70] Inder M. Soi and Krishnan K. Aggarwal. Reliability indices for topological design of computer communication networks. *IEEE Transactions on Reliability*, R-30(5):438–443, December 1981.

[71] Justin P. Rohrer, Abdul Jabbar, and James P.G. Sterbenz. Path Diversification for Future Internet End-to-End Resilience and Survivability. *Springer Telecommunication Systems*, 2012.

[72] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.

[73] Justin P. Rohrer, Abdul Jabbar, and James P.G. Sterbenz. Path Diversification: A Multipath Resilience Mechanism. In *Proceedings of the 7th IEEE International Workshop on the Design of Reliable Communication Networks (DRCN)*, pages 343–351, Washington, D.C., October 2009.

[74] William Liu, Harsha Sirisena, Krzysztof Pawlikowski, and Allan McInnes. Utility of algebraic connectivity metric in topology design of survivable networks. In *Proceedings of the 7th IEEE International Workshop on Design of Reliable Communication Networks (DRCN)*, pages 131–138, Washington, DC, October 2009.

[75] Egemen K. Çetinkaya, Mohammed J.F. Alenazi, Justin P. Rohrer, and James P.G. Sterbenz. Topology Connectivity Analysis of Internet Infrastructure Using Graph Spectra. In *Proceedings of the 4th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, pages 752–758, St. Petersburg, October 2012.

[76] Egemen K. Çetinkaya, Mohammed J.F. Alenazi, Andrew M. Peck, Justin P. Rohrer, and James P.G. Sterbenz. Multilevel Resilience Analysis of Transportation and Communication Networks. *Springer Telecommunication Systems Journal*, 2013. accepted on July 2013.

[77] John G. Apostolopoulos and Mitchell D. Trott. Path Diversity for Enhanced Media Streaming. *IEEE Communications Magazine*, 42(8):80–87, 2004.

[78] Aun Haider and Richard Harris. Recovery Techniques in Next Generation Networks. *IEEE Communications Surveys & Tutorials*, 9(3):2–17, 2007.

[79] Junghee Han, David Watson, and Farnam Jahanian. An Experimental Study of Internet Path Diversity. *IEEE Transactions on Dependable and Secure Computing*, 3(4):273–288, 2006.

[80] Jiayue He and Jennifer Rexford. Toward Internet-Wide Multipath Routing. *IEEE Network Magazine*, 22(2):16–21, 2008.

[81] J. W. Suurballe. Disjoint paths in a network. *Networks*, 4(2), 1974.

[82] J. W. Suurballe and R. E. Tarjan. A quick method for finding shortest pairs of disjoint paths. *Networks*, 14(2), 1984.

[83] Ramesh Bhandari. Optimal Diverse Routing in Telecommunication Fiber Networks. In *Proceedings of the IEEE INFOCOM*, volume 3, pages 1498–1508, Toronto, June 1994.

[84] Murtaza Motiwala, Megan Elmore, Nick Feamster, and Santosh Vempala. Path Splicing. In *Proceedings of the ACM SIGCOMM*, pages 27–38, Seattle, WA, August 2008.

[85] Xiaowei Yang and David Wetherall. Source Selectable Path Diversity via Routing Deflections. In *Proceedings of the ACM SIGCOMM*, pages 159–170, Pisa, September 2006.

[86] Ariel Orda and Alexander Sprintson. Efficient Algorithms for Computing Disjoint QoS Paths. In *Proceedings of the IEEE INFOCOM*, volume 1, pages 727–738, Hong Kong, March 2004.

[87] Yuchun Guo, Fernando Kuipers, and Piet Van Mieghem. Link-disjoint paths for reliable QoS routing. *International Journal of Communication Systems*, 16(9):779–798, 2003.

[88] Feng Wang and Lixin Gao. Path Diversity Aware Interdomain Routing. In *Proceedings of the IEEE INFOCOM*, pages 307–315, Rio de Janeiro, April 2009.

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

[90] Yufei Cheng, Junyan Li, and James P.G. Sterbenz. Path Geo-diversification: Design and Analysis. In *Proceedings of the 5th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, Almaty, September 2013.

[91] H. Han, S. Shakkottai, C. V. Hollot, R. Srikant, and D. Towsley. Multi-Path TCP: A Joint Congestion Control and Routing Scheme to Exploit Path Diversity in the Internet. *IEEE/ACM Transactions on Networking*, 14(6):1260–1271, 2006.

[92] Dahai Xu, Yang Chen, Yizhi Xiong, Chunming Qiao, and Xin He. On the complexity of and algorithms for finding the shortest path with a disjoint counterpart. *IEEE/ACM Transactions on Networking*, 14(1):147–158, 2006.

[93] Chung-Lun Li, S.Thomas McCormick, and David Simchi-Levi. The complexity of finding two disjoint paths with min-max objective function. *Discrete Applied Mathematics*, 26(1):105–115, 1990.

[94] Henry A. Malec. Communications Reliability: A Historical Perspective. *IEEE Transactions on Reliability*, 47(3):SP333 –SP345, 1998.

[95] Roy Billinton and Ronald N. Allan. *Reliability Evaluation of Engineering Systems*. Plenum Press, New York, 1992.

[96] Allen M. Johnson, Jr. and Miroslaw Malek. Survey of Software Tools for Evaluating Reliability, Availability, and Serviceability. *ACM Computing Surveys*, 20(4):227–269, 1988.

[97] Michael Grottke, Hairong Sun, Ricardo M. Fricks, and Kishor S. Trivedi. Ten Fallacies of Availability and Reliability Analysis. In Takashi Nanya, Fumihiro Maruyama, András Pataricza, and Miroslaw Malek, editors, *Service Availability*, volume 5017 of *Lecture Notes in Computer Science*, pages 187–206. Springer Berlin Heidelberg, 2008.

[98] John F. Meyer. On Evaluating the Performability of Degradable Computing Systems. *IEEE Transactions on Computers*, 100(29):720–731, 1980.

[99] John F. Meyer. Performability: a Retrospective and Some Pointers to the Future. *Performance Evaluation*, 14(3-4):139–156, 1992.

[100] John F. Meyer. Performability Evaluation: Where It Is and What Lies Ahead. In *Proceedings of the IEEE International Computer Performance and Dependability Symposium (IPDS)*, pages 334–343, Erlangen, April 1995.

[101] John F. Meyer. Defining and evaluating resilience: A performability perspective. In *International Workshop on Performability Modeling of Computer and Communication Systems (PMCCS)*, September 2009.

[102] Andrew L. Reibman and Malathi Veeraraghavan. Reliability Modeling: An Overview for System Designers. *IEEE Computer*, 24(4):49–57, 1991.

[103] Poul E. Heegaard and Kishor S. Trivedi. Network survivability modeling. *Computer Networks*, 53(8):1215–1234, 2009.

[104] Kishor S. Trivedi, Dong Seong Kim, Arpan Roy, and Deep Medhi. Dependability and Security Models. In *Proceedings of the 7th IEEE International Workshop on the Design of Reliable Communication Networks (DRCN)*, pages 11–20, Washington, D.C., October 2009.

[105] Tadao Murata. Petri Nets: Properties, Analysis and Applications. *Proceedings of the IEEE*, 77(4):541–580, 1989.

[106] Manish Malhotra and Kishor S. Trivedi. Dependability Modeling Using Petri-Nets. *IEEE Transactions on Reliability*, 44(3):428–440, 1995.

[107] Piotr Cholda, Anders Mykkeltveit, Bjarne E. Helvik, Otto J. Wittner, and Andrzej Jajszczyk. A Survey of Resilience Differentiation Frameworks in Communication Networks. *IEEE Communications Surveys & Tutorials*, 9(4):32–55, 2007.

[108] Mohamed Al-Kuwaiti, Nicholas Kyriakopoulos, and Sayed Hussein. A Comparative Analysis of Network Dependability, Fault-tolerance, Reliability, Security, and Survivability. *IEEE Communications Surveys & Tutorials*, 11(2):106–124, 2009.

[109] Michele Nogueira Lima, Aldri Luiz dos Santos, and Guy Pujolle. A Survey of Survivability in Mobile Ad Hoc Networks. *IEEE Communications Surveys & Tutorials*, 11(1):66–77, 2009.

[110] Abdul Jabbar Mohammad, David Hutchison, and James P.G. Sterbenz. Towards Quantifying Metrics for Resilient and Survivable Networks. In *Proceedings of the 14th IEEE International Conference on Network Protocols (ICNP)*, pages 17–18, Santa Barbara, CA, November 2006.

[111] Abdul Jabbar. *A Framework to Quantify Network Resilience and Survivability*. PhD thesis, The University of Kansas, Lawrence, KS, May 2010.

[112] Abdul Jabbar, Hemanth Narra, and James P.G. Sterbenz. An Approach to Quantifying Resilience in Mobile Ad hoc Networks. In *Proceedings of the 8th IEEE International Workshop on the Design of Reliable Communication Networks (DRCN)*, pages 140–147, Krakow, October 2011.

[113] Robert J. Ellison, David A. Fisher, Richard C. Linger, Howard F. Lipson, Thomas A. Longstaff, and Nancy R. Mead. Survivability: Protecting Your Critical Systems. *IEEE Internet Computing*, 3(6):55–63, 1999.

[114] John C. Knight, Elisabeth A. Strunk, and Kevin J. Sullivan. Towards a Rigorous Definition of Information System Survivability. In *Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX)*, pages 78–89, Washington, D.C., April 2003.

[115] Piet Demeester, Michael Gryseels, Achim Autenrieth, Carlo Brianza, Laura Castagna, Giulio Signorelli, Roberto Clemente, Mauro Ravera, Andrzej Jajszczyk, Dariusz Janukowicz, Kristof Van Doorselaere, and Yohnosuke Harada. Resilience in Multilayer Networks. *IEEE Communications Magazine*, 37(8):70–76, 1999.

[116] Ali Zolfaghari and Fred J. Kaudel. Framework for Network Survivability Performance. *IEEE Journal on Selected Areas in Communications*, 12(1):46–51, 1994.

[117] John Knight and Elisabeth Strunk. Achieving critical system survivability through software architectures. In Rogério de Lemos, Cristina Gacek, and Alexander Romanovsky, editors, *Architecting Dependable Systems II*, volume 3069 of *Lecture Notes in Computer Science*, pages 69–91. Springer Berlin / Heidelberg, 2004.

[118] Mark Lanus, Liang Yin, and Kishor S. Trivedi. Hierarchical Composition and Aggregation of State-Based Availability and Performability Models. *IEEE Transactions on Reliability*, 52(1):44–52, 2003.

[119] Kishor S. Trivedi. *Probability and Statistics with Reliability, Queuing, and Computer Science Applications*. John Wiley and Sons, New York, 2nd edition, 2001.

[120] Sally Floyd and Vern Paxson. Difficulties in Simulating the Internet. *IEEE/ACM Transactions on Networking*, 9(4):392–403, 2001.

[121] Victor S. Frost and Benjamin Melamed. Traffic Modeling For Telecommunications Networks. *IEEE Communications Magazine*, 32(3):70–81, 1994.

[122] Clémence Magnien, Matthieu Latapy, and Jean-Loup Guillaume. Impact of Random Failures and Attacks on Poisson and Power-Law Random Networks. *ACM Computing Surveys*, 43(3):13:1–13:31, 2011.

[123] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74(1):47–97, 2002.

[124] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D.-U. Hwang. Complex networks: Structure and dynamics. *Physics Reports*, 424(4–5):175–308, 2006.

[125] John C. Doyle, David L. Alderson, Lun Li, Steven Low, Matthew Roughan, Stanislav Shalunov, Reiko Tanaka, and Walter Willinger. The "robust yet fragile" nature of the Internet. *Proceedings of the National Academy of Sciences of the United States of America*, 102(41):14497–14502, 2005.

[126] Seung-Taek Park, Alexy Khrabrov, David M. Pennock, Steve Lawrence, C. Lee Giles, and Lyle H. Ungar. Static and Dynamic Analysis of the Internet's Susceptibility to Faults and Attacks. In *Proceedings of the 22nd IEEE Conference on*

Computer Communications (INFOCOM), volume 3, pages 2144–2154, San Francisco, CA, April 2003.

[127] Damien Magoni. Tearing Down the Internet. *IEEE Journal on Selected Areas in Communications*, 21(6):949–960, 2003.

[128] Ali Sydney, Caterina Scoglio, Mina Youssef, and Phillip Schumm. Characterising the Robustness of Complex Networks. *International Journal of Internet Technology and Secured Transactions*, 2(3/4):291–320, 2010.

[129] Wojciech Molisz and Jacek Rak. End-to-end service survivability under attacks on networks. *Journal of Telecommunications and Information Technology*, 3:19–26, 2006.

[130] Jacek Rak and Krzysztof Walkowiak. Survivability of Anycast and Unicast Flows under Attacks on Networks. In *Proceedings of the 2nd IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, pages 497–503, Moscow, October 2010.

[131] Sebastian Neumayer and Eytan Modiano. Network Reliability With Geographically Correlated Failures. In *Proceedings of the 29th IEEE Conference on Computer Communications (INFOCOM)*, pages 1–9, San Diego, CA, March 2010.

[132] Sebastian Neumayer, Gil Zussman, Reuven Cohen, and Eytan Modiano. Assessing the Vulnerability of the Fiber Infrastructure to Disasters. In *Proceedings of the 28th IEEE Conference on Computer Communications (INFOCOM)*, pages 1566–1574, Rio de Janeiro, April 2009.

[133] Sebastian Neumayer, Gil Zussman, Reuven Cohen, and Eytan Modiano. Assessing the Impact of Geographically Correlated Network Failures. In *Proceedings of the*

IEEE Military Communications Conference (MILCOM), pages 1–6, San Diego, CA, November 2008.

[134] Sebastian Neumayer, Gil Zussman, Reuven Cohen, and Eytan Modiano. Assessing the Vulnerability of the Fiber Infrastructure to Disasters. *IEEE/ACM Transactions on Networking*, 19(6):1610–1623, 2011.

[135] Sebastian Neumayer, Alon Efrat, and Eytan Modiano. Geographic Max-Flow and Min-Cut Under a Circular Disk Failure Model. In *Proceedings of the 31st IEEE Conference on Computer Communications (INFOCOM)*, pages 2736–2740, Orlando, FL, March 2012.

[136] M. Todd Gardner and Cory Beard. Evaluating Geographic Vulnerabilities in Networks. In *Proceedings of the IEEE International Workshop on Communications Quality and Reliability (CQR)*, pages 1–6, Naples, FL, May 2011.

[137] M. Todd Gardner, Cory Beard, and Deep Medhi. Avoiding high impacts of geospatial events in mission critical and emergency networks using linear and swarm optimization. In *Proceedings of the IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, pages 264–271, New Orleans, LA, March 2012.

[138] Bijan Bassiri and Shahram Shah Heydari. Network Survivability in Large-Scale Regional Failure Scenarios. In *Proceedings of the 2nd Canadian Conference on Computer Science and Software Engineering (C3S2E)*, pages 83–87, Montreal, May 2009.

[139] Alireza Izaddoost and Shahram Shah Heydari. A probabilistic model for network survivability in large scale failure scenarios. In *Proceedings of the 25th IEEE Cana-*

dian Conference on Electrical Computer Engineering (CCECE), pages 1–5, Montreal, May 2012.

[140] David J. Houck, Eunyoung Kim, Gerard P. O'Reilly, David D. Picklesimer, and Huseyin Uzunalioglu. A Network Survivability Model for Critical National Infrastructures. *Bell Labs Technical Journal*, 8(4):153–172, 2004.

[141] Gerard O'Reilly, Ahmad Jrad, Ramesh Nagarajan, Theresa Brown, and Stephen Conrad. Critical Infrastructure Analysis of Telecom for Natural Disasters. In *Proceedings of the 12th IEEE International Telecommunications Network Strategy and Planning Symposium (NETWORKS)*, pages 1–6, New Delhi, November 2006.

[142] F. Lau, S.H. Rubin, M.H. Smith, and L. Trajkovic. Distributed Denial of Service Attacks. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, volume 3, pages 2275–2280, Nashville, TN, October 2000.

[143] Salim Hariri, Guangzhi Qu, Tushneem Dharmagadda, Modukuri Ramkishore, and Cauligi S. Raghavendra. Impact Analysis of Faults and Attacks in Large-Scale Networks. *IEEE Security & Privacy*, 1(5):49–54, 2003.

[144] Igor Kotenko and Alexander Ulanov. Simulation of Internet DDoS Attacks and Defense. In Sokratis Katsikas, Javier López, Michael Backes, Stefanos Gritzalis, and Bart Preneel, editors, *Information Security*, volume 4176 of *Lecture Notes in Computer Science*, pages 327–342. Springer Berlin / Heidelberg, 2006.

[145] George F. Riley, Monirul L. Sharif, and Wenke Lee. Simulating Internet Worms. In *Proceedings of the 12th IEEE Annual International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems (MASCOTS)*, pages 268–274, Volendam, October 2004.

[146] Scalable Simulation Framework. `http://www.ssfnet.org/`, February 2012.

[147] Ioannis Chatzigiannakis, Athanasios Kinalis, Georgios Mylonas, Sotiris Nikoletseas, Grigoris Prasinos, and Christos Zaroliagis. TRAILS, a Toolkit for Efficient, Realistic and Evolving Models of Mobility, Faults and Obstacles in Wireless Networks. In *Proceedings of the 41st Annual Simulation Symposium (ANSS)*, pages 23–32, Ottawa, April 2008.

[148] Ioannis Chatzigiannakis, Georgios Mylonas, and Sotiris Nikoletseas. Modeling and Evaluation of the Effect of Obstacles on the Performance of Wireless Sensor Networks. In *Proceedings of the 39th Annual Simulation Symposium (ANSS)*, pages 50–60, Huntsville, AL, April 2006.

[149] Dan S. Broyles. Benchmarking Wireless Network Protocols: Threat and Challenge Analysis of the AeroRP. Master's thesis, The University of Kansas, Lawrence, KS, July 2011.

[150] Dongsheng Zhang, Santosh Ajith Gogi, Dan S. Broyles, Egemen K. Çetinkaya, and James P.G. Sterbenz. Modelling Attacks and Challenges to Wireless Networks. In *Proceedings of the 4th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, pages 806–812, St. Petersburg, October 2012.

[151] Dongsheng Zhang, Santosh Ajith Gogi, Dan S. Broyles, Egemen K. Çetinkaya, and James P.G. Sterbenz. Modelling Wireless Challenges. In *Proceedings of the 18th ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, pages 423–425, Istanbul, August 2012. Extended Abstract.

[152] David L. Tennenhouse, Jonathan M. Smith, W. David Sincoskie, David J. Wetherall, and Gary J. Minden. A Survey of Active Network Research. *IEEE Communications Magazine*, 35(1):80–86, 1997.

[153] Andrew T. Campbell, Herman G. De Meer, Michael E. Kounavis, Kazuho Miki, John B. Vicente, and Daniel Villela. A Survey of Programmable Networks. *ACM Computer Communication Review*, 29(2):7–23, 1999.

[154] Kurt Tutschku, Phuoc Tran-Gia, and Frank-Uwe Andersen. Trends in network and service operation for the emerging future Internet. *AEU - International Journal of Electronics and Communications*, 62(9):705–714, 2008.

[155] N.M. Mosharaf Kabir Chowdhury and Raouf Boutaba. A survey of network virtualization. *Computer Networks*, 54(5):862–876, 2010.

[156] GENI: Global Environment for Network Innovations. `http://www.geni.net`.

[157] FIRE: Future Internet Research & Experimentation. `http://cordis.europa.eu/fp7/ict/fire/`.

[158] Terry Benzel, Robert Braden, Dongho Kim, Clifford Neuman, Anthony Joseph, Keith Sklower, Ron Ostrenga, and Stephen Schwab. Design, Deployment, and Use of the DETER Testbed. In *Proceedings of the DETER Community Workshop on Cyber Security Experimentation and Test*, Boston, MA, August 2007.

[159] Jelena Mirkovic, Terry V. Benzel, Ted Faber, Robert Braden, John T. Wroclawski, and Stephen Schwab. The DETER Project: Advancing the Science of Cyber Security Experimentation and Test. In *Proceedings of the IEEE International Conference on Technologies for Homeland Security (HST)*, pages 1–7, Boston, MA, November 2010.

[160] Joan Calvet, Carlton R. Davis, José M. Fernandez, Wadie Guizani, Mathieu Kaczmarek, Jean-Yves Marion, and Pier-Luc St-Onge. Isolated virtualised clusters: testbeds for high-risk security experimentation and training. In *Proceedings of the*

3rd USENIX International Workshop on Cyber Security Experimentation and Test (CSET), pages 1–8, Washington, D.C., August 2010.

[161] Egemen K. Çetinkaya and James P.G. Sterbenz. A Taxonomy of Network Challenges. In *Proceedings of the 9th IEEE/IFIP International Conference on the Design of Reliable Communication Networks (DRCN)*, pages 322–330, Budapest, March 2013.

[162] Renesys – The Internet Intelligence Authority. `http://www.renesys.com`, January 2012.

[163] Arbor Networks. `http://www.arbornetworks.com/`, January 2012.

[164] BGPmon.net, a BGP monitoring and analyzer tool. `http://bgpmon.net/`, January 2012.

[165] RIPE Network Coordination Centre. `http://www.ripe.net/`, January 2012.

[166] ICANN – Internet Corporation for Assigned Names and Numbers. `http://www.icann.org/`, January 2012.

[167] NANOG – North American Network Operators' Group. `http://www.nanog.org/`, January 2012.

[168] ENISA – The European Network and Information Security Agency. `http://www.enisa.europa.eu/`, October 2012.

[169] Carl E. Landwehr, Alan R. Bull, John P. McDermott, and William S. Choi. A Taxonomy of Computer Program Security Flaws. *ACM Computing Surveys*, 26(3):211–254, 1994.

[170] Nicholas Weaver, Vern Paxson, Stuart Staniford, and Robert Cunningham. A Taxonomy of Computer Worms. In *Proceedings of the ACM Workshop on Rapid Malcode (WORM)*, pages 11–18, Washington, D.C., October 2003.

[171] Vinay M. Igure and Ronald D. Williams. Taxonomies of Attacks and Vulnerabilities in Computer Systems. *IEEE Communications Surveys & Tutorials*, 10(1):6–19, 2008.

[172] Anirban Chakrabarti and G. Manimaran. Internet Infrastructure Security: A Taxonomy. *IEEE Network*, 16(6):13–21, 2002.

[173] B. Harris and R. Hunt. TCP/IP security threats and attack methods. *Computer Communications*, 22(10):885–897, 1999.

[174] S. M. Bellovin. Security Problems in the TCP/IP Protocol Suite. *ACM Computer Communication Review*, 19(2):32–48, 1989.

[175] Carl E. Landwehr and David M. Goldschlag. Security Issues in Networks with Internet Access. *Proceedings of the IEEE*, 85(12):2034–2051, 1997.

[176] Jelena Mirkovic and Peter Reiher. A Taxonomy of DDoS Attack and DDoS Defense Mechanisms. *ACM Computer Communication Review*, 34(2):39–53, 2004.

[177] Christos Douligeris and Aikaterini Mitrokotsa. DDoS attacks and defense mechanisms: classification and state-of-the-art. *Computer Networks*, 44(5):643–666, 2004.

[178] Tao Peng, Christopher Leckie, and Kotagiri Ramamohanarao. Survey of Network-Based Defense Mechanisms Countering the DoS and DDoS Problems. *ACM Computing Surveys*, 39(1):3:1–3:42, 2007.

[179] Fred B. Schneider, Steven M. Bellovin, Martha Branstad, J. Randall Catoe, Stephen D. Crocker, Charlie Kaufman, Stephen T. Kent, John C. Knight, Steven

McGeady, Ruth R. Nelson, Allan M. Schiffman, George A. Spix, and Doug Tygar. *Trust in Cyberspace*. National Academy Press, Washington, D.C., 1999.

[180] Alin Popescu, Todd Underwood, and Earl Zmijewski. Quaking Tables: The Taiwan Earthquakes and the Internet Routing Table. In *APRICOT*, Bali, 2007. Renesys Corp.

[181] Yasuichi Kitamura, Youngseok Lee, Ryo Sakiyama, and Koji Okamura. Experience with Restoration of Asia Pacific Network Failures from Taiwan Earthquake. *IEICE Transactions on Communications*, E90-B(11):3095–3103, 2007.

[182] Yang Ran. Considerations and Suggestions on Improvement of Communication Network Disaster Countermeasures after the Wenchuan Earthquake. *IEEE Communications Magazine*, 49(1):44–47, 2011.

[183] Alexis Kwasinski and Alex K. Tang. Telecommunications Performance in the M=9.0 Off-shore East Coast of Japan Earthquake and Tsunami, March 11, 2011. In *Proceedings of the International Symposium on Engineering Lessons Learned from the 2011 Great East Japan Earthquake*, Tokyo, March 2012.

[184] James Cowie. Japan Quake. `http://www.renesys.com/blog/2011/03/japan-quake.shtml`, March 2011.

[185] Kensuke Fukuda, Michihiro Aoki, Shunji Abe, Yuseng Ji, Michihiro Koibuchi, Motonori Nakamura, Shigeki Yamada, and Shigeo Urushidani. Impact of Tohoku Earthquake on R&E Network in Japan. In *Proceedings of the ACM Special Workshop on Internet and Disasters (SWID)*, pages 1:1–1:6, Tokyo, December 2011.

[186] Kenjiro Cho, Cristel Pelsser, Randy Bush, and Youngjoon Won. The Japan earthquake: the impact on traffic and routing observed by a local ISP. In *Proceedings*

of the ACM Special Workshop on Internet and Disasters (SWID), pages 2:1–2:8, Tokyo, December 2011.

[187] Alexis Kwasinski. Effects of Notable Natural Disasters from 2005 to 2011 on Telecommunications Infrastructure: Lessons from on-site Damage Assessments. In *Proceedings of the 33rd IEEE International Telecommunications Energy Conference (INTELEC)*, pages 1–9, Amsterdam, October 2011.

[188] Rui Chen, John Coles, Jinkyu Lee, and H. Raghav Rao. Emergency Communication and System Design: The Case of Indian Ocean Tsunami. In *Proceedings of the International Conference on Information and Communication Technologies and Development (ICTD)*, pages 300–309, Doha, April 2009.

[189] Tom Davis et al. A Failure of Initiative: The Final Report of the Select Bipartisan Committee to Investigate the Preparation for and Response to Hurricane Katrina. Congressional Report H.Rpt. 109-377, US House of Representatives, Washington, D.C., 2006.

[190] James Cowie, Alin Popescu, and Todd Underwood. Impact of Hurricane Katrina on Internet Infrastructure. Technical report, Renesys, September 2005.

[191] Martin A. Brown. Ike Hammers Texas Internet. `http://www.renesys.com/blog/2008/09/ike-hammers-texas-internet.shtml`, September 2008.

[192] Martin A. Brown. Ike Brings Biggest Multi-State Internet Outage since 2003. `http://www.renesys.com/blog/2008/09/ike-brings-biggest-multistate.shtml`, September 2008.

[193] Supaporn Erjongmanee and Chuanyi Ji. Large-Scale Network-Service Disruption: Dependencies and External Factors. *IEEE Transactions on Network and Service Management*, 8(4):375–386, 2011.

[194] James Cowie. Irene Wallops US Internet. http://www.renesys.com/blog/2011/08/irene-wallops-us-internet.shtml, August 2011.

[195] Earl Zmijewski. ISPs Learn from Katrina, Survive Gustav. http://www.renesys.com/blog/2008/09/isps-learn-from-katrina-surviv.shtml, September 2008.

[196] Earl Zmijewski. Gustav: 3 days later. http://www.renesys.com/blog/2008/09/gustav-3-days-later.shtml, September 2008.

[197] Doug Madory. Hurricane Sandy: Initial Impact. http://www.renesys.com/blog/2012/10/hurricane-sandy-initial-impact.shtml, October 2012.

[198] James Cowie. Hurricane Sandy: Outage Animation. http://www.renesys.com/blog/2012/10/hurricane-sandy-outage-animati.shtml, October 2012.

[199] Doug Madory. Hurricane Sandy: Global Impacts. http://www.renesys.com/blog/2012/11/sandys-global-impacts.shtml, November 2012.

[200] Emile Aben. RIPE Atlas: Hurricane Sandy and How the Internet Routes Around Damage. https://labs.ripe.net/Members/emileaben/ripe-atlas-hurricane-sandy-global-effects, November 2012.

[201] Emile Aben. RIPE Atlas - Superstorm Sandy. https://labs.ripe.net/Members/emileaben/ripe-atlas-superstorm-sandy, October 2012.

[202] John Heidemann, Lin Quan, and Yuri Pradkin. A Preliminary Analysis of Network Outages During Hurricane Sandy. Technical Report ISI-TR-2008-685b, USC/Information Sciences Institute, November 2012. (correction Feb. 2013).

[203] Gerry Smith. Verizon Outages Strand Lower Manhattan Businesses 4 Months After Sandy. http://www.huffingtonpost.com/2013/02/28/verizon-sandy-new-york_n_2782664.html, February 2013.

[204] Aaron Schulman and Neil Spring. Pingin' in the Rain. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 19–28, Berlin, November 2011.

[205] Abdul Jabbar, Bharatwajan Raman, Victor S. Frost, and James P.G. Sterbenz. Weather disruption-tolerant self-optimising millimeter mesh networks. In *Proceedings of the 3rd International IFIP Workshop on Self-Organizing Systems (IWSOS)*, volume 5343 of *Lecture Notes in Computer Science*, pages 242–255, Vienna, December 2008. Springer.

[206] Abdul Jabbar, Justin P. Rohrer, Andrew Oberthaler, Egemen K. Çetinkaya, Victor S. Frost, and James P.G. Sterbenz. Performance Comparison of Weather Disruption-Tolerant Cross-Layer Routing Algorithms. In *Proceedings of the 28th IEEE Conference on Computer Communications (INFOCOM)*, pages 1143–1151, Rio de Janeiro, April 2009.

[207] Abdul Jabbar, Justin P. Rohrer, Victor S. Frost, and James P.G. Sterbenz. Survivable Millimeter-Wave Mesh Networks. *Computer Communications*, 34(16):1942–1955, 2011.

[208] Robert H. Johns, Jeffry S. Evans, and Stephen F. Corfidi. About Derechos. `http://www.spc.noaa.gov/misc/AbtDerechos/derechofacts.htm`, November 2012.

[209] Public Safety and Homeland Security Bureau. Impact of the June 2012 Derecho on Communications Networks and Services. Report and recommendations, Federal Communications Commission (FCC), January 2013.

[210] Kathleen F. Jones and Nathan D. Mulherin. An Evaluation of the Severity of the January 1998 Ice Storm in Northern New England. Technical report, Federal Emergency Management Agency, Hanover, NH, April 1998.

[211] Severe Space Weather Events: Understanding Societal and Economic Impacts. Workshop report, National Research Council, 2008.

[212] R. Sanders. Effect of Terrestrial Electromagnetic Storms on Wireline Communications. *IRE Transactions on Communication Systems*, 9(4):367–377, 1961.

[213] D. H. Boteler, R. J. Pirjola, and H. Nevanlinna. The effects of geomagnetic disturbances on electrical systems at the earth's surface. *Advances in Space Research*, 22(1):17–27, 1998.

[214] John Kappenman. A Perfect Storm of Planetary Proportions. *IEEE Spectrum Magazine*, 49(2):26–31, 2012.

[215] John S. Foster, Earl Gjelde, William R. Graham, Robert J. Hermann, Henry (Hank) M. Kluepfel, Richard L. Lawson, Gordon K. Soper, Lowell L. Wood, and Joan B. Woodard. Report of the Commission to Assess the Threat to the United States from Electromagnetic Pulse (EMP) Attack. Volume 1: Executive Report. Technical report, Electromagnetic Pulse (EMP) Commission, 2004.

[216] Colin R. Miller. Electromagnetic Pulse Threats in 2010. Technical report, Center for Strategy and Technology, Air War College, Air University, Maxwell AFB, AL, November 2005.

[217] John S. Foster, Earl Gjelde, William R. Graham, Robert J. Hermann, Henry (Hank) M. Kluepfel, Richard L. Lawson, Gordon K. Soper, Lowell L. Wood, and Joan B. Woodard. Report of the Commission to Assess the Threat to the United States from Electromagnetic Pulse (EMP) Attack. Critical National Infrastructures. Technical report, Electromagnetic Pulse (EMP) Commission, McLean, VA, April 2008.

[218] Clay Wilson. High Altitude Electromagnetic Pulse (HEMP) and High Power Microwave (HPM) Devices: Threat Assessments. Technical report, Congressional Research Service, Washington, D.C., July 2008.

[219] Satellite Collision Leaves Significant Debris Clouds. *NASA Orbital Debris Quarterly News*, 13(2), April 2009.

[220] Tom Clark. Iridium Satellite Collision Update on Satellite and Performance Status. `http://www.mssnews.com/archives/2009/02/entry_64.html`, February 2009.

[221] Iridium Communications Inc. Form 10-K Annual Report for Year Ending 2010. `http://investor.iridium.com/secfiling.cfm?filingID=1193125-11-57741`, March 2011.

[222] Jessica A. Steinberger. A Survey of Satellite Communications System Vulnerabilities. Master's thesis, Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, June 2008.

[223] Extended loss of GPS Impact on Reliability. White paper, North American Electric Reliability Corporation (NERC), December 2012.

[224] James A. Eibel. Fire Prevention in Telecommunications Facilities. Technical report, Network Reliability and Interoperability Council, 1993.

[225] John C. McDonald. Public Network Integrity–Avoiding a Crisis in Trust. *IEEE Journal on Selected Areas in Communications*, 12(1):5–12, 1994.

[226] Glenn Zorpette. Keeping the phone lines open. *IEEE Spectrum Magazine*, 26(6):32–36, 1989.

[227] Patrick Townson. The Great Fire. `http://massis.lcs.mit.edu/archives/history/fire.in.chicago.5-88`, May 1988. Online; accessed 15-June-2011.

[228] Hilary C. Styron. CSX Tunnel Fire: Baltimore, Maryland. US Fire Administration Technical Report USFA-TR-140, Federal Emergency Management Agency, Emmitsburg, MD, July 2001.

[229] CSX Freight Train Derailment and Subsequent Fire in the Howard Street Tunnel in Baltimore, Maryland, on July 18, 2001. Railroad accident brief, National Transportation Safety Board, Washington, D.C.

[230] Mark R. Carter, Mark P. Howard, Nicholas Owens, David Register, Jason Kennedy, Kelley Pecheux, and Aaron Newton. Effects of Catastrophic Events on Transportation System Management and Operations, Howard Street Tunnel Fire, Baltimore City, Maryland – July 18, 2001. Technical report, U.S. Department of Transportation, Washington, D.C., 2002.

[231] Xiaoliang Zhao, Daniel Massey, Mohit Lad, and Lixia Zhang. ON/OFF Model: A New Tool to Understand BGP Update Burst. USC-CSD Technical Report 04-819, August 2004.

[232] C. Chern and W. J. Tudor. OTEC Submarine Cable Environmental Characteristics and Hazards Analysis. In *Proceedings of the IEEE OCEANS Conference*, pages 502–507, Washington, D.C., September 1982.

[233] Louis J. Marra. Sharkbite on the SL Submarine Lightwave Cable System: History, Causes and Resolution. *IEEE Journal of Oceanic Engineering*, 14(3):230–237, 1989.

[234] Mayada Omer, Roshanak Nilchiani, and Ali Mostashari. Measuring the Resilience of the Trans-Oceanic Telecommunication Cable System. *IEEE Systems Journal*, 3(3):295–303, 2009.

[235] A. Antony, L. Cittadini, D. Karrenberg, R. Kisteleki, T. Refice, T. Vest, and R. Wilhelm. Mediterranean Fiber Cable Cut (January-February 2008) Analysis

of Network Dynamics. Technical Report RT-DIA-124-2008, Università degli Studi di Roma Tre, Roma, Italy, March 2008.

[236] Earl Zmijewski. Mediterranean Cable Break. http://www.renesys.com/blog/2008/01/mediterranean_cable_break.shtml, January 2008.

[237] Earl Zmijewski. Mediterranean Cable Break - Part II. http://www.renesys.com/blog/2008/01/mediterranean-cable-break-part-1.shtml, January 2008.

[238] Earl Zmijewski. Mediterranean Cable Break - Part III. http://www.renesys.com/blog/2008/02/mediterranean_cable_break_part.shtml, February 2008.

[239] Tomasz Bilski. Disaster's Impact on Internet Performance – Case Study. In Andrzej Kwiecień, Piotr Gaj, and Piotr Stera, editors, *Computer Networks*, volume 39 of *Communications in Computer and Information Science*, pages 210–217. Springer Berlin Heidelberg, 2009.

[240] Greg's Cable Map. http://www.cablemap.info/, November 2012.

[241] Daniel E. Crawford. Fiber Optic Cable Dig-ups: Causes and Cures. Technical report, Network Reliability and Interoperability Council, 1993.

[242] Harold T. Daugherty and William J. Klein. U.S. Network Reliability Issues and Major Outage Performance. In *Proceedings of the IEEE Symposium on Computers and Communications (ISCC)*, pages 114–119, Alexandria, June 1995.

[243] Tom Parfitt. Georgian woman cuts off web access to whole of Armenia. http://www.guardian.co.uk/world/2011/apr/06/georgian-woman-cuts-web-access, April 2011.

[244] Olav Lysne and Amund Kvalbein. Measurements from Telenor's outages. http://www.http://simula.no/news/measurements-from-telenors-outages, 2011.

[245] Fred Lawler. The 10 Most Bizarre and Annoying Causes of Fiber Cuts. http://blog.level3.com/2011/08/04/the-10-most-bizarre-and-annoying-causes-of-fiber-cuts/, August 2011.

[246] Pandemic Influenza Preparedness, Response, and Recovery Guide for Critical Infrastructure and Key Resources. Guide for Critical Infrastructure and Key Resources, Department of Homeland Security (DHS), September 2006.

[247] Pandemic Influenza Impact on Communications Networks Study. Unclassified, Department of Homeland Security (DHS), December 2007.

[248] Jean-Claude Laprie, Karama Kanoun, and Mohamed Kaâniche. Modelling Interdependencies Between the Electricity and Information Infrastructures. In Francesca Saglietti and Norbert Oster, editors, *Computer Safety, Reliability, and Security*, volume 4680 of *Lecture Notes in Computer Science*, pages 54–67. Springer Berlin / Heidelberg, 2007.

[249] Benoît Robert and Luciano Morabito. Dependency on Electricity and Telecommunications. In Zofia Lukszo, Geert Deconinck, and Margot P. C. Weijnen, editors, *Securing Electricity Supply in the Cyber Age*, volume 15 of *Topics in Safety, Risk, Reliability and Quality*, pages 33–52. Springer Netherlands, 2010.

[250] Dennis McGrath. Measuring the 4:11 Effect: The Power Failure and the Internet. *IEEE Security & Privacy*, 1(5):16–18, 2003.

[251] Sergey V. Buldyrev, Roni Parshani, Gerald Paul, H. Eugene Stanley, and Shlomo Havlin. Catastrophic cascade of failures in interdependent networks. *Nature*, 464(7291):1025–1028, 2010.

[252] S. Pahwa, A. Hodges, C. Scoglio, and S. Wood. Topological Analysis of the Power Grid and Mitigation Strategies Against Cascading Failures. In *4th Annual IEEE Systems Conference*, pages 272–276, San Diego, CA, April 2010.

[253] Andrey Bernstein, Daniel Bienstock, David Hay, Meric Uzunoglu, and Gil Zussman. Sensitivity Analysis of the Power Grid Vulnerability to Large-Scale Cascading Failures. *ACM SIGMETRICS Performance Evaluation Review*, 40(3):33–37, 2012.

[254] James Cowie. Lights Out in Rio. `http://www.renesys.com/blog/2009/11/lights-out-in-rio.shtml`, November 2009.

[255] James H. Cowie, Andy T. Ogielski, Brian J. Premore, Eric A. Smith, and Todd Underwood. Impact of the 2003 Blackouts on Internet Communications. Technical report, Renesys Corporation, 2003.

[256] Jun Li, Zhen Wu, and Eric Purpus. Toward Understanding the Behavior of BGP During Large-Scale Power Outages. In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*, pages 1–5, San Francisco, CA, November 2006.

[257] A. Snow, K. Chatanyam, G. Weckman, and P. Campbell. Power Related Network Outages: Impact, Triggering Events, And Root Causes. In *Proceedings of the 10th IEEE/IFIP Network Operations and Management Symposium (NOMS)*, pages 1–4, Vancouver, April 2006.

[258] Saswato R. Das. Northern India Recovering from Huge Blackout. `http://spectrum.ieee.org/tech-talk/energy/the-smarter-grid/northern-india-recovering-from-huge-blackout`, July 2012.

[259] Joshua Romero. Lack of Rain a Leading Cause of Indian Grid Collapse. `http://spectrum.ieee.org/energywise/energy/the-smarter-grid/`

disappointing-monsoon-season-wreaks-havoc-with-indias-grid/, July 2012.

[260] The Enquiry Committee. Grid Disturbance in Northern Region on 30th July and in Northern, Eastern, North-Eastern Region on 31st July 2012. Technical report, 2012.

[261] D. Richard Kuhn. Sources of Failure in the Public Switched Telephone Network. *IEEE Computer*, 30(4):31–36, 1997.

[262] Ghi Paul Im and Richard L. Baskerville. A Longitudinal Study of Information System Threat Categories: The Enduring Problem of Human Error. *ACM SIGMIS Database*, 36(4):68–79, 2005.

[263] Aaron B. Brown. Oops! Coping with Human Error in IT Systems. *ACM Queue*, 2(8):34–41, 2004.

[264] Garima Bajaj. A Pre and Post 9-11 Analysis of SS7 Outages in the Public Switched Telephone Network. Master's thesis, Ohio University, Athens, OH, March 2007.

[265] R. Bush and D. Meyer. Some Internet Architectural Guidelines and Philosophy. RFC 3439 (Informational), December 2002.

[266] Erik Romijn. RIPE NCC and Duke University BGP Experiment. https://labs.ripe.net/Members/erik/ripe-ncc-and-duke-university-bgp-experiment/, August 2010.

[267] Earl Zmijewski. Longer is not always better. http://www.renesys.com/blog/2009/02/longer-is-not-better.shtml, February 2009.

[268] Nick Heath. Global outage takes down sites and services across the internet. http://www.silicon.com/technology/networks/2011/11/07/

global-outage-takes-down-sites-and-services-across-the-internet-39748193/, November 2011.

[269] Jim Duffy. Juniper at the root of Internet outage? http://www.networkworld.com/news/2011/110711-internet-outage-252851.html?hpg1=bn, November 2011.

[270] Mark Imbriaco. Network problems last Friday. https://github.com/blog/1346-network-problems-last-friday, December 2012.

[271] Leonard Kleinrock, Cynthia H. Braddon, David D. Clark, William J. Emery, David J. Farber, A.G. Fraser, Russell D. Hensley, Lawrence H. Landweber, Robert W. Lucky, Susan K. Nutter, Radia Perlman, Susanna Schweizer, Connie Danner Stout, Charles Ellett Taylor, Thomas W. West, and Robert E. Kahn. *Realizing the Information Future: The Internet and Beyond*. The National Academy Press, Washington, D.C., 1994.

[272] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-To-End Arguments in System Design. *ACM Transactions on Computer Systems*, 2(4):277–288, 1984.

[273] Marjory S. Blumenthal and David D. Clark. Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World. *ACM Transactions on Internet Technology*, 1(1):70–109, 2001.

[274] Eric Schmidt, Terrence McGarty, Anthony S. Acampora, Walter S. Baer, Fred Baker, Andrew Blau, Deborah Estrin, Christian Huitema, Edward Jung, David A. Kettler, John C. Klensin, Milo Medin, Craig Partridge, and Daniel Schutzer. *The Internet's Coming of Age*. The National Academy Press, Washington, D.C., 2001.

[275] P.V. Mockapetris. Domain names - concepts and facilities. RFC 1034 (Standard), November 1987. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 2065, 2181, 2308, 2535, 4033, 4034, 4035, 4343, 4035, 4592.

[276] P.V. Mockapetris. Domain names - implementation and specification. RFC 1035 (Standard), November 1987. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 1995, 1996, 2065, 2136, 2181, 2137, 2308, 2535, 2845, 3425, 3658, 4033, 4034, 4035, 4343.

[277] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271 (Draft Standard), January 2006.

[278] Geoff Huston. Blurring the Lines. `http://www.potaroo.net/ispcol/2003-11/blurring.html`, November 2003.

[279] P. Traina. Experience with the BGP-4 protocol. RFC 1773 (Informational), March 1995.

[280] Dan Pei, Lixia Zhang, and Dan Massey. A Framework for Resilient Internet Routing Protocols. *IEEE Network*, 18(2):5–12, 2004.

[281] S. Murphy. BGP Security Vulnerabilities Analysis. RFC 4272 (Informational), January 2006.

[282] Ola Nordström and Constantinos Dovrolis. Beware of BGP Attacks. *ACM Computer Communication Review*, 34(2):1–8, 2004.

[283] Martin O. Nicholes and Biswanath Mukherjee. A Survey of Security Techniques for the Border Gateway Protocol (BGP). *IEEE Communications Surveys & Tutorials*, 11(1):52–65, 2009.

[284] Kevin Butler, Toni R. Farley, Patrick McDaniel, and Jennifer Rexford. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, 98(1):100–122, 2010.

[285] Geoff Huston, Mattia Rossi, and Grenville Armitage. Securing BGP – A Literature Survey. *IEEE Communications Surveys & Tutorials*, 13(2):199–222, 2011.

[286] Jun Li, Dejing Dou, Zhen Wu, Shiwoong Kim, and Vikash Agarwal. An Internet Routing Forensics Framework for Discovering Rules of Abnormal BGP Events. *ACM Computer Communication Review*, 35(5):55–66, 2005.

[287] James Cowie, Andy T. Ogielski, B. J. Premore, and Yougu Yuan. Internet worms and global routing instabilities. In *Proceedings of the SPIE Scalability and Traffic Control in IP Networks*, volume 4868, pages 195–199, Boston, MA, July 2002.

[288] Lan Wang, Xiaoliang Zhao, Dan Pei, Randy Bush, Daniel Massey, Allison Mankin, S. Felix Wu, and Lixia Zhang. Observation and Analysis of BGP Behavior under Stress. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement (IMW)*, pages 183–195, Marseille, November 2002.

[289] Mohit Lad, Xiaoliang Zhao, Beichuan Zhang, Dan Massey, and Lixia Zhang. Analysis of BGP Update Surge during Slammer Worm Attack. In Samir Das and Sajal Das, editors, *Distributed Computing - IWDC*, volume 2918 of *Lecture Notes in Computer Science*, pages 833–835. Springer Berlin / Heidelberg, 2003.

[290] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the Slammer Worm. *IEEE Security & Privacy*, 1(4):33–39, 2003.

[291] RIPE NCC Routing Information Service (RIS). `http://www.ripe.net/data-tools/stats/ris/routing-information-service`, 2011.

[292] MICHAEL Michael Bailey, Evan Cooke, Farnam Jahanian, David Watson, and Jose Nazario. The Blaster Worm: Then and Now. *IEEE Security & Privacy*, 3(4):26–31, 2005.

[293] Ricardo V. Oliveira, Rafit Izhak-Ratzin, Beichuan Zhang, and Lixia Zhang. Measurement of Highly Active Prefixes in BGP. In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*, volume 2, pages 894–898, St. Louis, MO, November 2005.

[294] Hitesh Ballani, Paul Francis, and Xinyang Zhang. A Study of Prefix Hijacking and Interception in the Internet. In *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 265–276, Kyoto, 2007.

[295] Mohit Lad, Dan Massey, Dan Pei, Yiguo Wu, Beichuan Zhang, and Lixia Zhang. PHAS: A Prefix Hijack Alert System. In *Proceedings of the 15th Conference on USENIX Security Symposium*, volume 15, pages 153–166, Vancouver, B.C., August 2006.

[296] Mohit Lad, Ricardo Oliveira, Beichuan Zhang, and Lixia Zhang. Understanding Resiliency of Internet Topology Against Prefix Hijack Attacks. In *Proceedings of the 37th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 368–377, Edinburgh, June 2007.

[297] Khin Thida Latt, Yasuhiro Ohara, Satoshi Uda, and Yoichi Shinoda. Analysis of IP Prefix Hijacking and Traffic Interception. *International Journal of Computer Science and Network Security*, 10(7):22–31, 2010.

[298] Xiaoliang Zhao, Dan Pei, Lan Wang, Dan Massey, Allison Mankin, S. Felix Wu, and Lixia Zhang. An Analysis of BGP Multiple Origin AS (MOAS) Conflicts. In *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement (IMW)*, pages 31–35, San Francisco, CA, November 2001.

[299] Ratul Mahajan, David Wetherall, and Tom Anderson. Understanding BGP Misconfiguration. In *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 3–16, Pittsburgh, PA, August 2002.

[300] Anirudh Ramachandran and Nick Feamster. Understanding the Network-Level Behavior of Spammers. In *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 291–302, Pisa, 2006.

[301] Laris Benkis. Practical BGP Security: Architecture, Techniques and Tools. White paper, Renesys, 2005.

[302] Stephen A. Misel. Wow, AS7007! `http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html`, April 1997.

[303] Avi Freedman. 7007: FROM THE HORSE'S MOUTH. `http://www.merit.edu/mail.archives/nanog/1997-04/msg00380.html`, April 1997.

[304] Vincent J. Bono. 7007 Explanation and Apology. `http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html`, April 1997.

[305] Andrew Partan. MAI.net filters. `http://www.merit.edu/mail.archives/nanog/1997-04/msg00469.html`, April 1997.

[306] Alin C. Popescu, Brian J. Premore, and Todd Underwood. The Anatomy of a Leak: AS9121. Presentation, Renesys Corporation, May 2005.

[307] Todd Underwood. Internet-Wide Catastrophe—Last Year. `http://www.renesys.com/blog/2005/12/internetwide_nearcatastrophela.shtml`, December 2005.

[308] Todd Underwood. Con-Ed Steals the 'Net. http://www.renesys.com/blog/2006/01/coned-steals-the-net.shtml, January 2006.

[309] Todd Underwood. What Really Caused the Panix Outage. http://www.renesys.com/blog/2006/01/what-really-caused-the-panix-o.shtml, January 2006.

[310] Martin A. Brown. Pakistan hijacks YouTube. http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml, February 2008.

[311] Martin A. Brown, Todd Underwood, and Earl Zmijewski. The Day the YouTube Died. Presentation, Renesys Corporation, Brooklyn, NY, June 2008.

[312] YouTube Hijacking: A RIPE NCC RIS case study . http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study, March 2008.

[313] Philip Hunter. Pakistan YouTube block exposes fundamental Internet security weakness: Concern that Pakistani action affected YouTube access elsewhere in world. *Computer Fraud & Security*, 2008(4):10–11, 2008.

[314] Danny McPherson. Internet Routing Insecurity::Pakistan Nukes YouTube? http://asert.arbornetworks.com/2008/02/internet-routing-insecuritypakistan-nukes-youtube/, February 2008.

[315] Danny McPherson. Africa Online Kenya Latest Internet Routing Insecurity Casualty. http://asert.arbornetworks.com/2008/03/africa-online-kenya-latest-internet-routing-insecurity-casuality/, March 2008.

[316] Andree Toonk. Prefix hijack by AS16735. http://bgpmon.net/blog/?p=80, November 2008.

[317] James Cowie. Brazil Leak: If a tree falls in the rainforest.... http://www.renesys.com/blog/2008/11/brazil-leak-if-a-tree-falls-in.shtml, November 2008.

[318] Danny McPherson. When Hijacking the Internet.... http://asert.arbornetworks.com/2008/11/when-hijacking-the-internet/, November 2008.

[319] Andree Toonk. Chinese ISP hijacks the Internet. http://bgpmon.net/blog/?p=282, April 2010.

[320] Andree Toonk. Chinese BGP hijack, putting things into perspective. http://bgpmon.net/blog/?p=323, November 2010.

[321] James Cowie. China's 18-Minute Mystery. http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml, November 2010.

[322] U.S.-China Economic and Security Review Commission. 2010 Annual Report to Congress. Technical report, One Hundred Eleventh Congress, Washington, D.C., November 2010.

[323] Craig Labovitz. China Hijacks 15% of Internet Traffic? http://asert.arbornetworks.com/2010/11/china-hijacks-15-of-internet-traffic/, November 2010.

[324] Cisco Security Advisory: Cisco IOS XR Software Border Gateway Protocol Vulnerability. http://www.cisco.com/warp/public/707/cisco-sa-20100827-bgp.shtml, September 2010.

[325] Earl Zmijewski. Reckless Driving on the Internet. http://www.renesys.com/blog/2009/02/the-flap-heard-around-the-worl.shtml, February 2009.

[326] Danny McPherson. Ahh, The Ease of Introducing Global Routing Instability. http://asert.arbornetworks.com/2009/02/ahh-the-ease-of-introducing-global-routing-instability/, February 2009.

[327] Andree Toonk. Long AS paths causing commotion. http://bgpmon.net/blog/?p=125, February 2009.

[328] Tao Wan and Paul C. van Oorschot. Analysis of BGP Prefix Origins During Google's May 2005 Outage. In *Proceedings of the 20th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, pages 1–8, Rhodes Island, April 2006.

[329] Andree Toonk. How the Internet in Australia went down under. http://www.bgpmon.net/how-the-internet-in-australia-went-down-under/, February 2012.

[330] Paul V. Mockapetris and Kevin J. Dunlap. Development of the Domain Name System. *ACM Computer Communication Review*, 25(1):112–122, 1995.

[331] Vasileios Pappas, Zhiguo Xu, Songwu Lu, Daniel Massey, Andreas Terzis, and Lixia Zhang. Impact of Configuration Errors on DNS Robustness. In *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SICGOMM)*, pages 319–330, Portland, OR, 2004.

[332] Vasileios Pappas, Duane Wessels, Daniel Massey, Songwu Lu, Andreas Terzis, and Lixia Zhang. Impact of Configuration Errors on DNS Robustness. *IEEE Journal on Selected Areas in Communications*, 27(3):275–290, 2009.

[333] Philip Hunter. Verisign attack highlights where the real risks and fears lie. *Computer Fraud & Security*, 2007(3):16–17, 2007.

[334] George Lawton. Stronger Domain Name System Thwarts Root-Server Attacks. *IEEE Computer*, 40(5):14–17, 2007.

[335] Collin Jackson, Adam Barth, Andrew Bortz, Weidong Shao, and Dan Boneh. Protecting Browsers from DNS Rebinding Attacks. *ACM Transactions on the Web*, 3(1):2:1–2:26, 2009.

[336] D. Atkins and R. Austein. Threat Analysis of the Domain Name System (DNS). RFC 3833 (Informational), August 2004.

[337] Suranjith Ariyapperuma and Chris J. Mitchell. Security vulnerabilities in DNS and DNSSEC. In *Proceedings of the 2nd International Conference on Availability, Reliability and Security (ARES)*, pages 335–342, Vienna, April 2007.

[338] Nikolaos Chatzis. Motivation for Behaviour-Based DNS Security: A Taxonomy of DNS-Related Internet Threats. In *Proceedings of the International Conference on Emerging Security Information, Systems, and Technologies (SecureWare)*, pages 36–41, Valencia, October 2007.

[339] Paul Vixie, Gerry Sneeringer, and Mark Schleifer. 21 Oct 2002 Root Server Denial of Service Attack - Report. `http://www.isc.org/f-root-denial-of-service-21-oct-2002`, November 2002.

[340] CAIDA. Nameserver DoS Attack October 2002 . `http://www.caida.org/projects/dns/dns-root-gtld/oct02dos.xml`, May 2008.

[341] ICANN Security and Stability Advisory Committee (SSAC). DNS Distributed Denial of Service (DDoS) Attacks. SSAC Advisory (SAC008), ICANN, March 2006.

[342] Vern Paxson. An analysis of using reflectors for distributed denial-of-service attacks. *ACM Computer Communication Review*, 31(3):38–47, 2001.

[343] J. Damas and F. Neves. Preventing Use of Recursive Nameservers in Reflector Attacks. RFC 5358 (Best Current Practice), October 2008.

[344] Ken Silva, Frank Scalzo, and Piet Barber. Anatomy of Recent DNS Reflector Attacks from the Victim and Reflector Point of View. White paper, VeriSign, April 2006.

[345] Root server attack on 6 February 2007. Factsheet, ICANN, March 2007.

[346] Danny McPherson. February 2007 Root Server Attacks – A Qualitative Report. `http://asert.arbornetworks.com/2007/06/february-2007-root-server-attacks-a-qualitative-report/`, June 2007.

[347] C. Partridge, T. Mendez, and W. Milliken. Host Anycasting Service. RFC 1546 (Informational), November 1993.

[348] J. Abley and K. Lindqvist. Operation of Anycast Services. RFC 4786 (Best Current Practice), December 2006.

[349] Earl Zmijewski. Accidentally Importing Censorship. `http://www.renesys.com/blog/2010/03/fouling-the-global-nest.shtml`, March 2010.

[350] Earl Zmijewski. Two Strikes For the I-root. `http://www.renesys.com/blog/2010/06/two-strikes-i-root.shtml`, June 2010.

[351] Martin A. Brown, Doug Madory, Alin Popescu, and Earl Zmijewski. DNS Tampering and Root Servers. Presentation, Renesys Corporation, November 2010.

[352] i.root-servers.net. `http://www.netnod.se/dns/iroot`, July 2011.

[353] Earl Zmijewski. DNS: When Governments Lie (1). http://www.renesys.com/blog/2010/11/dns-when-governments-lie-1.shtml, November 2010.

[354] Earl Zmijewski. DNS: When Governments Lie (2). http://www.renesys.com/blog/2010/12/dns-when-governments-lie-2.shtml, December 2010.

[355] The OpenNet Initiative. http://opennet.net/.

[356] Kim Davies. Advisory – "L Root" changing IP address on 1st November. http://blog.icann.org/2007/10/advisory-%E2%80%94-l-root-changing-ip-address-on-1st-november/, October 2007.

[357] Earl Zmijewski. Identity Theft Hits the Root Name Servers. http://www.renesys.com/blog/2008/05/identity_theft_hits_the_root_n_1.shtml, May 2008.

[358] Danny McPherson. Uprooting of the DNS Root. http://asert.arbornetworks.com/2008/05/uprooting-of-the-dns-root/, May 2008.

[359] David Conrad. Ghosts of Root Servers Past. http://blog.icann.org/2008/05/ghosts-of-root-servers-past/, May 2008.

[360] Martin A. Brown, Alin Popescu, and Earl Zmijewski. Who's Manning the L root? Presentation, Renesys Corporation, Brooklyn, NY, June 2008.

[361] Earl Zmijewski. Securing the Root. http://www.renesys.com/blog/2008/06/securing-the-root-1.shtml, June 2008.

[362] Earl Zmijewski. Root Reprise: More Questions than Answers. http://www.renesys.com/blog/2008/06/root-reprise-more-questions-th-1.shtml, June 2008.

[363] ICANN Security and Stability Advisory Committee (SSAC). Domain Name Hijacking: Incidents, Threats, Risks, and Remedial Actions. Technical Report (SAC007), ICANN, July 2005.

[364] Steven Wright. Cybersquatting at the Intersection of Internet Domain Names and Trademark Law. *IEEE Communications Surveys & Tutorials*, 14(1):193–205, 2012.

[365] David Dagon, Niels Provos, Christopher P. Lee, and Wenke Lee. Corrupted DNS Resolution Paths: The Rise of a Malicious Resolution Authority. In *Proceedings of the 16th Annual Network & Distributed System Security Symposium (NDSS)*, San Diego, CA, February 2008.

[366] ICANN Security and Stability Advisory Committee (SSAC). Measures to Protect Domain Registration Services Against Exploitation or Misuse. Technical Report (SAC040), ICANN, August 2009.

[367] Danny McPherson. Your DNS is an Asset (Twitter DNS Woes...). `http://asert.arbornetworks.com/2009/12/your-dns-is-an-asset-twitter-dns-woes/`, December 2009.

[368] Robin Gandhi, Anup Sharma, William Mahoney, William Sousan, Qiuming Zhu, and Phillip Laplante. Dimensions of Cyber-Attacks: Cultural, Social, Economic, and Political. *IEEE Technology and Society Magazine*, 30(1):28–38, 2011.

[369] Craig Partridge, Paul Barford, David D. Clark, Sean Donelan, Vern Paxson, Jennifer Rexford, and Mary K. Vernon. *The Internet Under Crisis Conditions: Learning from September 11.* The National Academy Press, Washington, D.C., 2003.

[370] Andy Ogielski and Jim Cowie. Internet Routing Behavior on 9/11. Presentation, Renesys Corporation, March 2002.

[371] Declan McCullagh. Egypt's Internet disconnect reaches 24 hours. http://news.cnet.com/8301-31921_3-20029973-281.html, January 2011.

[372] Craig Labovitz. Middle East Internet Scorecard (February 12 – 20). http://asert.arbornetworks.com/2011/02/middle-east-internet-scorecard-february-12-%E2%80%93-20/, February 2011.

[373] Yuval Shavitt and Noa Zilberman. Arabian Nights: Measuring the Arab Internet During the 2011 Events. *IEEE Network*, 26(6):75–80, 2012.

[374] James Cowie. Egypt Leaves the Internet. http://www.renesys.com/blog/2011/01/egypt-leaves-the-internet.shtml, January 2011.

[375] Andree Toonk. Internet in Egypt offline. http://bgpmon.net/blog/?p=450, January 2011.

[376] Craig Labovitz. Egypt Loses the Internet. http://asert.arbornetworks.com/2011/01/egypt-loses-the-internet/, January 2011.

[377] Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly C. Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapé. Analysis of Country-wide Internet Outages Caused by Censorship. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 1–18, Berlin, November 2011.

[378] Earl Zmijewski. Egypt's Net on Life Support. http://www.renesys.com/blog/2011/01/egypts-net-on-life-support.shtml, January 2011.

[379] James Cowie. Egypt Returns To The Internet. http://www.renesys.com/blog/2011/02/egypt-returns-to-the-internet.shtml, February 2011.

[380] Andree Toonk. Egypt Back Online. http://bgpmon.net/blog/?p=480, February 2011.

[381] Craig Labovitz. Egypt Returns to the Internet. http://asert.arbornetworks.com/2011/02/egypt-returns-to-the-internet/, February 2011.

[382] James Cowie. Libyan Disconnect. http://www.renesys.com/blog/2011/02/libyan-disconnect-1.shtml, February 2011.

[383] James Cowie. What Libya Learned from Egypt. http://www.renesys.com/blog/2011/03/what-libya-learned-from-egypt.shtml, March 2011.

[384] James Cowie. Libyan Internet Instability. http://www.renesys.com/blog/2011/08/libyan-internet-instability.shtml, August 2011.

[385] James Cowie. The Battle for Tripoli's Internet. http://www.renesys.com/blog/2011/08/the-battle-for-tripolis-intern.shtml, August 2011.

[386] James Cowie. Syrian Internet Shutdown. http://www.renesys.com/blog/2011/06/syrian-internet-shutdown.shtml, June 2011.

[387] Andree Toonk. Internet Syria offline. http://bgpmon.net/blog/?p=511, June 2011.

[388] James Cowie. Tracing the Syrian Blackout. http://www.renesys.com/blog/2011/06/tracing-the-syrian-blackout.shtml, June 2011.

[389] Doug Madory. PCCW Keeps Syria Connected. http://www.renesys.com/blog/2012/08/china-keeps-syria-connected.shtml, August 2012.

[390] James Cowie. Syrian Internet Is Off The Air. http://www.renesys.com/blog/2012/11/syria-off-the-air.shtml, November 2012.

[391] The Internet Society. The Internet Society on Syria's Internet Shutdown. http://www.internetsociety.org/news/internet-society-syria's-internet-shutdown, November 2012.

[392] Matthew Prince. How Syria Turned Off the Internet. http://blog.cloudflare.com/how-syria-turned-off-the-internet, November 2012.

[393] James Cowie. Restoration in Syria. http://www.renesys.com/blog/2012/12/restoration-in-syria-1.shtml, December 2012.

[394] Geoff Huston. Interconnection, Peering, and Settlements. In *Proceedings of the 9th Annual Conference of the Internet Society (INET)*, San Jose, CA, June 1999.

[395] William B. Norton. The Art of Peering: The Peering Playbook. http://drpeering.net/white-papers/Art-Of-Peering-The-Peering-Playbook.html, August 2010.

[396] Todd Underwood. Peering—The Fundamental Architecture of the Internet. http://www.renesys.com/blog/2005/12/peering-the-fundamental-archit.shtml, December 2005.

[397] Alin Popescu and Todd Underwood. D(3)peered: Just the Facts Ma'am – A technical review of Level (3)'s Depeering of Cogent. Presentation, Renesys Corp, Los Angeles, CA, October 2005.

[398] Earl Zmijewski. You can't get there from here. http://www.renesys.com/blog/2008/03/you-cant-get-there-from-here-1.shtml, March 2008.

[399] Earl Zmijewski. He said, she said: Cogent vs. Telia. http://www.renesys.com/blog/2008/03/he-said-she-said-cogent-vs-tel.shtml, March 2008.

[400] Earl Zmijewski. Telia and Cogent Kiss and Make Up. http://www.renesys.com/blog/2008/03/telia-and-cogent-kiss-and-make-1.shtml, March 2008.

[401] Martin A. Brown, Alin Popescu, and Earl Zmijewski. Peering Wars – Lessons Learned from the Cogent-Telia De-peering. Presentation, Renesys Corp, New York, NY, June 2008.

[402] Todd Underwood. Wrestling With the Zombie: Sprint Depeers Cogent, Internet Partitioned. http://www.renesys.com/blog/2008/10/wrestling-with-the-zombie-spri.shtml, October 2008.

[403] Todd Underwood. Sprint and Cogent Repeer–For Now. http://www.renesys.com/blog/2008/11/sprint-and-cogent-repeerfor-no.shtml, November 2008.

[404] Jon Crowcroft. Net Neutrality: The Technical Side of the Debate: A White Paper. *ACM Computer Communication Review*, 37(1):49–56, 2007.

[405] Peter Cochrane. Net Neutrality or Suicide? *Proceedings of the IEEE*, 94(10):1779–1780, 2006.

[406] Scott Jordan. Implications of Internet Architecture on Net Neutrality. *ACM Transactions on Internet Technology*, 9(2):5:1–5:28, 2009.

[407] Aaron Weiss. Net Neutrality?: There's Nothing Neutral About It. *ACM netWorker Magazine*, 10(2):18–25, 2006.

[408] Greg Goth. ISP Traffic Management: Will Innovation or Regulation Ensure Fairness? *IEEE Distributed Systems Online*, 9(9), September 2008. art. no. 0809-o9002.

[409] Daniel J. Weitzner. Net Neutrality... Seriously this Time. *IEEE Internet Computing*, 12(3):86–89, 2008.

[410] Marguerite Reardon. What Verizon's FCC tethering settlement means to you (FAQ). http://news.cnet.com/8301-1035_3-57485518-94/what-verizons-fcc-tethering-settlement-means-to-you-faq, August 2012.

[411] James P.G. Sterbenz and Joseph D. Touch. *High-Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication*. Wiley, 1st edition, May 2001.

[412] Steven Cherry. The net effect: as China's Internet gets a much-needed makeover, will the new network promote freedom or curtail it? *IEEE Spectrum Magazine*, 42(6):38–44, 2005.

[413] George Danezis and Ross Anderson. The Economics of Resisting Censorship. *IEEE Security & Privacy*, 3(1):45–50, 2005.

[414] Jonathan Zittrain and Benjamin Edelman. Internet Filtering in China. *IEEE Internet Computing*, 7(2):70–77, 2003.

[415] Jedidiah R. Crandall, Daniel Zinn, Michael Byrd, Earl Barr, and Rich East. ConceptDoppler: A Weather Tracker for Internet Censorship. In *Proceedings of the 14th ACM Conference on Computer and Communications Security (CCS)*, pages 352–365, Alexandria, VA, October 2007.

[416] Heejin Lee and Jaeho Hwang. ICT Development in North Korea: Changes and Challenges. *Inf. Technol. Int. Dev.*, 2(1):75–88, 2004.

[417] Paul Tjia. Inside the hermit kingdom: IT and outsourcing in North Korea. *Communications of the ACM*, 55(8):22–25, 2012.

[418] David M. Kristol. HTTP Cookies: Standards, Privacy, and Politics. *ACM Transactions on Internet Technology*, 1(2):151–198, 2001.

[419] Joon S. Park and Ravi Sandhu. Secure Cookies on the Web. *IEEE Internet Computing*, 4(4):36–44, 2000.

[420] Huaiqing Wang, Matthew K. O. Lee, and Chen Wang. Consumer Privacy Concerns about Internet Marketing. *Communications of the ACM*, 41(3):63–70, 1998.

[421] Susan Landau. Security, Wiretapping, and the Internet. *IEEE Security & Privacy*, 3(6):26–33, 2005.

[422] Steven M. Bellovin, Matt Blaze, Whitfield Diffie, Susan Landau, Peter G. Neumann, and Jennifer Rexford. Risking Communications Security: Potential Hazards of the Protect America Act. *IEEE Security & Privacy*, 6(1):24–33, 2008.

[423] Jim W. Roberts. Traffic Theory and the Internet. *IEEE Communications Magazine*, 39(1):94–99, 2001.

[424] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja, and X. Xiao. Overview and Principles of Internet Traffic Engineering. RFC 3272 (Informational), May 2002. Updated by RFC 5462.

[425] Paul Barford, Jeffery Kline, David Plonka, and Amos Ron. A Signal Analysis of Network Traffic Anomalies. In *Proceedings of the 2nd ACM Workshop on Internet Measurment (IMW)*, pages 71–82, Marseille, November 2002.

[426] Jaeyeon Jung, Balachander Krishnamurthy, and Michael Rabinovich. Flash Crowds and Denial of Service Attacks: Characterization and Implications for CDNs and Web Sites. In *Proceedings of the 11th ACM International Conference on World Wide Web (WWW)*, pages 293–304, Honolulu, HI, May 2002.

[427] William LeFebvre. CNN.com: Facing a World Crisis. In *Proceedings of the 15th USENIX Conference on Systems Administration (LISA)*, San Diego, CA, December 2001. Invited Talk.

[428] Report of the 7 July Review Committee. Report, Greater London Authority, June 2006.

[429] US-CERT Vulnerability Notes Database. `http://www.kb.cert.org/vuls/`.

[430] NIST National Vulnerability Database. `http://nvd.nist.gov/`.

[431] Michael Lesk. The New Front Line: Estonia under Cyberassault. *IEEE Security & Privacy*, 5(4):76–79, 2007.

[432] Thomas M. Chen and Saeed Abu-Nimeh. Lessons from Stuxnet. *IEEE Computer*, 44(4):91–93, 2011.

[433] Ralph Langner. Stuxnet: Dissecting a Cyberwarfare Weapon. *IEEE Security & Privacy*, 9(3):49–51, 2011.

[434] John Laprise. Cyber-Warfare Seen through a Mariner's Spyglass. *IEEE Technology and Society Magazine*, 25(3):26–33, 2006.

[435] Scott D. Applegate. Cyber Militias and Political Hackers: Use of Irregular Forces in Cyber Warfare. *IEEE Security & Privacy*, 9(5):16–22, 2011.

[436] Richard Power and Dario Forte. Ten years in the wilderness – a retrospective Part 2: Cyber Security = National Security. *Computer Fraud & Security*, 2006(2):16–20, 2006.

[437] Maura Conway. Hackers as terrorists? why it doesn't compute. *Computer Fraud & Security*, 2003(12):10–13, 2003.

[438] Steve Mansfield-Devine. Hacktivism: assessing the damage. *Network Security*, 2011(8):5–13, 2011.

[439] Keiran Hardy. WWWMDs: Cyber-attacks against infrastructure in domestic anti-terror laws. *Computer Law & Security Review*, 27(2):152–161, 2011.

[440] Craig Labovitz. Iranian Traffic Engineering. `http://asert.arbornetworks.com/2009/06/iranian-traffic-engineering/`, June 2009.

[441] James Cowie. Strange Changes in Iranian Transit. `http://www.renesys.com/blog/2009/06/strange-changes-in-iranian-int.shtml`, June 2009.

[442] Craig Labovitz. A Deeper Look at The Iranian Firewall. `http://asert.arbornetworks.com/2009/06/a-deeper-look-at-the-iranian-firewall/`, June 2009.

[443] Craig Labovitz. Return to the Iranian Firewall. `http://asert.arbornetworks.com/2009/08/return-to-the-iranian-firewall/`, August 2009.

[444] Craig Labovitz. Behind the Firewall – A Look at Six Iranian ISPs Forty Days Later. `http://asert.arbornetworks.com/2009/08/1132/`, August 2009.

[445] James Cowie. The Proxy Fight for Iranian Democracy. `http://www.renesys.com/blog/2009/06/the-proxy-fight-for-iranian-de.shtml`, June 2009.

[446] David D. Clark, John Wroclawski, Karen R. Sollins, and Robert Braden. Tussle in Cyberspace: Defining Tomorrow's Internet. *IEEE/ACM Transactions on Networking*, 13(3):462–475, 2005.

[447] Leonard Kleinrock. Nomadicity: Anytime, anywhere in a disconnected world. *Mobile Networks and Applications*, 1(4):351–357, 1996.

[448] George H. Forman and John Zahorjan. The Challenges of Mobile Computing. *IEEE Computer*, 27(4):38–47, 1994.

[449] David Tipper, Teresa Dahlberg, Hyundoo Shin, and Charlermpol Charnsripinyo. Providing Fault Tolerance in Wireless Access Networks. *IEEE Communications Magazine*, 40(1):58–64, 2002.

[450] Andrew P. Snow, Upkar Varshney, and Alisha D. Malloy. Reliability and Survivability of Wireless and Mobile Networks. *IEEE Computer*, 33(7):49–55, 2000.

[451] Zhensheng Zhang. Routing in Intermittently Connected Mobile Ad Hoc Networks and Delay Tolerant Networks: Overview and Challenges. *IEEE Communications Surveys & Tutorials*, 8(1):24–37, 2006.

[452] Luciana Pelusi, Andrea Passarella, and Marco Conti. Opportunistic Networking: Data Forwarding in Disconnected Mobile Ad Hoc Networks. *IEEE Communications Magazine*, 44(11):134–141, 2006.

[453] Haitao Liu, Baoxian Zhang, Hussein T. Mouftah, Xiaojun Shen, and Jian Ma. Opportunistic Routing for Wireless Ad Hoc and Sensor Networks: Present and Future Directions. *IEEE Communications Magazine*, 47(12):103–109, 2009.

[454] Maurice J. Khabbaz, Chadi M. Assi, and Wissam F. Fawaz. Disruption-Tolerant Networking: A Comprehensive Survey on Recent Developments and Persisting Challenges. *IEEE Communications Surveys & Tutorials*, 14(2):607–640, 2012.

[455] Ricardo Sánchez, Joseph Evans, and Gary Minden. Networking on the Battlefield: Challenges in Highly Dynamic Multi-hop Wireless Networks. In *Proceedings of the IEEE Military Communications Conference (MILCOM)*, volume 2, pages 751–755, Atlantic City, NJ, November 1999.

[456] Justin P. Rohrer, Abdul Jabbar, Erik Perrins, and James P.G. Sterbenz. Cross-Layer Architectural Framework for Highly-Mobile Multihop Airborne Telemetry Networks. In *Proceedings of the IEEE Military Communications Conference (MILCOM)*, pages 1–9, San Diego, CA, November 2008.

[457] Justin P. Rohrer, Abdul Jabbar, Egemen K. Çetinkaya, Erik Perrins, and James P.G. Sterbenz. Highly-Dynamic Cross-Layered Aeronautical Network Architecture. *IEEE Transactions on Aerospace and Electronic Systems*, 47(4):2742–2765, 2011.

[458] Ian F. Akyildiz, Weilian Su, Yogesh Sankarasubramaniam, and Erdal Cayirci. Wireless sensor networks: a survey. *Computer Networks*, 38(4):393–422, 2002.

[459] Chris Karlof and David Wagner. Secure routing in wireless sensor networks: attacks and countermeasures. *Ad Hoc Networks*, 1(2-3):293–315, 2003.

[460] Wenyuan Xu, Wade Trappe, Yanyong Zhang, and Timothy Wood. The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks. In *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 46–57, Urbana-Champaign, IL, May 2005.

[461] Konstantinos Pelechrinis, Marios Iliofotou, and Srikanth V. Krishnamurthy. Denial of Service Attacks in Wireless Networks: The Case of Jammers. *IEEE Communications Surveys & Tutorials*, 13(2):245–257, 2011.

[462] Aristides Mpitziopoulos, Damianos Gavalas, Charalampos Konstantopoulos, and Grammati Pantziou. A Survey on Jamming Attacks and Countermeasures in WSNs. *IEEE Communications Surveys & Tutorials*, 11(4):42–56, 2009.

[463] Wenyuan Xu, Ke Ma, Wade Trappe, and Yanyong Zhang. Jamming Sensor Networks: Attack and Defense Strategies. *IEEE Network*, 20(3):41–47, 2006.

[464] Ahmad Al Hanbali, Eitan Altman, and Philippe Nain. A Survey of TCP over Ad Hoc Networks. *IEEE Communications Surveys & Tutorials*, 7(3):22–36, 2005.

[465] Ka-Cheong Leung and Victor O.K. Li. Transmission Control Protocol (TCP) in Wireless Networks: Issues, Approaches, and Challenges. *IEEE Communications Surveys & Tutorials*, 8(4):64–79, 2006.

[466] Nasir Ghani and Sudhir Dixit. TCP/IP Enhancements for Satellite Networks. *IEEE Communications Magazine*, 37(7):64–72, 1999.

[467] Rajesh Krishnan, James P.G. Sterbenz, Wesley M. Eddy, Craig Partridge, and Mark Allman. Explicit Transport Error Notification (ETEN) for Error-Prone Wireless and Satellite Networks. *Computer Networks*, 46(3):343–362, 2004.

[468] Ian F. Akyildiz, Xudong Wang, and Weilin Wang. Wireless mesh networks: a survey. *Computer Networks*, 47(4):445–487, 2005.

[469] Sally Floyd. TCP and Explicit Congestion Notification. *ACM SIGCOMM Computer Communication Review*, 24(5):8–23, 1994.

[470] K. Ramakrishnan, S. Floyd, and D. Black. The Addition of Explicit Congestion Notification (ECN) to IP. RFC 3168 (Proposed Standard), September 2001.

[471] Hari Balakrishnan, Venkata N. Padmanabhan, Srinivasan Seshan, and Randy H. Katz. A Comparison of Mechanisms for Improving TCP Performance over Wireless Links. *IEEE/ACM Transactions on Networking (TON)*, 5(6):756–769, 1997.

[472] Saad Biaz and Nitin H. Vaidya. Distinguishing Congestion Losses from Wireless Transmission Losses: A Negative Result. In *Proceedings of the 7th IEEE International Conference on Computer Communications and Networks (ICCCN)*, pages 722–731, Lafayette, LA, October 1998.

[473] Ye Tian, Kai Xu, and N. Ansari. TCP in Wireless Environments: Problems and Solutions. *IEEE Communications Magazine*, 43(3):S27–S32, 2005.

[474] Ruhai Wang, Tarik Taleb, Abbas Jamalipour, and Bo Sun. Protocols for Reliable Data Transport in Space Internet. *IEEE Communications Surveys & Tutorials*, 11(2):21–32, 2009.

[475] Kevin Fall. A Delay-Tolerant Network Architecture for Challenged Internets. In *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 27–34, Karlsruhe, August 2003.

[476] Kevin Fall and Stephen Farrell. DTN: An Architectural Retrospective. *IEEE Journal on Selected Areas in Communications*, 26(5):828–836, 2008.

[477] Scott Burleigh, Adrian Hooke, Leigh Torgerson, Kevin Fall, Vint Cerf, Bob Durst, Keith Scott, and Howard Weiss. Delay-Tolerant Networking: An Approach to Interplanetary Internet. *IEEE Communications Magazine*, 41(6):128–136, 2003.

[478] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss. Delay-Tolerant Networking Architecture. RFC 4838 (Informational), April 2007.

[479] Stephen Farrell, Vinny Cahill, Dermot Geraghty, Ivor Humphreys, and Paul McDonald. When TCP Breaks: Delay- and Disruption- Tolerant Networking. *IEEE Internet Computing*, 10(4):72–78, 2006.

[480] Robert C. Durst, Gregory J. Miller, and Eric J. Travis. TCP extensions for space communications. In *Proceedings of the ACM 2nd Annual International Conference on Mobile Computing and Networking (MobiCom)*, pages 15–26, Rye, NY, November 1996.

[481] Robert C. Durst, Gregory J. Miller, and Eric J. Travis. TCP extensions for space communications. *Wireless Networks*, 3(5):389–403, 1997.

[482] K. Scott and S. Burleigh. Bundle Protocol Specification. RFC 5050 (Experimental), November 2007.

[483] S. Burleigh, M. Ramadas, and S. Farrell. Licklider Transmission Protocol - Motivation. RFC 5325 (Informational), September 2008.

[484] M. Ramadas, S. Burleigh, and S. Farrell. Licklider Transmission Protocol - Specification. RFC 5326 (Experimental), September 2008.

[485] Rahul C. Shah, Sumit Roy, Sushant Jain, and Waylon Brunette. Data MULEs: modeling and analysis of a three-tier architecture for sparse sensor networks. *Ad Hoc Networks*, 1(2-3):215–233, 2003.

[486] Anthony Ephremides. Energy Concerns in Wireless Networks. *IEEE Wireless Communications*, 9(4):48–59, 2002.

[487] Andrea J. Goldsmith and Stephen B. Wicker. Design Challenges for Energy-Constrained Ad Hoc Wireless Networks. *IEEE Wireless Communications*, 9(4):8–27, 2002.

[488] Richard Cocchiara, Hugh Davis, and Doug Kinnaird. Data center topologies for mission-critical business systems. *IBM Systems Journal*, 47(4):695–706, 2008.

[489] George Loukas, Diane Gan, and Tuan Vuong. A taxonomy of cyber attack and defence mechanisms for emergency management networks. In *Proceedings of the 3rd IEEE International Workshop on Pervasive Networks for Emergency Management (PerNEM)*, San Diego, CA, March 2013.

[490] Doug Madory. Syria Briefly Disconnects. http://www.renesys.com/blog/2012/07/syria-leaves-the-internet.shtml, July 2012.

[491] https://twitter.com/twittercomms/status/30377205695647744, January 2011.

[492] Thomas M. Chen. Governments and the Executive 'Internet Kill Switch'. *IEEE Network*, 25(2):2–3, 2011.

[493] Vinay M. Igure, Sean A. Laughter, and Ronald D. Williams. Security issues in SCADA networks. *Computers & Security*, 25(7):498–506, 2006.

[494] E.Eugene Schultz. A framework for understanding and predicting insider attacks. *Computers & Security*, 21(6):526–531, 2002.

[495] Jeffrey M. Stanton, Kathryn R. Stam, Paul Mastrangelo, and Jeffrey Jolton. Analysis of end user security behaviors. *Computers & Security*, 24(2):124–133, 2005.

[496] Matthew Roughan, Walter Willinger, Olaf Maennel, Debbie Perouli, and Randy Bush. 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems. *IEEE Journal on Selected Areas in Communications*, 29(9):1810–1821, 2011.

[497] David Alderson, Lun Li, Walter Willinger, and John C. Doyle. Understanding Internet Topology: Principles, Models, and Validation. *IEEE/ACM Transactions on Networking*, 13(6):1205–1218, 2005.

[498] John C. Doyle, David L. Alderson, Lun Li, Steven Low, Matthew Roughan, Stanislav Shalunov, Reiko Tanaka, and Walter Willinger. The 'robust yet fragile' nature of the Internet. *Proceedings of the National Academy of Sciences of the United States of America*, 102(41):14497–14502, 2005.

[499] Egemen K. Çetinkaya, Andrew M. Peck, and James P.G. Sterbenz. Flow Robustness of Multilevel Networks. In *Proceedings of the 9th IEEE/IFIP International Conference on the Design of Reliable Communication Networks (DRCN)*, pages 274–281, Budapest, March 2013.

[500] Egemen K. Çetinkaya, Mohammed J.F. Alenazi, Yufei Cheng, Andrew M. Peck, and James P.G. Sterbenz. On the Fitness of Geographic Graph Generators for Modelling Physical Level Topologies. In *Proceedings of the 5th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, Almaty, September 2013.

[501] American Association of State Highway and Transportation Officials. Guidelines for the Selection of Supplemental Guide Signs for Traffic Generators Adjacent to Freeways. Washington, DC, 2001.

[502] M. T. Gastner and M. E.J. Newman. The spatial structure of networks. *The European Physical Journal B - Condensed Matter and Complex Systems*, 49(2):247–252, 2006.

[503] Rocketfuel: An ISP topology mapping engine. `http://www.cs.washington.edu/research/networking/rocketfuel/interactive`, 2002.

[504] Level 3 network map. `http://maps.level3.com`.

[505] KMI Corporation. North American Fiberoptic Long-haul Routes Planned and in Place, 1999.

[506] TeliaSonera. `http://www.teliasoneraic.com`.

[507] Internet2. `http://www.internet2.edu`.

[508] The Next Generation Core Optical Networks (CORONET). http://www.darpa.mil/Our_Work/STO/Programs/Dynamic_Multi-Terabit_Core_Optical_Networks_(CORONET).aspx.

[509] George Clapp, Ronald A. Skoog, Ann C. Von Lehmen, and Brian Wilson. Management of Switched Systems at 100 Tbps: the DARPA CORONET Program. In *International Conference on Photonics in Switching (PS)*, pages 1–4, Pisa, September 2009.

[510] John Strand, Angela L. Chiu, and Robert Tkach. Issues for routing in the optical layer. *IEEE Communications Magazine*, 39(2):81–87, 2001.

[511] Ellen W. Zegura, Kenneth L. Calvert, and Michael J. Donahoo. A Quantitative Comparison of Graph-Based Models for Internet Topology. *IEEE/ACM Transactions on Networking*, 5(6):770–783, 1997.

[512] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring Network Structure, Dynamics, and Function using NetworkX. In *7th Python in Science Conference (SciPy)*, pages 11–15, Pasadena, CA, August 2008.

[513] Linton C. Freeman. A Set of Measures of Centrality Based on Betweenness. *Sociometry*, 40(1):35–41, 1977.

[514] Priya Mahadevan, Dmitri Krioukov, Marina Fomenkov, Xenofontas Dimitropoulos, K. C. Claffy, and Amin Vahdat. The Internet AS-Level Topology: Three Data Sources and One Definitive Metric. *ACM Computer Communication Review*, 36(1):17–26, 2006.

[515] Linton C. Freeman. Centrality in social networks conceptual clarification. *Social Networks*, 1(3):215–239, 1978–1979.

[516] M. E. J. Newman. Assortative mixing in networks. *Phys. Rev. Lett.*, 89(20):208701, October 2002.

[517] A. Jamaković and S. Uhlig. On the relationship between the algebraic connectivity and graph's robustness to node and link failures. In *Proceedings of the 3rd EuroNGI Conference on Next Generation Internet Networks*, pages 96–102, Trondheim, May 2007.

[518] A. Jamaković and P. Van Mieghem. On the Robustness of Complex Networks by Using the Algebraic Connectivity. In *Proceedings of the 7th International IFIP Networking Conference*, volume 4982 of *Lecture Notes in Computer Science*, pages 183–194. Singapore, May 2008.

[519] Fan R. K. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997.

[520] P. Van Mieghem. *Graph Spectra for Complex Networks*. Cambridge University Press, 2011.

[521] Norman Biggs. *Algebraic Graph Theory*. Cambridge University Press, 2nd edition, 1993.

[522] Andries E. Brouwer and Willem H. Haemers. *Spectra of Graphs*. Springer New York, 2012.

[523] Dragoš Cvetković, Peter Rowlinson, and Slobodan Simić. *An Introduction to the Theory of Graph Spectra*. London Mathematical Society, 2009.

[524] Christos Gkantsidis, Milena Mihail, and Ellen Zegura. Spectral Analysis of Internet Topologies. In *Proceedings of the IEEE INFOCOM*, volume 1, pages 364–374, San Francisco, CA, April 2003.

[525] A. Jamaković and P. Van Mieghem. The Laplacian Spectrum of Complex Networks. In *Proceedings of the European Conference on Complex Systems (ECCS)*, Oxford, September 2006.

[526] Danica Vukadinović, Polly Huang, and Thomas Erlebach. On the Spectrum and Structure of Internet Topology Graphs. In *Proceedings of the Second International Workshop on Innovative Internet Computing Systems (IICS)*, volume 2346 of *Lecture Notes in Computer Science*, pages 83–95. Kühlungsborn, June 2002.

[527] Anirban Banerjee and Jürgen Jost. Spectral characterization of network structures and dynamics. In Niloy Ganguly, Andreas Deutsch, and Animesh Mukherjee, editors, *Dynamics On and Of Complex Networks*, Modeling and Simulation in Science, Engineering and Technology, pages 117–132. Birkhäuser Boston, 2009.

[528] R. J. Ellison, D. A. Fisher, R. C. Linger, H. F. Lipson, T. Longstaff, and N. R. Mead. Survivable network systems: An emerging discipline. Technical Report CMU/SEI-97-TR-013, Carnegie-Mellon Sorfware Engineering Institute, PA, 1999.

[529] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the Internet topology. In *Proceedings of the ACM SIGCOMM*, pages 251–262, Cambridge, MA, 1999.

[530] Tom Lehman, Xi Yang, Nasir Ghani, Feng Gu, Chin Guok, Inder Monga, and Brian Tierney. Multilayer Networks: An Architecture Framework. *IEEE Communications Magazine*, 49(5):122–130, 2011.

[531] Ulrik Brandes and Daniel Fleischer. Centrality measures based on current flow. In *Proceedings of the 22nd Annual Conference on Theoretical Aspects of Computer Science (STACS)*, volume 3404 of *LNCS*, pages 533–544. Springer Berlin / Heidelberg, Stuttgart, February 2005.

[532] AtlantaNAP. http://www.atlantanap.com.

[533] Equinix. http://www.equinix.com.

[534] Terremark. http://www.terremark.com.

[535] MAE-East. http://en.wikipedia.org/wiki/MAE-East.

[536] Mahmood A. Hameed, Abdul Jabbar, Egemen K. Çetinkaya, and James P.G. Sterbenz. Deriving Network Topologies from Real World Constraints. In *Proceedings of IEEE GLOBECOM Workshop on Complex and Communication Networks (CCNet)*, pages 400–404, Miami, FL, December 2010.

[537] James P.G. Sterbenz, Egemen K. Çetinkaya, Mahmood A. Hameed, Abdul Jabbar, and Justin P. Rohrer. Modelling and analysis of network resilience (invited paper). In *Proceedings of the Third IEEE International Conference on Communication Systems and Networks (COMSNETS)*, pages 1–10, Bangalore, January 2011.

[538] Mohammed J.F. Alenazi, Egemen K. Çetinkaya, and James P.G. Sterbenz. Network Design and Optimisation Based on Cost and Algebraic Connectivity. In *Proceedings of the 5th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, Almaty, September 2013.

[539] K. Ruben Gabriel and Robert R. Sokal. A New Statistical Approach to Geographic Variation Analysis. *Systematic Zoology*, 18(3):259–278, 1969.

[540] David W. Matula and Robert R. Sokal. Properties of Gabriel Graphs Relevant to Geographic Variation Research and the Clustering of Points in the Plane. *Geographical Analysis*, 12(3):205–222, 1980.

[541] Mathew Penrose. *Random Geometric Graphs*. Oxford Studies in Probability 5, 2003.

[542] N.V.R. Mahadev and U.N. Peled. *Threshold Graphs and Related Topics*, volume 56 of *Annals of Discrete Mathematics*. Elsevier North-Holland, Inc., 1995.

[543] Milan Bradonjić, Aric Hagberg, and Allon G. Percus. The Structure of Geographical Threshold Graphs. *Internet Mathematics*, 5(1-2):113–139, 2008.

[544] US Census Bureau Population Estimates. `http://www.census.gov/popest/data/cities/totals/2011/files/SUB-EST2011_ALL.csv`, 2013.

[545] Bernard M. Waxman. Routing of Multipoint Connections. *IEEE Journal on Selected Areas in Communications*, 6(9):1617–1622, 1988.

[546] James P.G. Sterbenz, Deep Medhi, Byrav Ramamurthy, Caterina Scoglio, Justin P. Rohrer, Egemen K. Çetinkaya, Ramkumar Cherukuri, Xuan Liu, Pragatheeswaran Angu, Andy Bavier, and Cort Buffington. The GpENI Testbed: Network Infrastructure, Implementation Experience, and Experimentation. *Computer Networks*, July 2013. accepted with minor modifications.

[547] The ns-3 Network Simulator. `http://www.nsnam.org/`.

[548] The ns-2 Network Simulator. `http://www.isi.edu/nsnam/ns/`.

[549] Egemen K. Çetinkaya, Abdul Jabbar, Dan Broyles, Amit Dandekar, Rabat Mahmood, and James P.G. Sterbenz. Challenge modelling. `https://wiki.ittc.ku.edu/resilinets/Challenge_Modelling`, August 2008.

[550] CGAL, Computational Geometry Algorithms Library. http://www.cgal.org.

[551] Abdul Jabbar, Qian Shi, Egemen Çetinkaya, and James P.G. Sterbenz. KU-LocGen: Location and Cost-Constrained Network Topology Generator. ITTC Technical Report ITTC-FY2009-TR-45030-01, The University of Kansas, Lawrence, KS, December 2008.

[552] Sprint Network Maps. https://www.sprint.net/network_maps.php.

[553] Michael E. Whitman. Enemy at the Gate: Threats to Information Security. *Communications of the ACM*, 46(8):91–95, 2003.

[554] Daniel Turner, Kirill Levchenko, Alex C. Snoeren, and Stefan Savage. California Fault Lines: Understanding the Causes and Impact of Network Failures. In *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 315–326, New Delhi, September 2010.

[555] tinc wiki. http://www.tinc-vpn.org/, 2010.

[556] Egemen K. Çetinkaya, Dongsheng Zhang, Mohammed J. F. Alenazi, Yufei Cheng, Parker Riley, and James P.G. Sterbenz. Resilience Experiments on the GpENI Testbed (poster). In *17th GENI Engineering Conference (GEC 17) Demo and Poster Session*, Madison, WI, July 2013.

[557] David Oppenheimer, Archana Ganapathi, and David A. Patterson. Why Do Internet Services Fail, and What Can Be Done About It? In *Proceedings of the 4th Conference on USENIX Symposium on Internet Technologies and Systems (USITS)*, pages 1–16, Seattle, WA, March 2003.

[558] Athina Markopoulou, Gianluca Iannaccone, Supratik Bhattacharyya, Chen-Nee Chuah, Yashar Ganjali, and Christophe Diot. Characterization of Failures in an Operational IP Backbone Network. *IEEE/ACM Transactions on Networking*, 16(4):749–762, 2008.

[559] US Federal Communications Commission (FCC) Network Outage Reporting System (NORS). http://www.fcc.gov/pshs/services/cip/nors/nors.html.

[560] Doug Madory. Blast in Turkey Impacts Iran, Iraq. http://www.renesys.com/blog/2012/10/blast-in-turkey-impacts-iran-i.shtml, October 2012.

[561] Kari Huus. Gas line explodes in West Virginia; homes burn, freeway damaged. http://usnews.nbcnews.com/_news/2012/12/11/15845530-gas-line-explodes-in-west-virginia-homes-burn-freeway-damaged?lite, December 2012.

Page left intentionally blank.

Appendix A

Multilevel Flow Robustness Plots

This appendix contains a full set of flow robustness plots for the topologies used in the analysis of the *multilevel networks*.

A.1 Node Deletions

A.1.1 AT&T

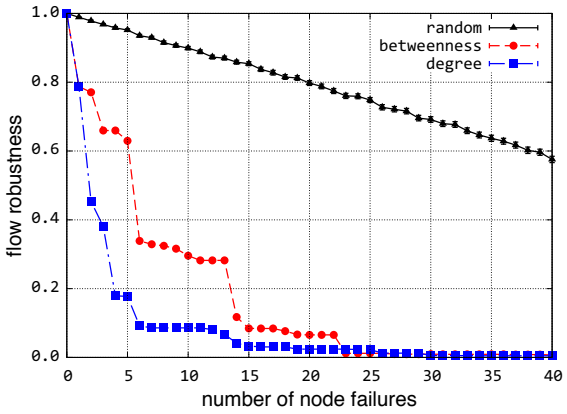

Figure A.1: Flow robustness for dynamic routing during adaptive node deletions

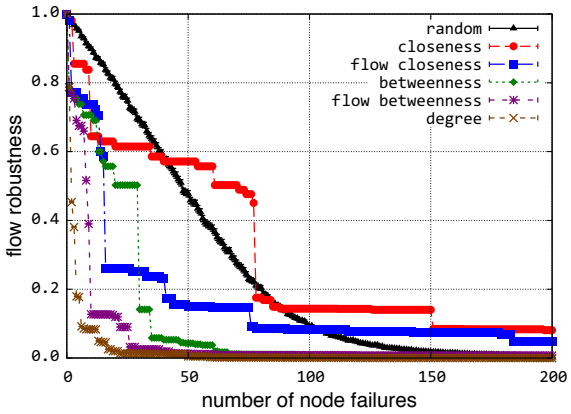

Figure A.2: Flow robustness for dynamic routing during non-adaptive node deletions

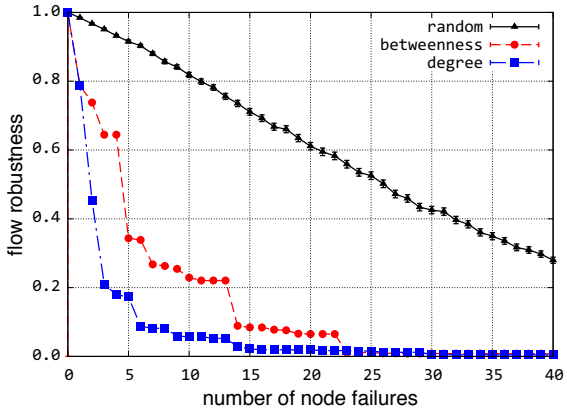

Figure A.3: Flow robustness for static routing during adaptive node deletions

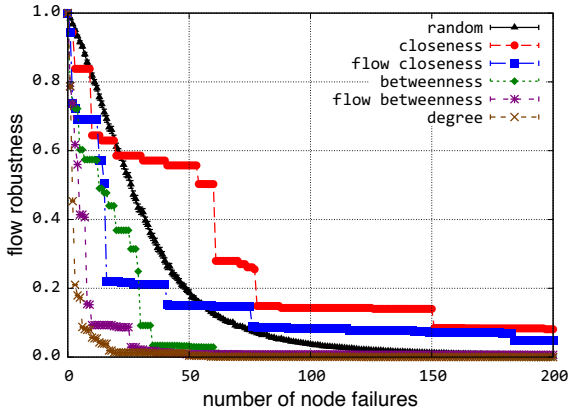

Figure A.4: Flow robustness for static routing during non-adaptive node deletions

A.1.2 Level 3

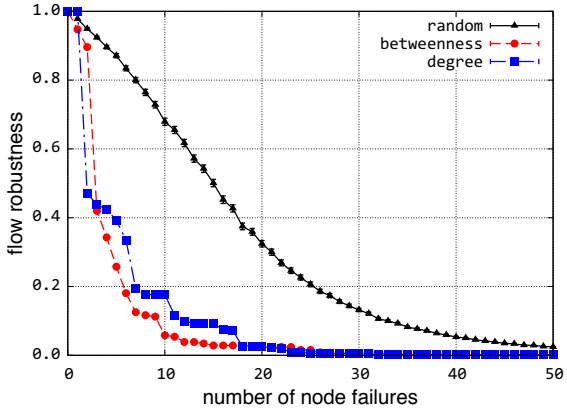

Figure A.5: Flow robustness for dynamic routing during adaptive node deletions

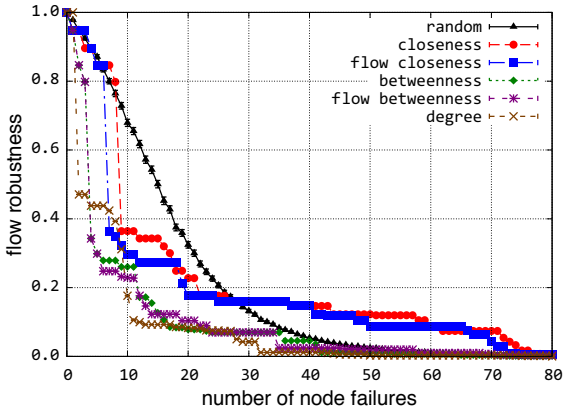

Figure A.6: Flow robustness for dynamic routing during non-adaptive node deletions

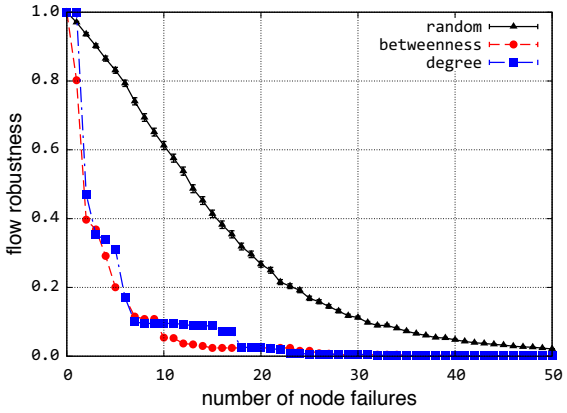

Figure A.7: Flow robustness for static routing during adaptive node deletions

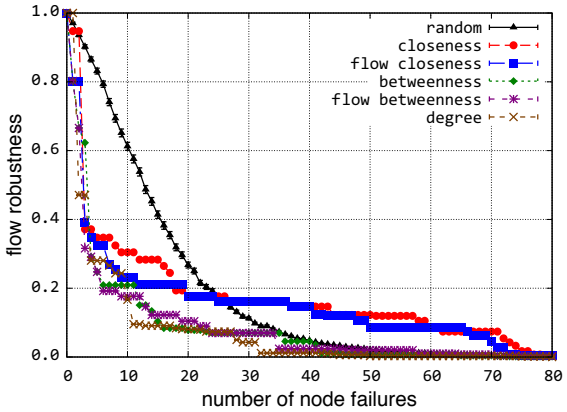

Figure A.8: Flow robustness for static routing during non-adaptive node deletions

A.1.3 Sprint

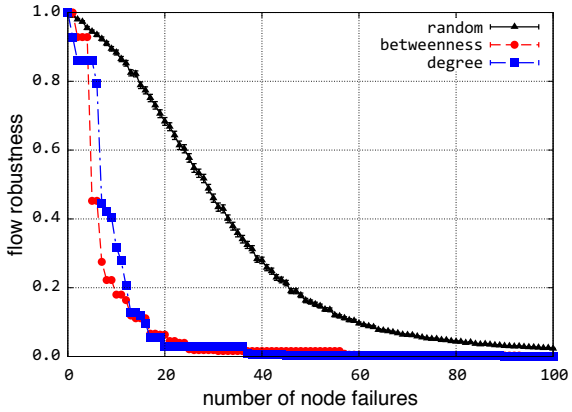

Figure A.9: Flow robustness for dynamic routing during adaptive node deletions

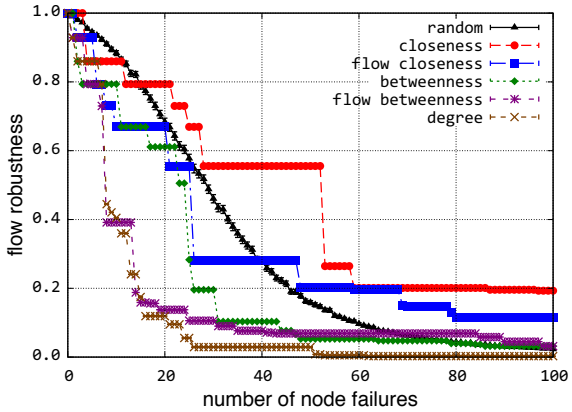

Figure A.10: Flow robustness for dynamic routing during non-adaptive node deletions

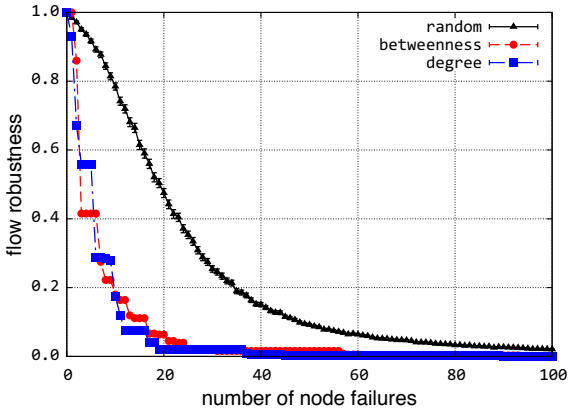

Figure A.11: Flow robustness for static routing during adaptive node deletions

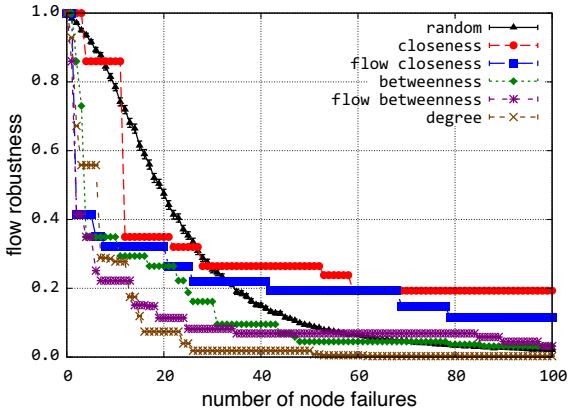

Figure A.12: Flow robustness for static routing during non-adaptive node deletions

A.1.4 TeliaSonera

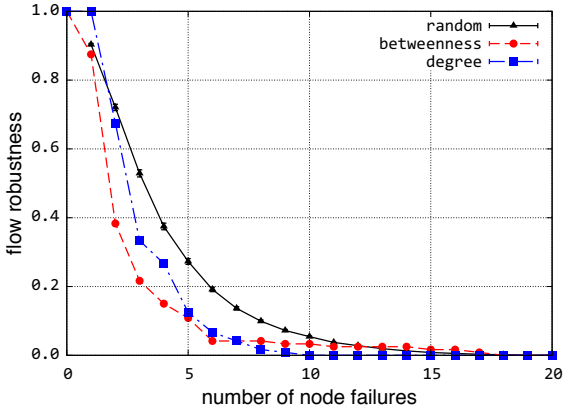

Figure A.13: Flow robustness for dynamic routing during adaptive node deletions

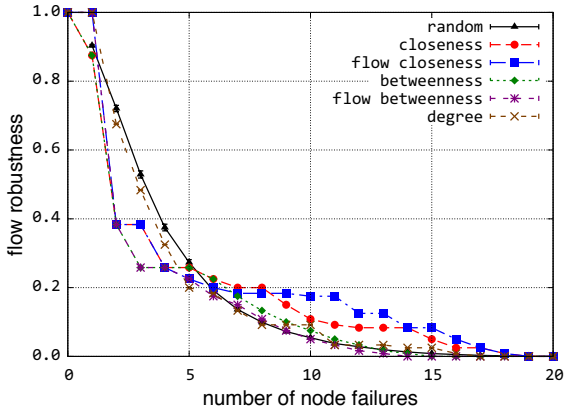

Figure A.14: Flow robustness for dynamic routing during non-adaptive node deletions

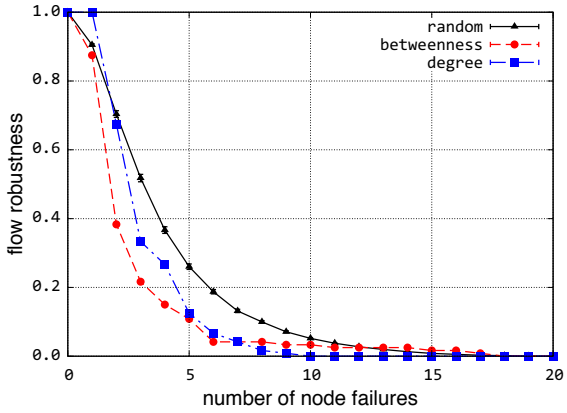

Figure A.15: Flow robustness for static routing during adaptive node deletions

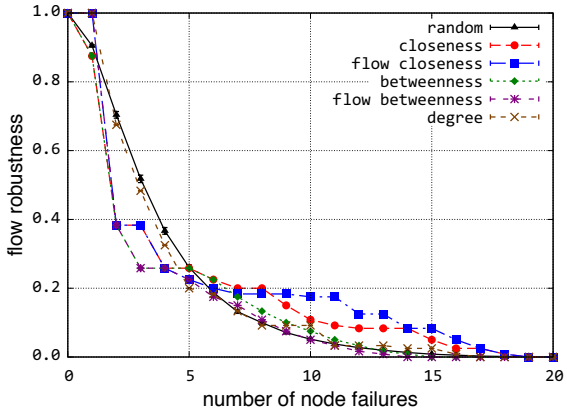

Figure A.16: Flow robustness for static routing during non-adaptive node deletions

A.1.5 Internet2

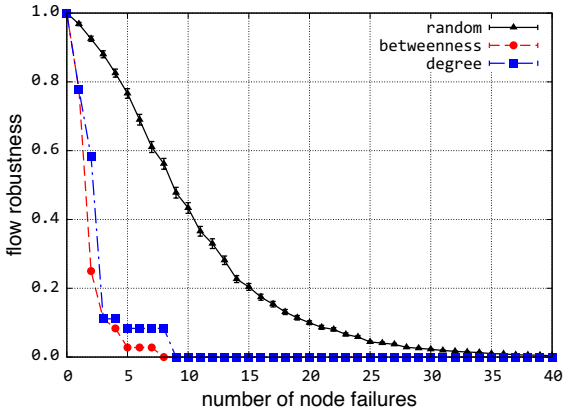

Figure A.17: Flow robustness for dynamic routing during adaptive node deletions

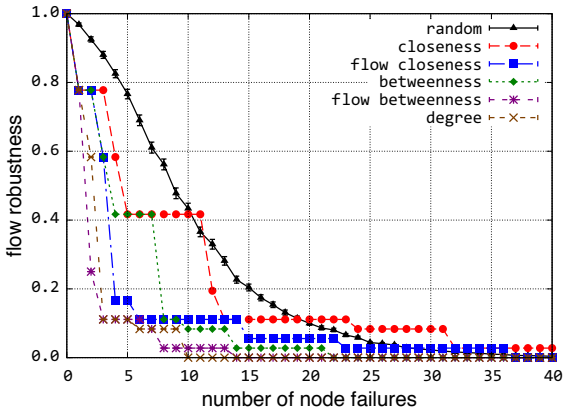

Figure A.18: Flow robustness for dynamic routing during non-adaptive node deletions

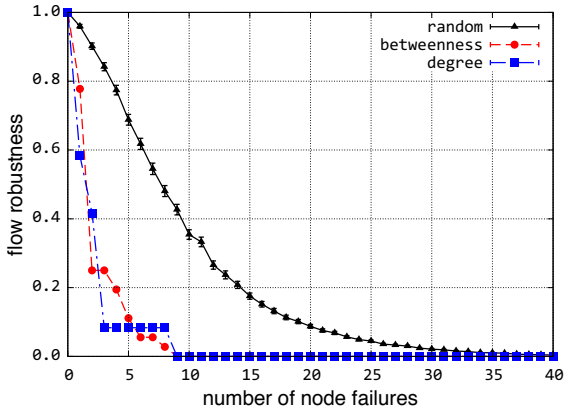

Figure A.19: Flow robustness for static routing during adaptive node deletions

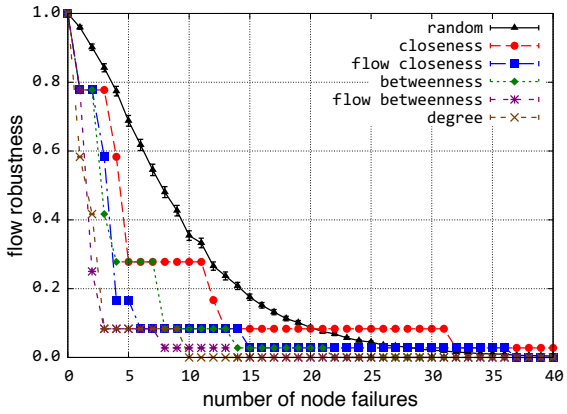

Figure A.20: Flow robustness for static routing during non-adaptive node deletions

A.2 Link Deletions

A.2.1 AT&T

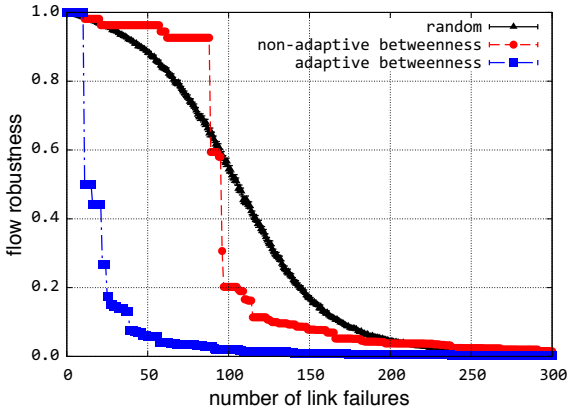

Figure A.21: Flow robustness for dynamic routing during adaptive and non-adaptive link deletions

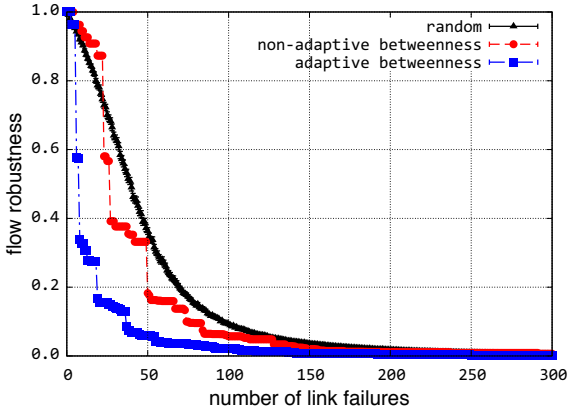

Figure A.22: Flow robustness for static routing during adaptive and non-adaptive link deletions

A.2.2 Level 3

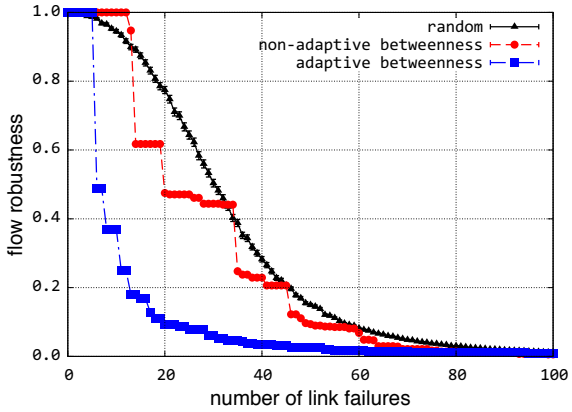

Figure A.23: Flow robustness for dynamic routing during adaptive and non-adaptive link deletions

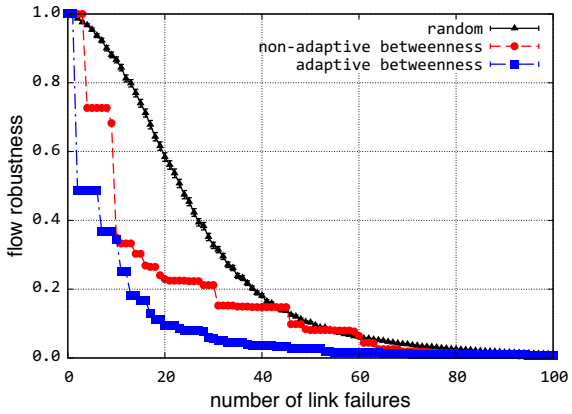

Figure A.24: Flow robustness for static routing during adaptive and non-adaptive link deletions

A.2.3 Sprint

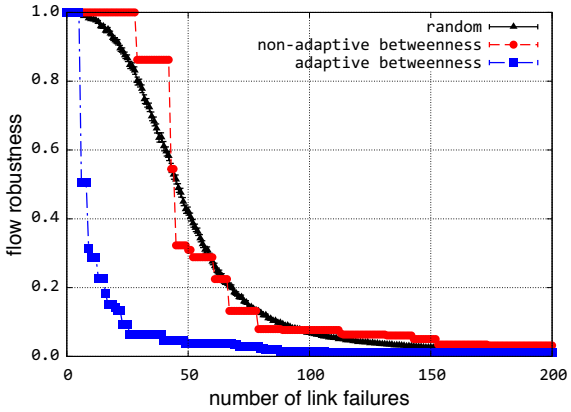

Figure A.25: Flow robustness for dynamic routing during adaptive and non-adaptive link deletions

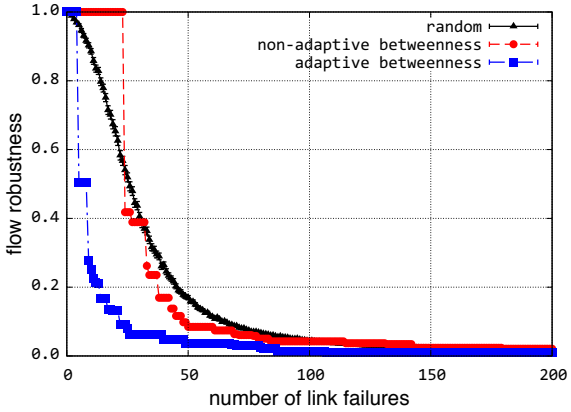

Figure A.26: Flow robustness for static routing during adaptive and non-adaptive link deletions

A.2.4 TeliaSonera

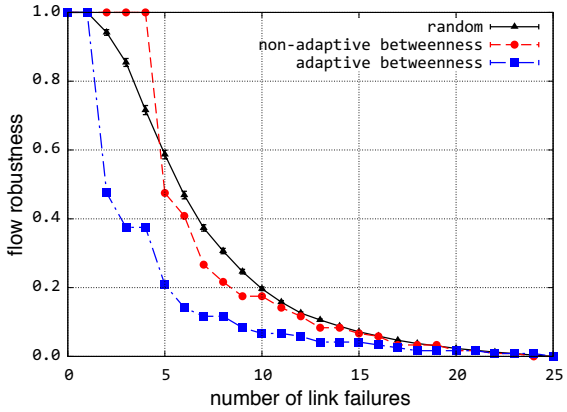

Figure A.27: Flow robustness for dynamic routing during adaptive and non-adaptive link deletions

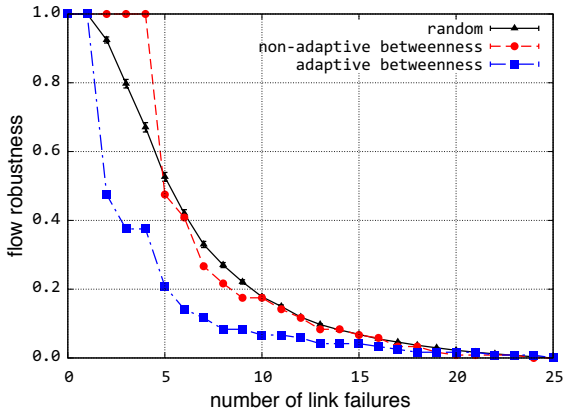

Figure A.28: Flow robustness for static routing during adaptive and non-adaptive link deletions

A.2.5 Internet2

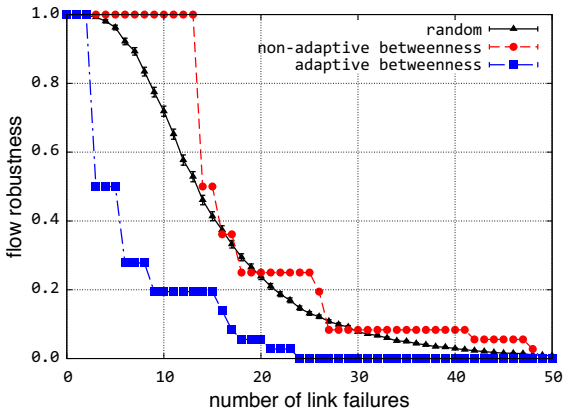

Figure A.29: Flow robustness for dynamic routing during adaptive and non-adaptive link deletions

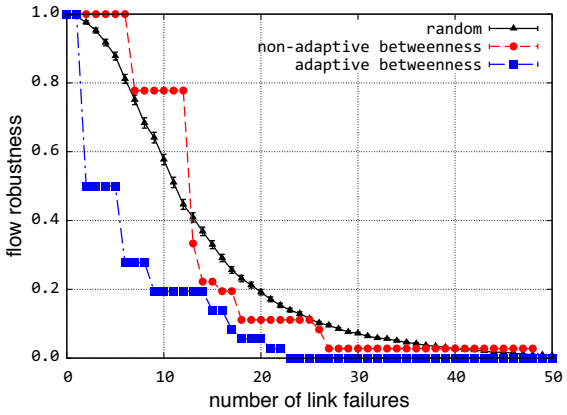

Figure A.30: Flow robustness for static routing during adaptive and non-adaptive link deletions

Appendix B

Graph Optimisation Plots

This appendix contains a full set of plots for the topologies used in the analysis of the *graph optimisation algorithms*.

B.1 Graph Optimisation via Algebraic Connectivity

B.1.1 Physical-Level Graphs

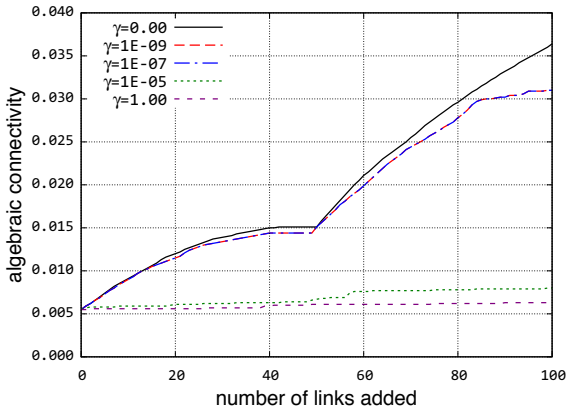

Figure B.1: AT&T connectivity improvement

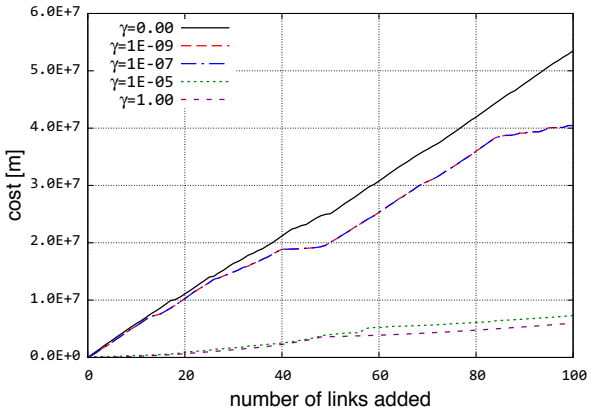

Figure B.2: AT&T cost incurred with adding links

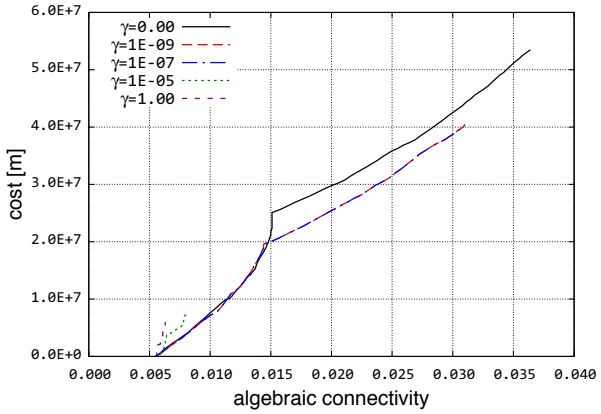

Figure B.3: Connectivity and cost trade-offs for AT&T

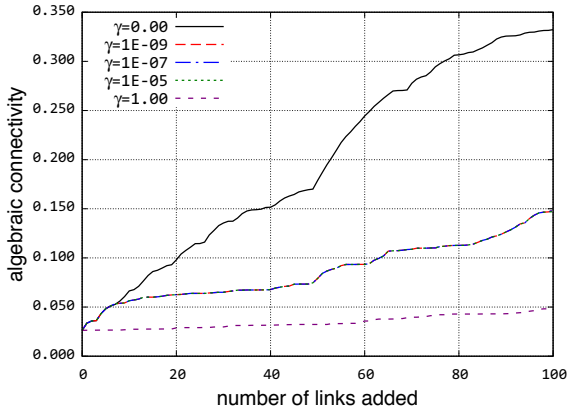

Figure B.4: Level 3 connectivity improvement

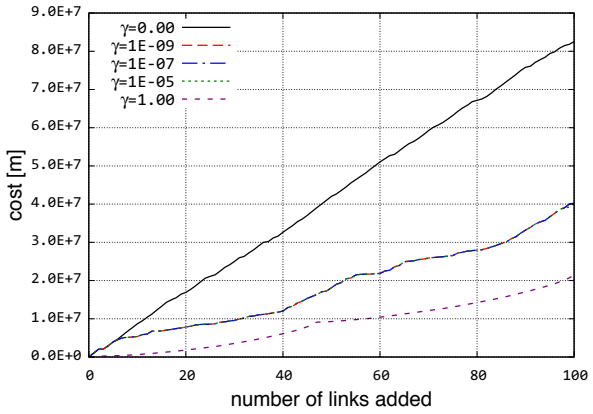

Figure B.5: Level 3 cost incurred with adding links

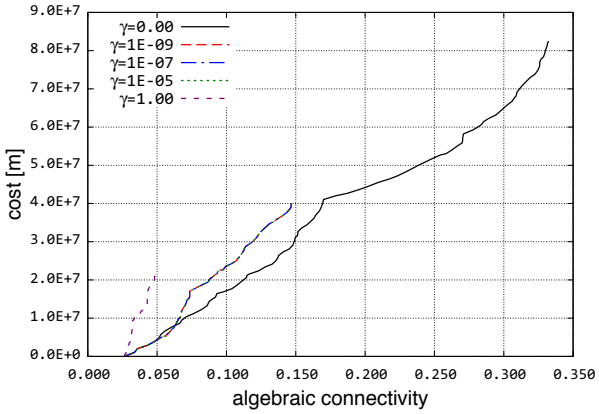

Figure B.6: Connectivity and cost trade-offs for Level 3

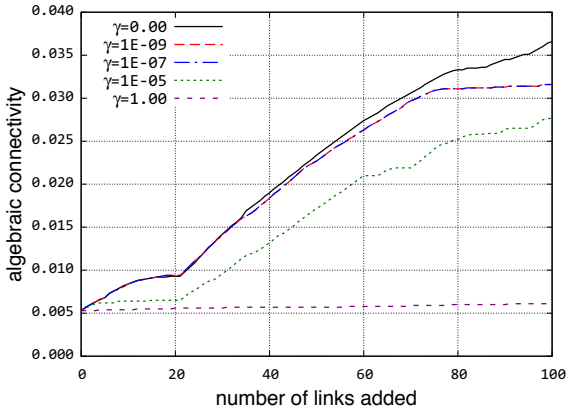

Figure B.7: Sprint connectivity improvement

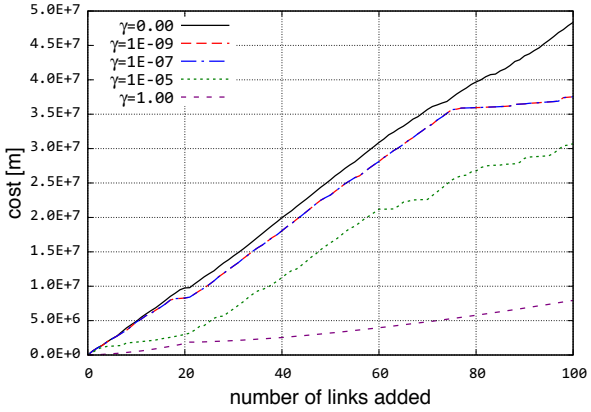

Figure B.8: Sprint cost incurred with adding links

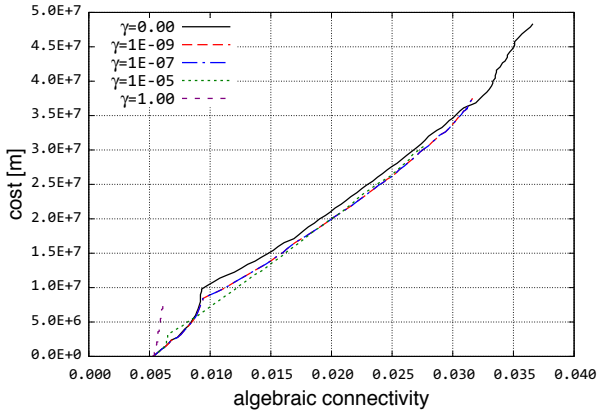

Figure B.9: Connectivity and cost trade-offs for Sprint

B.1.2 Logical-Level Graphs

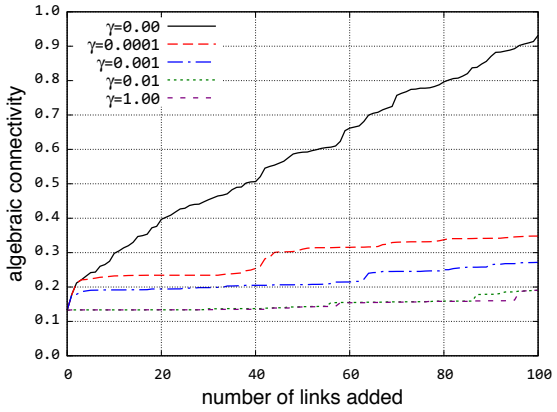

Figure B.10: AT&T connectivity improvement

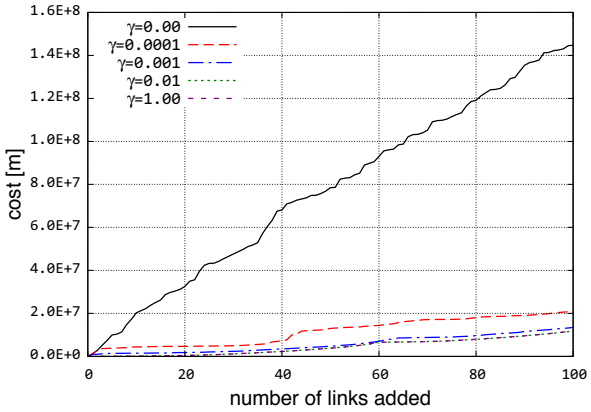

Figure B.11: AT&T cost incurred with adding links

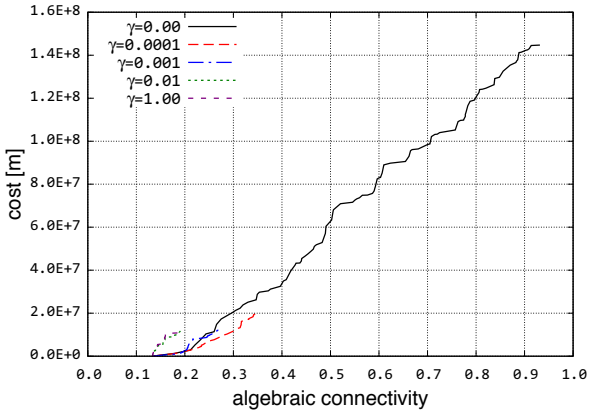

Figure B.12: Connectivity and cost trade-offs for AT&T

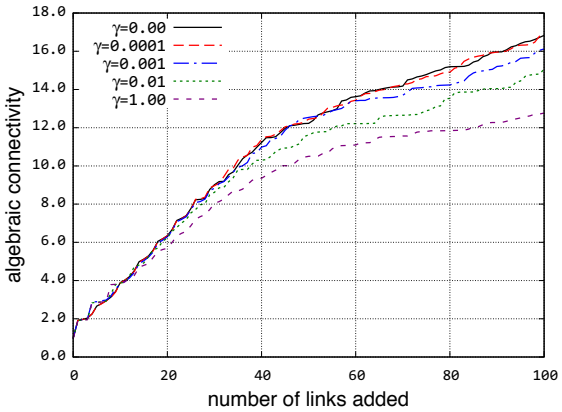

Figure B.13: Level 3 connectivity improvement

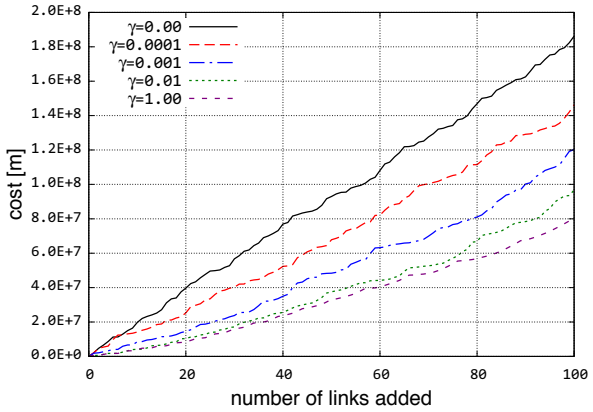

Figure B.14: Level 3 cost incurred with adding links

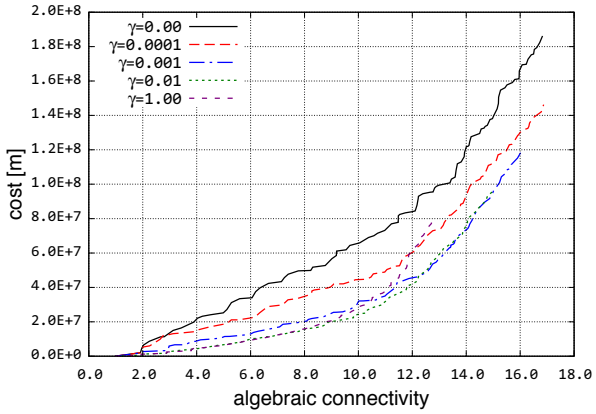

Figure B.15: Connectivity and cost trade-offs for Level 3

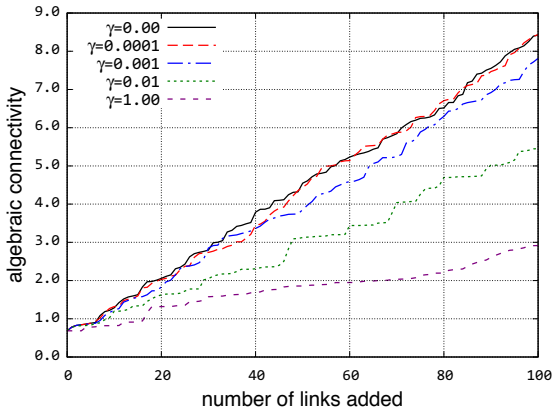

Figure B.16: Sprint connectivity improvement

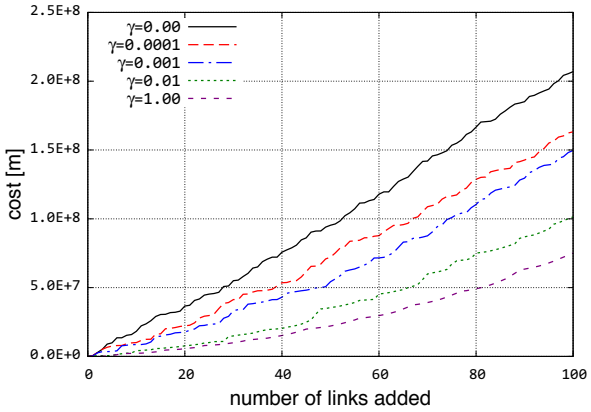

Figure B.17: Sprint cost incurred with adding links

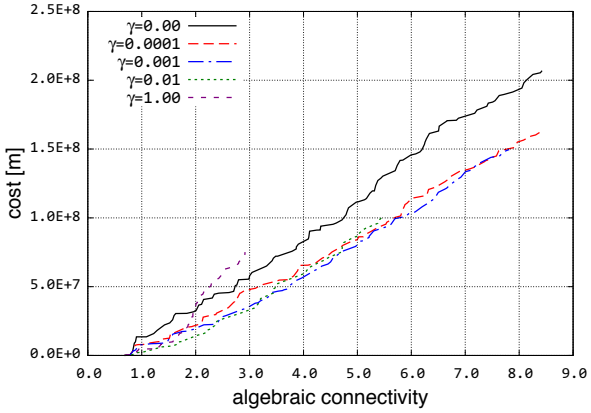

Figure B.18: Connectivity and cost trade-offs for Sprint

B.1.3 Comparison of Providers

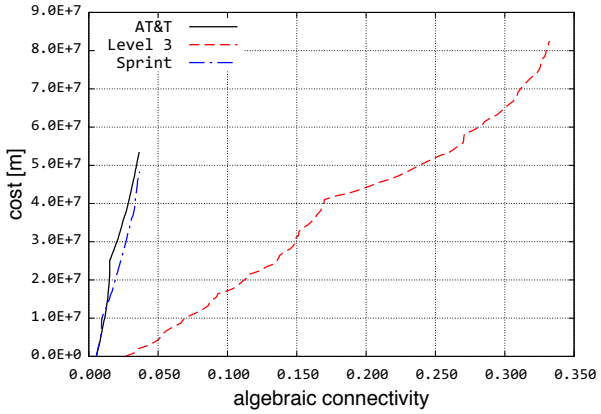

Figure B.19: Algebraic connectivity and cost effect for $\gamma = 0$ for physical level topologies

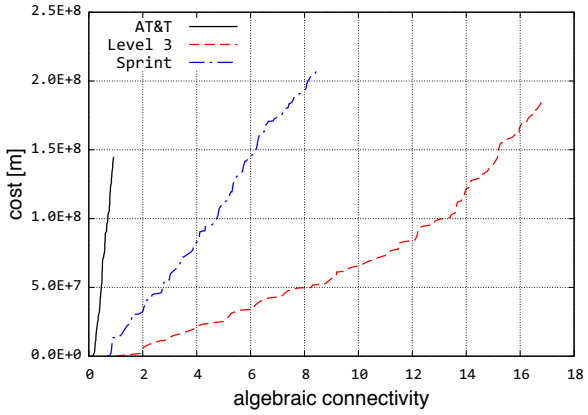

Figure B.20: Algebraic connectivity and cost effect for $\gamma = 0$ for logical level topologies

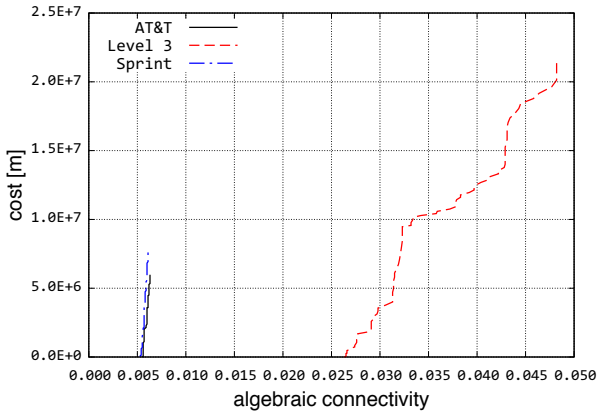

Figure B.21: Algebraic connectivity and cost effect for $\gamma = 1$ for physical level topologies

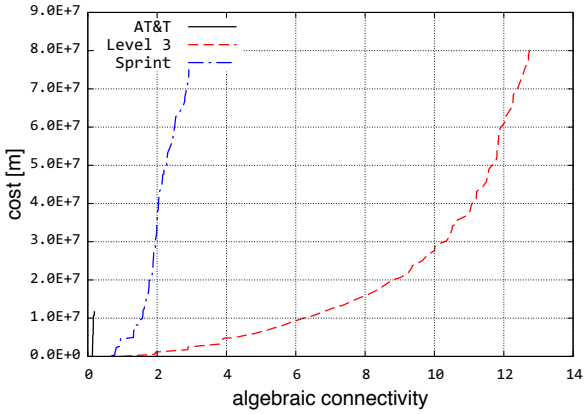

Figure B.22: Algebraic connectivity and cost effect for $\gamma = 1$ for logical level topologies

B.2 Graph Optimisation via Path Diversity

B.2.1 Impact of Varying Hop Count Threshold on TGD and Cost

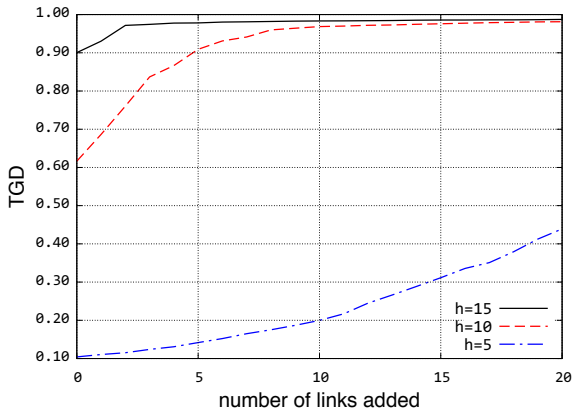

Figure B.23: CORONET TGD improvement

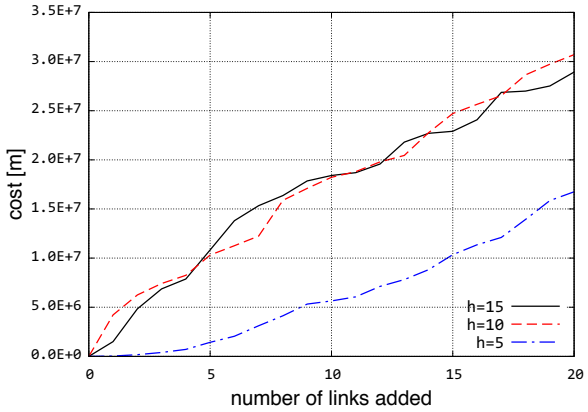

Figure B.24: CORONET cost incurred

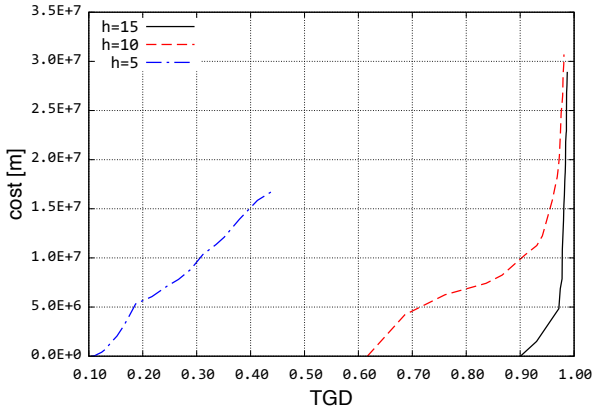

Figure B.25: CORONET cost and TGD

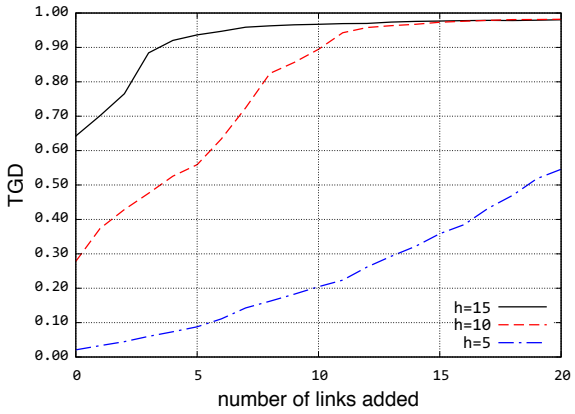

Figure B.26: Internet2 TGD improvement

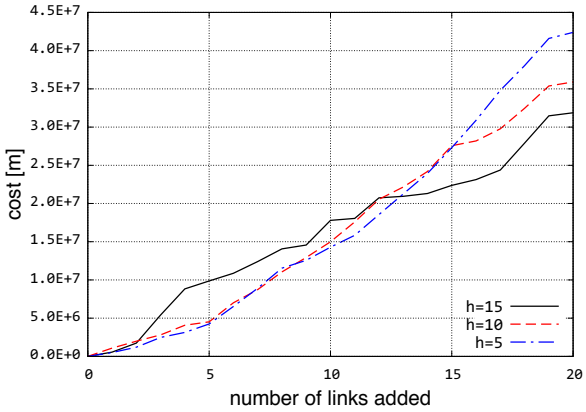

Figure B.27: Internet2 cost incurred

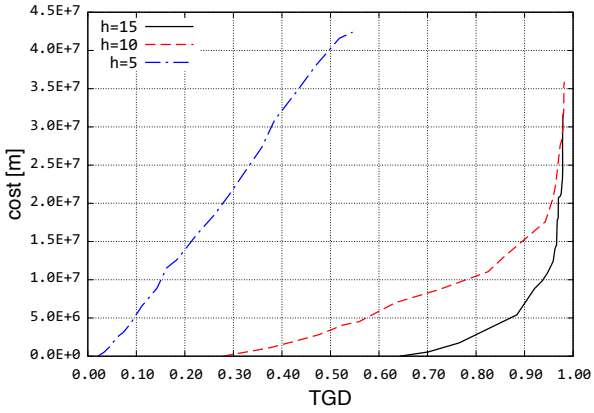

Figure B.28: Internet2 cost and TGD

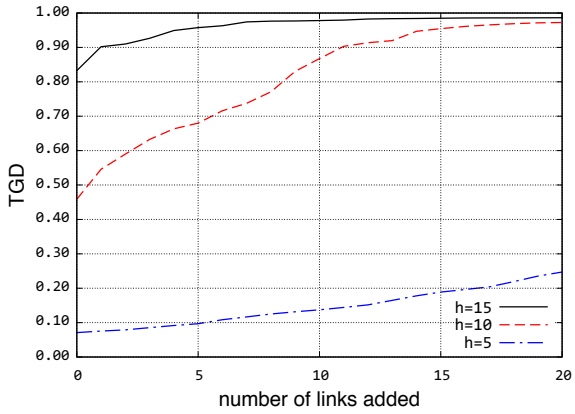

Figure B.29: Level 3 TGD improvement

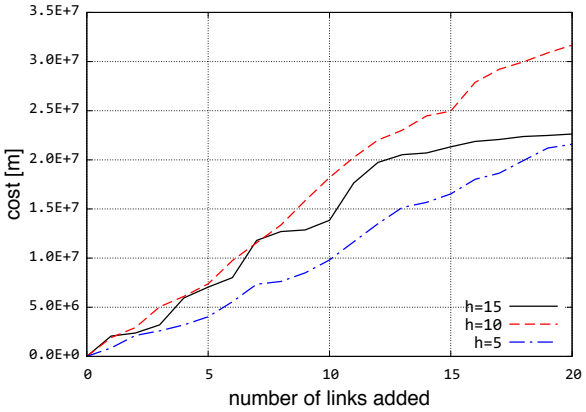

Figure B.30: Level 3 cost incurred

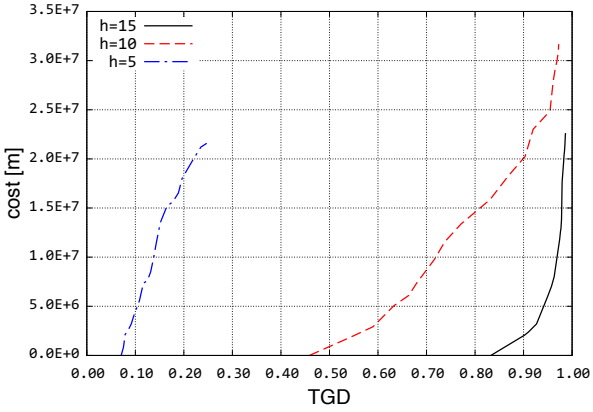

Figure B.31: Level 3 cost and TGD

B.2.2 Impact of Varying k on TGD and Cost

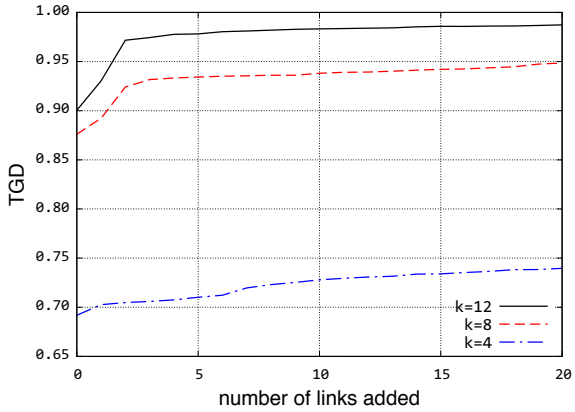

Figure B.32: CORONET TGD improvement

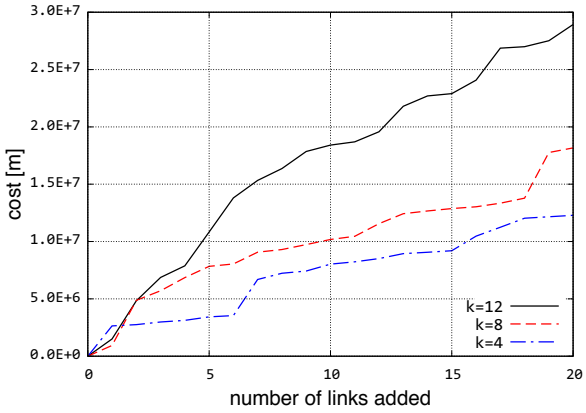

Figure B.33: CORONET cost incurred

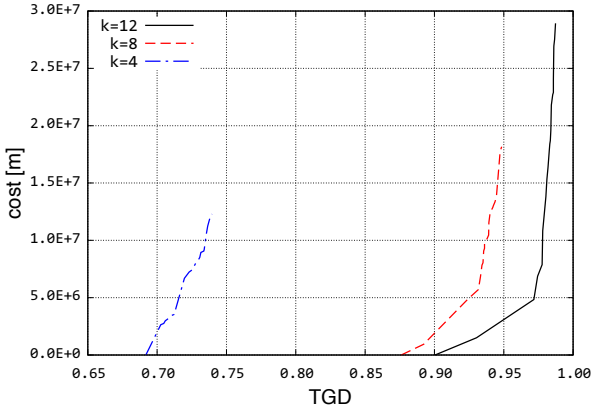

Figure B.34: CORONET cost and TGD

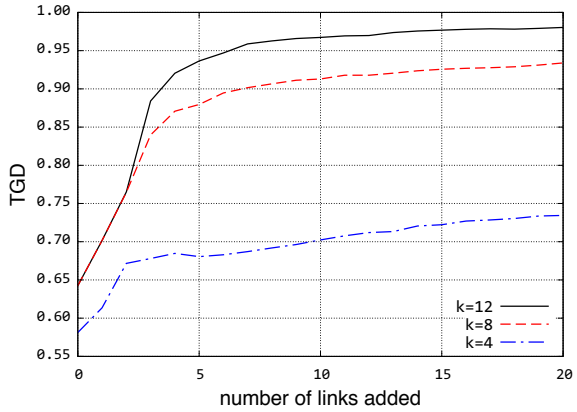

Figure B.35: Internet2 TGD improvement

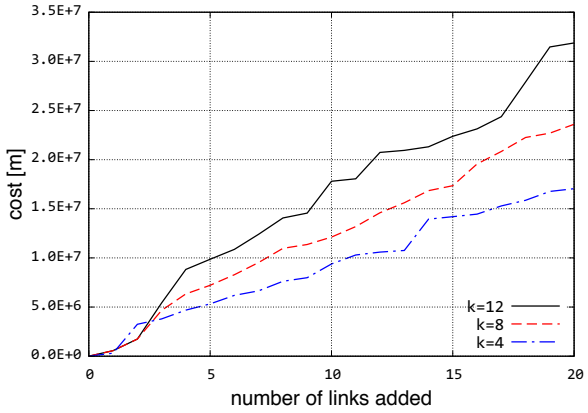

Figure B.36: Internet2 cost incurred

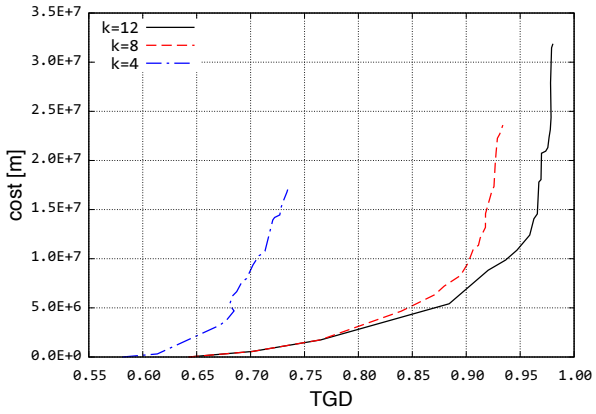

Figure B.37: Internet2 cost and TGD

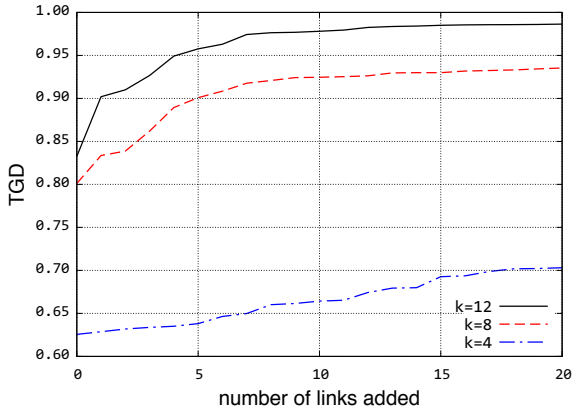

Figure B.38: Level 3 TGD improvement

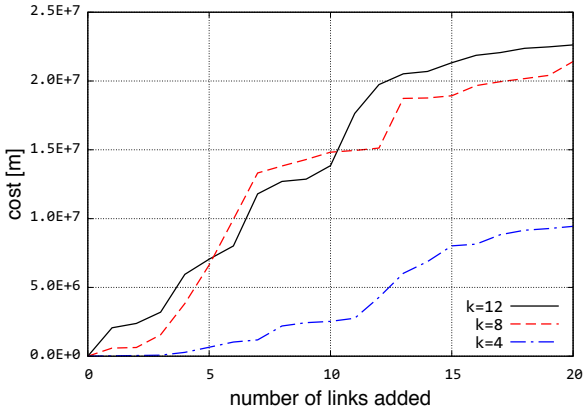

Figure B.39: Level 3 cost incurred

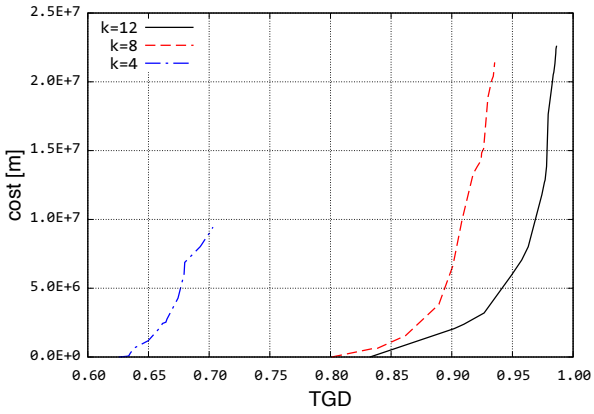

Figure B.40: Level 3 cost and TGD

B.2.3 Flow Robustness Analysis of Graph Optimisation

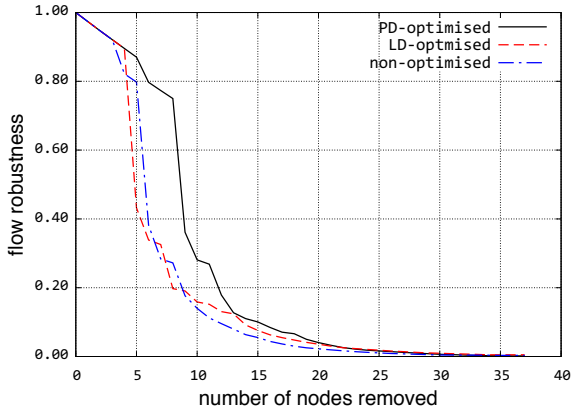

Figure B.41: CORONET betweenness-based attack

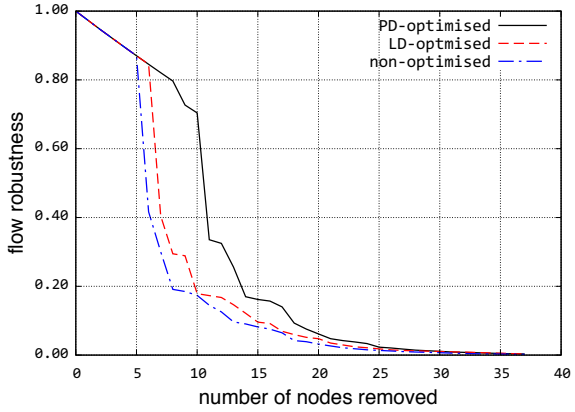

Figure B.42: CORONET closeness-based attack

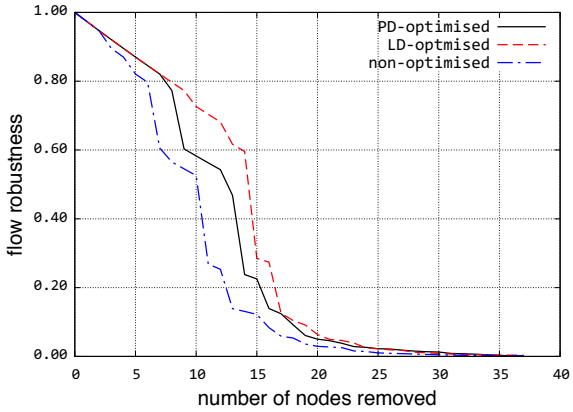

Figure B.43: CORONET degree-based attack

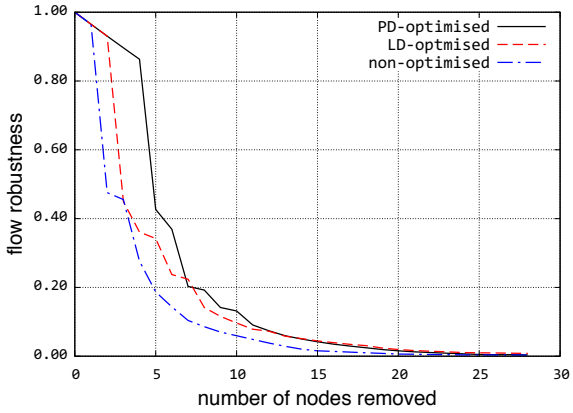

Figure B.44: Internet2 betweenness-based attack

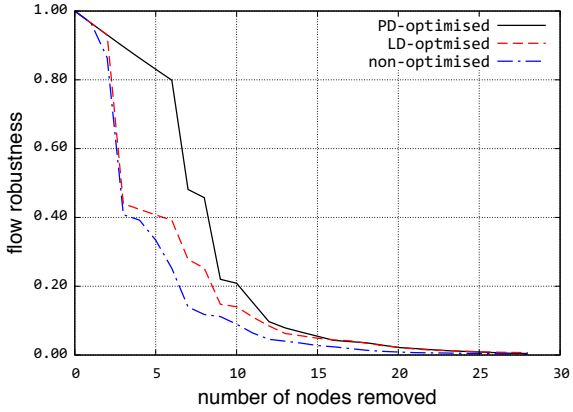

Figure B.45: Internet2 closeness-based attack

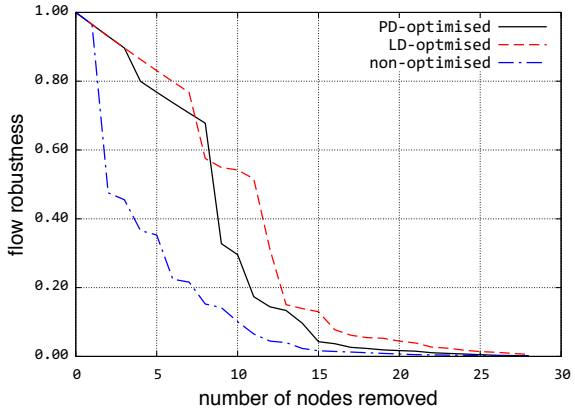

Figure B.46: Internet2 degree-based attack

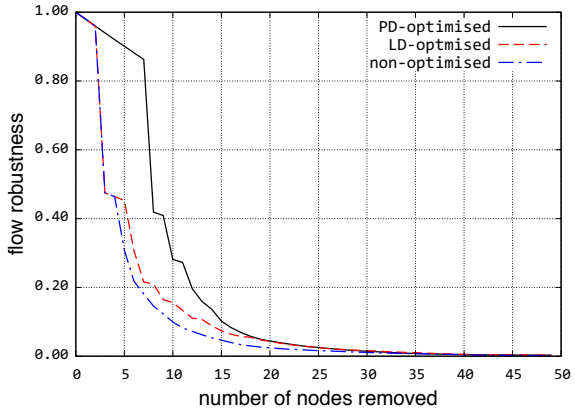

Figure B.47: Level 3 betweenness-based attack

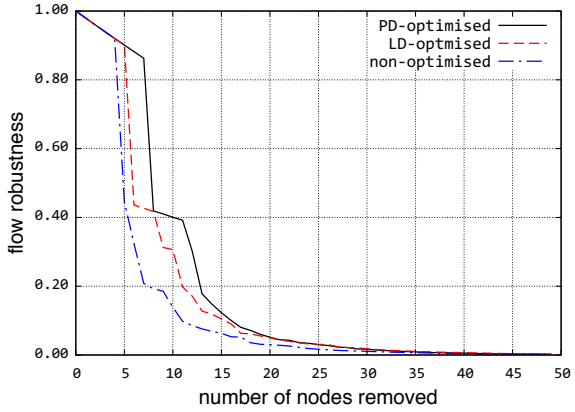

Figure B.48: Level 3 closeness-based attack

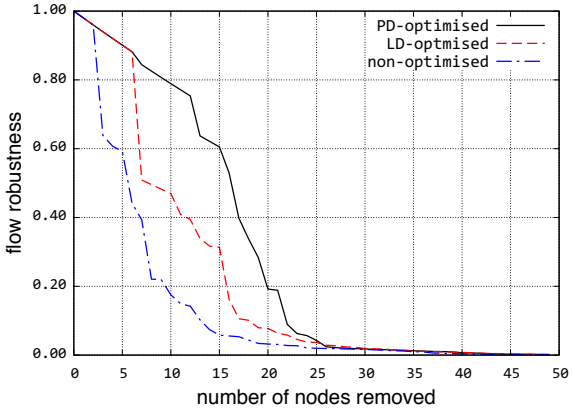

Figure B.49: Level 3 degree-based attack