TU Delft

Delft University of Technology

A Data-Driven Approach to Disaster Resilience in Communication Networks

Oostenbrink, J.

DOI
10.4233/uuid:323b614e-c766-4c15-8fff-8d4890c61806

Publication date
2023

Document Version
Final published version

Citation (APA)
Oostenbrink, J. (2023). *A Data-Driven Approach to Disaster Resilience in Communication Networks*. [Dissertation (TU Delft), Delft University of Technology]. https://doi.org/10.4233/uuid:323b614e-c766-4c15-8fff-8d4890c61806

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.
For technical reasons the number of authors shown on this cover page is limited to a maximum of 10.

A Data-Driven Approach to Disaster Resilience in Communication Networks



Jorik Oostenbrink

A Data-Driven Approach to Disaster Resilience in Communication Networks

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus, prof. dr. ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates
to be defended publicly on Wednesday 8 February 2023 at 12:30 o'clock

by

Jorik Oostenbrink

Master of Science in Computer Science,
Delft University of Technology, the Netherlands,
born in Groningen, the Netherlands.

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Prof. dr. ir. F.A. Kuipers	Delft University of Technology, promotor
Prof. dr. K.G. Langendoen	Delft University of Technology, promotor
Independent members:	
Prof. dr. ir. P.H.A.J.M. van Gelder	Delft University of Technology
Prof. dr. M.M. de Weerdt	Delft University of Technology
Prof. dr. C.E.W. Hesselman	University of Twente
Dr hab. inż. J. Rak	Gdańsk University of Technology
Dr. ir. E. Vriezekolk	Rijksinspectie Digitale Infrastructuur

Keywords: Network Resilience, Natural Disasters, Regional Failures, Geographically Correlated Failures

Printed by: Ipskamp Printing

Cover design: Niek Vergeer

Copyright © 2023 by J. Oostenbrink

ISBN 978-94-6384-410-9

An electronic version of this dissertation is available at
http://repository.tudelft.nl/.

CONTENTS

Summary vii

Samenvatting ix

1 Introduction 1
- 1.1 Regional Failures . . . 2
- 1.2 Resilience . . . 3
  - 1.2.1 Metric . . . 4
- 1.3 The Cost of Resilience . . . 5
  - 1.3.1 Protecting Against All Disasters . . . 5
  - 1.3.2 Random Disaster Models . . . 8
  - 1.3.3 Data-Driven Approaches . . . 10
- 1.4 Data-Driven Disaster Resilience . . . 11
  - 1.4.1 Problem Statement . . . 11
  - 1.4.2 Our Approach . . . 11

2 A Global Study of the Risk of Earthquakes to IXPs 13
- 2.1 Introduction . . . 13
- 2.2 Related Work . . . 15
- 2.3 Datasets . . . 16
  - 2.3.1 IXPs . . . 16
  - 2.3.2 OpenQuake Engine . . . 17
- 2.4 Facilities at Risk . . . 18
  - 2.4.1 Methods . . . 18
  - 2.4.2 Results . . . 21
  - 2.4.3 Country-Level Analysis . . . 24
  - 2.4.4 Conterminous United States . . . 26
- 2.5 Combined Failures . . . 27
  - 2.5.1 Disruption . . . 27
  - 2.5.2 Increasing Redundancy - A Novel Metric . . . 32
  - 2.5.3 Overlap Between Hazard Models . . . 33
- 2.6 Discussion . . . 34
- 2.7 Conclusion . . . 34

3 Computing the Impact of Disasters on Networks 37
- 3.1 Introduction . . . 38
- 3.2 Model . . . 38
- 3.3 J-SHIS Earthquake Scenarios . . . 39
- 3.4 Vulnerability Distributions . . . 40
  - 3.4.1 Failure States . . . 40

3.4.2 metrics . . . 40
3.5 Disaster Impact Visualization . . . 41
3.6 Experimental Results . . . 42
3.7 Conclusion . . . 44

4 The Risk of Successive Disasters: A Blow-by-Blow Network Vulnerability Analysis 45
4.1 Introduction . . . 46
4.2 Network and Disaster Model . . . 47
4.2.1 Example Network and Disasters Instance . . . 47
4.3 Problem Statement . . . 48
4.3.1 Number of Successive Disasters N . . . 49
4.3.2 Impact . . . 49
4.3.3 Total Time to Full Recovery . . . 49
4.4 Analysis . . . 49
4.4.1 Markov Chain . . . 49
4.4.2 Number of Successive Disasters N . . . 50
4.4.3 Impact . . . 52
4.4.4 Total Time to Full Recovery . . . 53
4.5 Monte Carlo . . . 54
4.6 Experiments . . . 54
4.6.1 Dataset . . . 54
4.6.2 The Effect of Component Repair Time . . . 57
4.6.3 Concurrent Repair . . . 59
4.7 Related Work . . . 60
4.8 Conclusion . . . 61

5 A Moment of Weakness: Protecting Against Targeted Attacks Following a Natural Disaster 63
5.1 Introduction . . . 64
5.2 Framework . . . 64
5.3 Experiments . . . 66
5.3.1 Impact . . . 67
5.3.2 Repair Strategies . . . 68
5.4 Conclusion . . . 71

6 Evaluating Local Disaster Recovery Strategies 73
6.1 Introduction . . . 74
6.2 Evaluation Model . . . 74
6.2.1 Local Area . . . 75
6.2.2 Evaluation Metrics . . . 76
6.3 Recovery Strategies . . . 78
6.3.1 Optimal Strategy . . . 78
6.3.2 Simple Strategies . . . 79

6.4 Algorithm . . . 80
6.5 Experiments . . . 81
6.6 Related Work . . . 83
6.7 Conclusion . . . 84

7 Going the Extra Mile with Disaster-Aware Network Augmentation 87
7.1 Introduction . . . 88
7.2 Problem Statement . . . 89
7.2.1 Cable Costs . . . 90
7.2.2 Disaster-Aware Network Augmentation Problem . . . 90
7.3 NP-hardness . . . 91
7.4 Branch and Bound . . . 93
7.4.1 Branching . . . 95
7.4.2 Bounding . . . 95
7.4.3 Shortest Route Computations . . . 96
7.4.4 Global Optimization . . . 96
7.5 Heuristic . . . 97
7.5.1 Simulated Annealing . . . 97
7.6 Experiments . . . 98
7.6.1 Connecting Node Pairs . . . 99
7.6.2 Global Solution . . . 100
7.6.3 Resilience Against New Disasters . . . 101
7.6.4 Number of Representative Disasters . . . 102
7.7 Related Work . . . 103
7.8 Conclusion . . . 104

8 Conclusion 105
8.1 Data-Driven Disaster Resilience Assessment . . . 105
8.2 Data-Driven Disaster Resilience Improvement . . . 106
8.3 Scalability . . . 107
8.4 Future Work . . . 108

Acknowledgements 111

Curriculum Vitæ 113

List of Publications 115

References 116

SUMMARY

Communication networks are critical in business, government, and even our day-to-day life. A prolonged communication outage can have devastating effects, particularly during and after a disaster. Unfortunately, our communication infrastructure is still vulnerable to natural disasters and other events that damage multiple network components within a confined area. In this thesis, we study the disaster resilience of communication networks. We propose scalable, data-driven methods to help stakeholders both assess and improve the resilience of networks to disasters.

We first study the global risk of earthquakes to Internet Exchange Points (IXPs). We find that many facilities are at risk of earthquakes and that, when an earthquake occurs, it is not unlikely that multiple facilities will fail simultaneously. Fortunately, our analysis also shows that larger IXPs tend to be located in less earthquake-prone areas, and that peering at multiple facilities significantly reduces the impact of earthquakes to IXPs and autonomous systems. To help network operators in reducing the impact of earthquakes on their autonomous systems, we propose a novel metric for selecting peering facilities, based on the probability of simultaneous facility failures. We show that applying our metric can significantly increase the resilience of individual autonomous systems, as well as that of the Internet as a whole.

To effectively improve the resilience of communication networks to natural disasters, stakeholders need to make well-informed trade-offs between costs, network performance, and network resilience. To help stakeholders make these decisions, we propose a single-disaster and a successive-disaster framework for assessing the resilience of a network to natural disasters. These frameworks can help stakeholders anticipate potential disasters, and compare the effects of any trade-off on the resilience of their networks.

The main principle behind both frameworks is to assess the disaster resilience of a network based on a large set of representative disaster scenarios (called the disaster set). This approach is flexible with respect to the underlying disaster dataset, and can be applied to datasets of widely varying sizes and properties. Our single-disaster framework allows one to efficiently compute the distribution of a network performance metric, assuming that a single, random disaster strikes the network and damages one or more network components in a confined area. Our method speeds up computation by first computing the distribution of the state of the network after a random disaster (the number of possible states tends to be much smaller than the disaster set itself), and only then computing the performance of the network in each of these states.

In addition to studying the impact of a single disaster on a network, we also address the issue of successive disasters. We first define the concept of successive disasters: a subsequent disaster that strikes the network while the damage due to a previous disaster is still being repaired. We then propose a framework capable of modeling a sequence of disasters in time, while taking into account recovery operations. We develop both an exact and a Monte Carlo method to compute the vulnerability of a network to successive

disasters and find that the probability of a second disaster striking the network during recovery can be significant even for short repair times.

Our successive disaster framework can not only be applied to subsequent disasters, but also to potential follow-up attacks. Experiments on two network topologies show that even small targeted attacks can greatly aggravate the network disruption caused by a natural disaster. Fortunately, we find that this effect can be mitigated - at almost no cost to network performance - by adopting a calculated repair strategy that takes into account the possibility of follow-up attacks.

In addition to providing methods for assessing the resilience of networks, we also provide algorithms for improving the resilience of networks to natural disasters. These algorithms can help stakeholders (1) recover network functionality more effectively in the initial period after a disaster, and (2) reduce the initial impact of a disaster on network performance.

After a disaster, a network operator can quickly restore some functionality by replacing nodes with temporary emergency nodes. These emergency nodes should be deployed as soon as possible. However, selecting an optimal set of replacement nodes is computationally intensive, and the complete state of the network might still be unknown after the disaster. Thus, we propose selecting a disaster strategy a priori - before the occurrence of the disaster. We give an algorithm for evaluating such strategies, by extending our single-disaster assessment framework.

An effective, but costly, method of improving the disaster resilience of a network is to add new, geographically redundant, cable connections. These redundant connections ensure that more areas remain connected after a disaster strikes the network, and thus reduce the initial impact of the disaster on the network. We provide algorithms for finding cable routes that minimize a function of disaster impact and cable cost under any disaster set. Since this problem is NP-hard, we give an exact algorithm, as well as a heuristic, for solving it.

SAMENVATTING

Communicatienetwerken zijn cruciaal voor bedrijven, overheden, en ons dagelijks leven. Een langdurige communicatiestoring kan desastreuze gevolgen hebben, in het bijzonder na en tijdens een ramp. Helaas is onze communicatie-infrastructuur nog steeds kwetsbaar voor natuurrampen en andere gebeurtenissen die meerdere netwerkcomponenten in een begrensd gebied beschadigen.
In dit proefschrift bestuderen we de weerbaarheid van communicatienetwerken tegen natuurrampen. We geven schaalbare, datagedreven methodes om stakeholders te assisteren in het beoordelen en verbeteren van de weerbaarheid van netwerken tegen rampen.

We bestuderen eerst het wereldwijde risico van aardbevingen voor Internet Exchange Points (IXPs). We bevinden dat veel datacenters risico lopen op aardbevingen, en dat het niet onwaarschijnlijk is dat, als er een aardbeving plaatsvindt, meerdere datacenters tegelijkertijd zullen uitvallen. Gelukkig laat onze analyse ook zien dat grotere IXPs veelal in gebieden liggen die minder vaak door aardbevingen geraakt worden, en dat de impact van aardbevingen op IXPs en autonome systemen significant verlaagd wordt als deze verpreid zijn over meerdere datacenters. Om netwerkbeheerders te assisteren in het verminderen van de impact van rampen op hun autonome systemen geven we een nieuwe metriek voor het selecteren van peering locaties, die gebaseerd is op de kans van simultane datacenter uitval. We laten zien dat het toepassen van onze metriek de weerbaarheid van individuele autonome systemen en het Internet als geheel significant kunnen verhogen.

Om de weerbaarheid van communicatienetwerken tegen natuurrampen effectief te kunnen verbeteren, moeten stakeholders goed geïnformeerde compromissen maken tussen kosten, netwerk prestatie, en weerbaarheid. Om stakeholders te helpen met het maken van deze beslissingen, geven wij een enkelvoudige-ramp en een opeenvolgende-ramp raamwerk voor het beoordelen van de weerbaarheid van een netwerk tegen natuurrampen. Deze raamwerken kunnen stakeholders assisteren in het anticiperen van mogelijke rampen, en in het vergelijken van de effecten van beslissingen op de weerbaarheid van hun netwerken.

De hoeksteen van beide raamwerken is het beoordelen van de weerbaarheid van communicatienetwerken tegen natuurrampen op basis van een grote groep representatieve rampscenario's (de "disaster set"). Deze aanpak is flexibel met betrekking tot de onderliggende dataset van rampen, en kan toegepast worden op datasets met sterk verschillende groottes en eigenschappen. Ons enkelvoudige-ramp raamwerk maakt het mogelijk om de kansverdeling van een impactsmetriek efficiënt te berekenen, onder de aanname dat een enkele, willekeurige ramp het netwerk raakt en één of meerdere netwerkcomponenten in een begrensd gebied beschadigt. Om de rekensnelheid van onze methode te verhogen berekenen we eerst de kansverdeling van alle mogelijke combinaties van uitval (normaliter zijn er veel minder mogelijke combinaties van uitval dan dat

er rampscenario's zijn), waarna we de impact van elke combinatie van uitval berekenen, in plaats van de impact van elk rampscenario.

Naast het bestuderen van de impact van enkelvoudige rampen op een netwerk, richten we ons ook op opeenvolgende rampen. We definiëren eerst het concept "opeenvolgende ramp": een ramp die plaatsvindt terwijl het netwerk nog hersteld wordt van de schade van een vorige ramp. Vervolgens geven we een raamwerk dat een reeks rampen over tijd kan modelleren en daarbij hersteloperaties meeneemt. We ontwikkelen een exacte en een Monte Carlo methode voor het berekenen van de kwetsbaarheid van een netwerk voor opeenvolgende rampen, en bevinden dat de kans op opeenvolgende rampen zelfs met korte reparatietijden significant kan zijn.

Ons raamwerk voor opeenvolgende rampen kan niet alleen op rampen toegepast worden, maar ook op potentiële aanvallen. Onze experimenten op twee netwerktopologiëen laten zien dat zelfs kleine gerichte aanvallen de impact van een natuurramp op een netwerk sterk kunnen verergeren. Gelukkig bevinden we dat dit effect gemitigeerd kan worden - met bijna geen nadelige effecten op het netwerk - door het meenemen van de mogelijkheid van gerichte aanvallen in de reparatie-strategie.

Naast methodes om de weerbaarheid van netwerken te beoordelen, geven we ook algoritmes om de weerbaarheid van netwerken tegen natuurrampen te verbeteren. Deze algoritmes kunnen stakeholders assisteren in (1) het effectiever herstellen van netwerk functionaliteit in de aanvankelijke periode na een ramp, en (2) het verminderen van de aanvankelijke impact van een ramp op de netwerk prestatie.

Na een ramp kan een netwerkbeheerder snel functionaliteit herstellen door netwerknodes te vervangen door tijdelijke noodnodes. Deze noodnodes moeten zo snel mogelijk ingezet worden. Het berekenen van een optimale selectie van nodes om te vervangen is echter tijdrovend, en de volledige staat van het netwerk hoeft nog niet meteen bekend te zijn na de ramp. Wij stellen daarom voor om de reparatiestrategie a priori te bepalen - voordat de ramp heeft plaatsgevonden. We geven een algoritme voor het evalueren van zulke reparatiestrategieën dat bouwt op ons raamwerk voor enkelvoudige rampen.

Een effectieve, maar dure, methode voor het verbeteren van de weerbaarheid van netwerken tegen natuurrampen is het toevoegen van nieuwe, geografisch redundante, kabelverbindingen. Deze redundante verbindingen verzekeren dat meer gebieden na de ramp verbonden zullen blijven, en verminderen dus de aanvankelijke impact van een ramp op het netwerk. We geven algoritmes voor het vinden van kabelroutes die, gegeven een groep rampscenario's, een functie van rampimpact en kabelkosten minimalizeren. Omdat dit probleem NP-hard is, geven we naast een exact algoritme ook een heuristiek om het op te lossen.

1

INTRODUCTION

Communication networks have become integral to businesses, governments, and even our day-to-day life. As of 2019, more than 50 percent of the world's population is connected to the Internet [1] - itself a network of communication networks. This pervasiveness gives communication (together with energy) infrastructure somewhat of a unique position among all of our critical infrastructures: Communication systems provide "enabling functions" across all other critical infrastructures [2]. The failure of our communication networks can cause havoc on our critical infrastructures, economy, as well as our personal lives.

The Netherlands experienced the danger of our dependency on communication networks firsthand in 2019: On June 24, KPN (a Dutch telecommunications provider) experienced a 3.5 hour outage [3]. As the Dutch government relied on KPN to forward emergency calls, citizens in the entire country were unable to call for emergency services. At the same time, another outage at KPN also prevented the government from reaching KPN customers through the Public Warning System, which made it more difficult to keep all citizens informed.

Society's dependency on communication networks only increases in the run-up to and aftermath of a natural disaster. Citizens rely on communication networks to obtain information, contact their loved ones, and call for help; emergency services rely on communication networks to organize an effective disaster response; and government agencies rely on communication networks to warn and inform citizens, as well as communicate internally.

Unfortunately, while society's dependency on communication networks increases during and after a natural disaster, the same disaster often has devastating effects on the communication networks themselves. For example, in 2018, an earthquake and resulting tsunami struck the Indonesian island of Sulawesi, killing thousands of people. Just before the tsunami struck, multiple tsunami warnings were sent out to the public. However, in addition to a number of other issues with the early warning system, the communication and power infrastructures were already damaged by the earthquake itself, and these warnings may not have been received by local residents [4]. The earthquake

and tsunami severely damaged the infrastructure of national telecom operators [5]. In particular, the area of Donggala was completely cut off from the outside world.

Hurricane Irma knocked out communications in Sint-Maarten in 2017. The lack of communication delayed the international humanitarian response and severely disrupted the ability of the government to inform their citizens [6, 7].

In 2021, hurricane Ida made landfall in Louisiana, U.S. The combination of flooding, storm damage, and power outages had a massive impact on local communication networks [8, 9]. As a result, numerous areas lost cell, landline, broadband, and/or cable service. In addition, emergency services were unreachable in multiple areas, due to the failure of 911 systems.

These are just some of the recent examples where a natural disaster both knocked out communication infrastructure, while also increasing the need for functional communication networks. There is evidence that an increase in cell phone access can save lives during a disaster: Statistically, an increase in a country's cell phone usage has been found to reduce disaster fatalities [10]. With the rise of "the Internet of Things", our dependency on communication networks only keeps growing. At the same time, the frequency of weather and climate disasters, such as extreme precipitation and tropical cyclones, has increased as well [11]. It is thus becoming more and more important to protect our networks to these and other types of natural disasters.

1.1. REGIONAL FAILURES

In this thesis, we consider the resilience of communication networks to physical damage. Within this context, the main driver of resilience is *redundancy*.

A communication network is often modeled as a graph. The nodes in the graph represent devices (e.g. routers, servers, base stations, or even mobile phones), while the links in the graph represent the physical (or in some cases virtual) connections between these nodes. For example, Fig. 1.1 shows a simple line topology of 6 nodes. If node 1 wants to send a message to node 6, it first passes through nodes 2 to 5, and vice-versa.

A line topology is not very resilient to failures; if node 2, 3, 4, or 5 fails, the network is split into two parts (called *connected components*). Suddenly, node 1 is not able to communicate with node 6 anymore, even though both nodes themselves are still functioning.

A ring topology, as shown in Fig. 1.2, is already much more resilient. By adding a single link between nodes 1 and 6, we have made the network resilient to any single hardware failure. Whichever node or link fails, there will always be a possible path between any of the surviving nodes (albeit with a reduction in total bandwidth capacity and a potential increase in latency).

From a topological perspective, the ring topology in Fig. 1.2 seems quite resilient. However, we have completely ignored the location of each network component. It is cheaper to host nodes in the same datacenter, and to route links through the same cables or ducts. Suppose nodes 1 and 4 share a datacenter, as illustrated in Fig. 1.3. The network would still be able to handle any random node or link failure. However, it would not be able to handle the failure of a whole datacenter (for example, due to a loss of power). This kind of failures, where multiple components within a region fail together, are called *regional failures* or geographically correlated challenges.

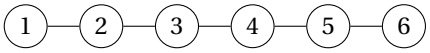


Figure 1.1: 6-node line topology.

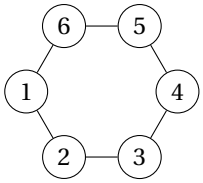
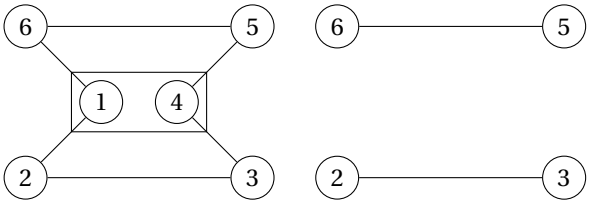


Figure 1.2: 6-node ring topology.



(a) Before regional failure. (b) After regional failure.

Figure 1.3: 6-node ring topology. Nodes 1 and 4 share a common location.

Regional failures can range from the very small (e.g., the failure of 2 cables sharing a common duct) to the very large (e.g., due to a natural disaster). Past events have repeatedly shown that many communication networks are vulnerable to regional failures all across this scale [12–15]. Thus, we need approaches to network resilience engineering that take into account both the structure of a network, as well as the locations of its components.

1.2. RESILIENCE

In the previous section, we referred to resilience in an intuitive manner. In fact, most works on the disaster resilience of communication networks do not explicitly define what "resilience" means, simply relying on a common understanding of the word. It turns out, however, that resilience is quite a broad concept. As a result, resilience (and related terms) has many different definitions across many different disciplines [16].

There does seem to be a general consensus that resilience is not just about mitigating the initial impact of an adverse event. According to Hickford et al., four main principles emerge across many of the definitions of resilience: (1) anticipate, (2) absorb, (3) adapt, and (4) recover. In a disaster-resilient network, the operator *anticipates* and *adapts* to disasters, with the goal of both increasing the ability of communication networks to *absorb* and *adapt* to the impact of a disaster (as to "maintain an acceptable level of service" [17]), as well their own ability to quickly *recover* the network.

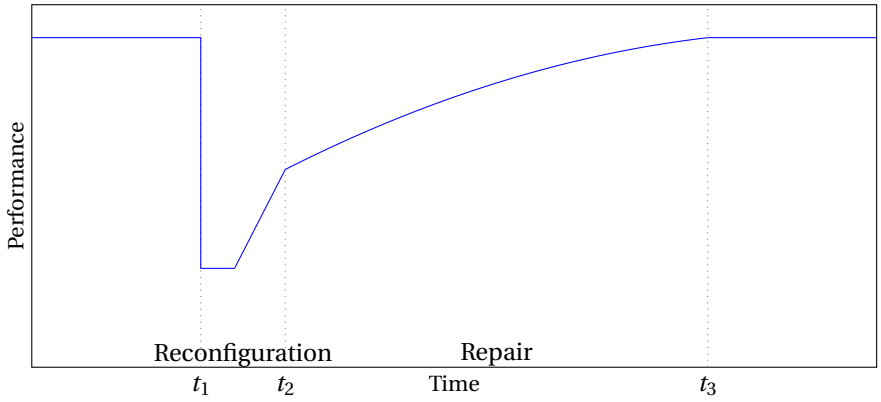


Figure 1.4: A simplified example of a resilience curve.

To measure and improve the resilience of a network, one needs to quantify the network's performance after a disaster. Fig. 1.4 shows a simplified resilience curve of the performance of a network after a natural disaster. At $t = t_1$, the network is struck by a disaster, causing an immediate drop in performance (absorb). Either manually or automatically, traffic is rerouted around the disaster area, thus allowing the network to regain some performance (adapt/recover). Finally, the much slower process of network repair can begin (recover). At $t = t_3$, the network is fully recovered. Hopefully, by learning from their response to the disaster, the network operator is able to increase the resilience of the network to the next disaster (adapt and anticipate).

1.2.1. METRIC

The selection of a metric is crucial for assessing and increasing disaster resilience. The choice of metric almost completely determines one's focus. For example, some authors choose to quantify the impact of a disaster by the number of damaged components. While this might be a good indicator of the monetary repair costs of a disaster, it is not a good resilience metric. The metric completely shifts the operator's focus to placing network components away from hazardous areas. This seems like a good idea, but it ignores the network's purpose of connecting end-users. In fact, it could punish the addition of redundant components to the network! Essentially, a resilience metric should reflect the requirements, or "mission", of the network. These requirements can differ from stakeholder to stakeholder. For example, a network operator might prioritize meeting its service level agreements, an individual business its own connection quality, and a government the ability to maintain communication with its citizens.

Since there is no one correct metric, the algorithms and approaches we propose in this thesis are designed to be applicable to almost any choice of resilience metric.

As a proof of concept, we demonstrate our approaches on the Average Two-Terminal Reliability (ATTR) metric. The ATTR is the ratio of the number of node pairs that remain connected after a regional failure, and the total number of node pairs in the original network. The ATTR of a connected network is 1, while the ATTR of a network without

any surviving links is 0. The ATTR is one of the more powerful metrics for measuring the ability of the network to maintain connections between different areas.

As an example, the number of connected node pairs in Fig. 1.3b is 4. The total number of node pairs in the original network (Fig. 1.3a) was $6 \times 5 = 30$. Thus, the ATTR of the network after this regional failure is $\frac{4}{30} \approx 0.133$[1]. The loss of a single datacenter disconnected almost all node pairs!

ATTR is named after two-terminal reliability [18]. The more general definition of ATTR, which incorporates component failure probabilities, is the average of the probabilities that pairs of nodes remain connected. If we only allow link failure probabilities, this is exactly the average of the two-terminal reliability of all node pairs.

1.3. THE COST OF RESILIENCE

Completely negating the impact of natural disasters on a network is not only nearly impossible, but also tremendously expensive. Any practical approach to raising the disaster resilience of a network is a trade-off between the cost of enhancing resilience and the impact this has on the resilience of the network. Of course, stakeholders can only make these trade-offs if they are well-informed on both the risk of disasters to the network, as well as the cost of increasing disaster resilience.

The difficulty here lies in that the number of choices a network operator can make to increase the resilience of the network, as well as the number of potential disasters that can strike the network, are large and diverse. In addition, data on the impact of these potential disasters on the network is quite sparse, as individual disasters are rare[2]. It is simply unfeasible to take all these potentialities into account manually. Thus, we are in need of algorithms that can help assess, as well as potentially advice on, the disaster resilience of a network.

Fig. 1.5 shows a seismic hazard map of the conterminous United States. Clearly, seismic hazard is not uniform. Some areas are at a much higher risk of earthquakes than other areas. The same is true for many other, if not all, types of natural disasters.

In general, different areas will experience disasters at different frequencies, and with different properties. If we want to support stakeholders in making well-informed decisions, it is crucial that our solutions take this into account. The only way to do so is to create algorithms that can act on currently available disaster data.

In the next sections, we give a detailed overview of past work on the disaster resilience of communication networks. We focus on how different approaches model disasters, and how well these approaches allow stakeholders to incorporate disaster data. Readers without interest in a broad overview of the field can safely skip to Section 1.3.3 (for a discussion of data-driven approaches) or Section 1.4 (for the problem statement).

1.3.1. PROTECTING AGAINST ALL DISASTERS

In 1991, Bienstock proposed a polynomial-time algorithm for a generalized variant of the min-cut problem for plane graphs, called the min-break problem [19]. Bienstock

[1] For comparison, the ATTR of a 6-node ring topology after a single node failure is $\frac{5*4}{30} = \frac{2}{3} \approx 0.667$

[2] In fact, a natural disaster is an example of a so-called High Impact, Low Frequency (HILF) event.

[3] https://www.usgs.gov/media/images/2018-long-term-national-seismic-hazard-map

1

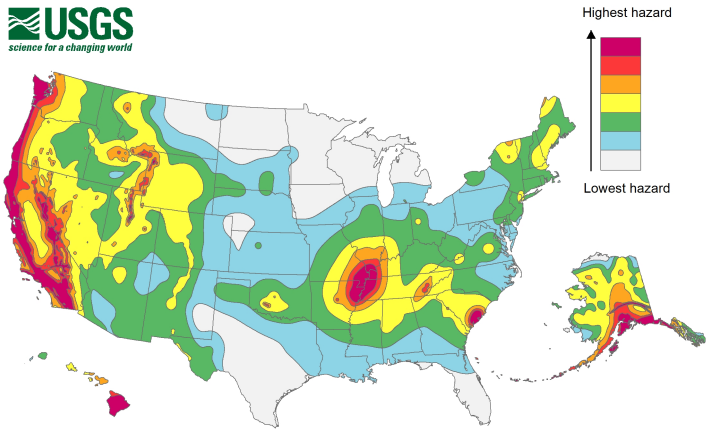


Figure 1.5: 2018 Long-term National Seismic Hazard Map of the conterminous United States[3].

modeled a regional failure as a “hole” - a subset of the plane homeomorphic to an open disk. Given a finite set of holes and nodes s and t, the min-break problem is to compute the minimum number of holes whose selection would disconnect s from t. Bienstock did not discuss how to obtain the list of holes. Years later, when research on the disaster-resilience of networks had begun in earnest, Neumayer et al. defined the Geographical Min-Cut By Circular Disasters Problem (GMCCD) [20]. Instead of only considering a finite set of potential holes, the GMCCD problem considers all possible disk failures of radius r (except those centered inside a protective disk around s and t). The goal is the same as the min-break problem: find a minimum cardinality set of disks that disconnect s from t. While the min-cut approach could have potential for assessing the risk of targeted attacks to a network, it is less applicable to natural disasters, since it essentially looks for the most impactful combination of regional failures.

A number of works consider the problem of finding the regional failure that does the most damage to the network [21–28] or drops the performance of the network below a specified level [29]. While suitable for assessing and preparing networks for targeted attacks (such as an Electromagnetic Pulse (EMP) attack), these approaches are less effective in assessing disaster resiliency. Effectively, they only convey information on a single, probably extremely rare, disaster scenario.

A Shared Risk Group (SRG) is a set of network components that may fail simultaneously. A Shared Risk Link Group (SRLG) is a SRG of links, and a Shared Risk Node Group (SRNG) a SRG of nodes. Note that, since we can capture any node failure as a combined failure of multiple links, SRLGs are essentially a generalization of SRNGs. By selecting a group of SRGs, and ensuring that primary and backup paths are not routed through the same SRG, operators can enhance the resilience of their network to the combined failure of components that, e.g., share common physical infrastructure.

The concept of SRGs has also been applied to disaster resilience; put simply, if we are able to model each potential disaster as an SRG, and *if* the operator can protect their network against all these SRGs, we have achieved full disaster resilience.

Tapolcai et al. defined a regional failure as any non-empty set of links, which can all be covered by a disk of radius r [30, 31]. They noted that protecting the network against all possible maximal regional failures is sufficient to protect the network against all possible regional failures[4]. Thus, they proposed a polynomial-time algorithm for finding all maximal regional failures, as well as a polynomial-time algorithm for finding the union of all maximal regional failures for all radii below r (leaving the exact selection of regional failures to the operator). Vass et al. also proposed over-estimating the impact of a regional failure by a disk, and gave an algorithm for finding all maximal link sets that can be hit by a disk hitting exactly k nodes [32]. They later extended their work to the case where the exact paths of each link are not known [33]. Recently, Vass et al. gave algorithms for finding maximal regional failures for networks on a sphere, instead of on the plane (as is commonly assumed) [34]. Depending on the radius, applying these techniques could lead to both over-protecting and under-protecting paths in different areas of the network. Furthermore, since no probabilities are assigned to each regional failure, stakeholders lack the necessary information to make well-informed trade-offs.

Cheng et al. took a similar approach [35]. They modeled a regional failure as a disk with fixed radius, and proposed scanning potential failure coordinates to find all possible combinations of node failures. Cheng et al. also considered a simplified probabilistic regional failure model, consisting of two disks with the same center. Nodes in the inner disk are assigned a failure probability of 1, and nodes in the outer circle a failure probability of 0.5. To simplify scanning, in this model, they only allow regional failures centered at network nodes. These regional failures could then be used to find vulnerable areas and assess the disaster resilience of the network.

Gardner et al. proposed using Geographic Multi-Topology Routing to quickly route traffic away from a regional failure [36]. The main idea behind their approach is that, in the event of a failure, the network switches to a virtual topology that routes traffic around the disaster region. To achieve this, they generate a set of topologies based on disk failures, where link weights are adjusted depending on the distance between the link and the center of the disk failure. Their algorithms can create a set of topologies by covering the network area by disk failures, or by being explicitly provided a set of disk centers. Although Gardner et al. assumed fixed disk sizes, they did demonstrate that the performance of their approach greatly depends on how well the actual regional failures match the model used to generate the virtual topology. This suggests that methods such as these could significantly benefit from the application of disaster data.

Zhang et al. proposed strengthening (or shielding) critical links against regional failures [37]. They gave an MILP formulation and heuristics for finding the minimum-cost combination of shielded links that ensures the network remains connected or partially connected under any SRLG. In addition, they also gave an algorithm and MILP formulations for ensuring connectivity between two nodes under any disk failure of fixed radius, or any SRLG. Allawi et al. proposed adding wireless backup connections to a subset of network links, and gave algorithms for selecting these links under a random regional

[4] In other words, the maximal regional failures can be used as Shared Risk Link Groups.

failure model [38]. Essentially, this is another variant of the shielding problem.

Habib et al. and Ju et al. proposed algorithms for placing content, and finding primary and backup paths, in datacenter networks [39, 40]. Given a set of SRGs, their approaches ensure each node remains connected to their requested content. Liu et al. discussed a similar problem, but proposed using cooperative storage instead of cloning data [41]. In their approach, content is encoded into several fragments. To recover the content, only a subset of fragments is required. Since a single fragment is much smaller than a whole clone of the content, this approach requires less storage space, while still protecting the network against any SRG.

Grebla et al. studied the distribution of replicas or fragments of a single file [42]. They proposed algorithms to place replicas (or fragments, when erasure coding is used) such that any node could still access the file after any potential regional failure. Regional failures were modeled as arbitrary disks, line segments, or a union of b disks and line segments.

Cai et al. proposed a disaster protection scheme for service function chain embedding using multi-path routing [43]. Their approach provisions primary and backup virtual network functions and paths such that the network remains functioning under any single SRG failure.

SRLGs and SRNGs can also be used to map virtual machines and their backup resources in such a way that the virtual machines and their connections are protected against any disaster of interest [44–46].

Banerjee et al. proposed algorithms for augmenting a network with additional links such that, after any disk failure of radius r, the size of the largest connected component [25] or the number of connected components [26] stays above a given level. Tapolcai et al. proposed algorithms for computing additional links, and the routes of these links, such that all surviving nodes remain connected after any disk failure of radius r [47].

A large number of approaches do not explicitly model regional failures, but consider the distance between network components instead [21, 35, 48–58]. The geodiversity between a pair of paths is the minimum distance between any intermediate node or edge of one path and any node or edge of the other path [49, 51, 57]. Clearly, any single regional failure with a diameter smaller than the geodiversity (that does not strike one of the end-nodes) cannot take out both paths simultaneously. Implicitly, these methods make the same assumptions as approaches that model regional failures as disk failures. However, they do allow for some flexibility by setting different geodiversity constraints per region [49, 52].

Girão-Silva et al. combined geodiversity with Shared Risk Link Groups (SRLG) [59]. They formulated integer linear programs for finding maximally SRLG-disjoint paths, as well as finding geodiverse maximally SRLG-disjoint paths.

1.3.2. RANDOM DISASTER MODELS

In contrast to the previous approaches, a random failure model allows stakeholders to assess the overall disaster resilience of the network, and to make informed trade-offs based on disaster risk. Neumayer and Modiano gave polynomial-time algorithms for computing the expected average two-terminal reliability and other metrics under random disk or line failures [60]. They model the network as a graph on the plane, and

assume a single random disk failure (of fixed radius) or line failure strikes the network. Disaster occurrence probabilities are assumed to be uniform over the network area. A number of other works take a similar approach by assuming uniform disaster probability and characteristics, and computing the expected value of a metric, or the probability that a metric exceeds some given value, after a random disaster [61–64].

An interesting, related approach, is to find the regions around the network that are vulnerable to a regional failure. Gardner and Beard proposed two methods to find the regions where a disk failure of a given radius would either disconnect a pre-selected pair of nodes, or disconnect any part of the network [65]. Note that, under the assumption of uniform occurrence probabilities, the total area of these regions is proportional to the probability that a random disk failure would disconnect part of the network. Similarly, Gardner et al. proposed an algorithm to find the regions where a disk failure of a given radius could cause the network to fail to perform its mission [66].

Assessing the resilience of the network to uniform disk failures can be a powerful tool, since it essentially measures if components are located too closely together. In fact, disks are one of, if not the, most popular models for regional failures. Computing distances is a cheap operation, and a disk can essentially be seen as an overestimate of any disaster shape with a smaller or equal diameter. That being said, when evaluating the disaster resilience of a network, assuming both uniform occurrence probabilities and disaster sizes can and will lead to severe over- or under-estimations of disaster impact in different areas of the network. As such, these approaches are best used in conjunction with a disaster assessment based on actual disaster data.

Saito proposed a regional failure model where a randomly placed line splits the network area in two [67]. Any network component intersecting the right-half plane formed by the line are assigned a failure probability. Saito also considered a more generic convex disaster area, constructed by randomly sampling a reference point and angle of two reference lines intersecting this point [68]. He obtained some theoretical results for tree, ring, and a combination of ring networks, and validated these results on a selection of historical earthquakes.

Rahnamay-Naeini et al. proposed a model for multiple correlated random regional failures [69]. In their model, the centers of regional failures are placed according to a Strauss point process. Link failure probabilities depend on the distances to these centers and are computed using Gaussian functions. The impact of random disasters following this model can be computed by running a Monte Carlo simulation. Das et al. extended the approach to include node and link failure probability functions based on Gaussian functions, lines and, circles [70]. Neither of these approaches take disaster data into account.

Cao et al. discussed a number of optimization problems with the goal of finding cable paths that minimize the disconnection probability of cities under a random disaster [71]. They modeled a disaster as a disk with uniform center distribution and an exponentially distributed radius.

A number of algorithms simply model a regional failure as a combination of component failures or failure probabilities, and take a list of such regional failures as input [72–80]. This bears a resemblance to SRGs, but also (either implicitly or explicitly) assigns an occurrence probability to each SRG. The problem of obtaining the list of regional fail-

ures and their related component failures is outside the scope of these approaches. In this thesis, we call the combination of component failures after a regional failure a *failure state*. Our approach, discussed in Chapter 3, computes lists of failures states as an intermediate step.

Another approach to assessing the resilience of a network is to select a specific disaster, and then perform an in-depth simulation or computation of the impact of this disaster on the network [81–84]. Given the variety and number of potential disasters that can strike a network (and the extremely low frequency of any individual disaster scenario), it is ill-advisable to solely assess network resilience on a single, or few, potential disaster scenarios. Thus, this approach is best used in conjunction with other methods, to further study a selection of interesting scenarios.

1.3.3. DATA-DRIVEN APPROACHES

By themselves, none of the previous approaches can assess or improve the disaster resilience of a network based on actual disaster data. In fact, many approaches explicitly ignore disaster data by either protecting the network against any disk failure, or assuming uniform disaster properties across the network area.

Ma et al. improved on the uniform random disaster model[5] by dividing the network region into a grid, and assigning an occurrence probability and intensity to each grid cell [85]. Using this model, the impact of a random disaster on the network can be estimated by generating one regional failure per grid cell, and computing its impact on the network (making the assumption that the impact of any other disaster centered on the same grid cell will be roughly the same). The disadvantage of this model is that it assumes all disasters occurring in each grid cell have the exact same properties. Furthermore, Ma et al. failed to demonstrate how to extract both occurrence probabilities and disaster intensities from actual disaster data.

Tran and Saito proposed two algorithms for enhancing the robustness of a network that *do* take actual disaster data into account [86, 87]. Their algorithms aim to optimize the weighted average of the sum of end-to-end disconnections of all node pairs under a set of earthquake scenarios, by either adding new links to the network, or changing the routes of existing links. They assume that, for each scenario, they are given a grid of ground motion intensities. The failure probability of each link is then computed by evaluating the length of the segments of the link passing through each grid cell and the intensities at these grid cells. Ground motion grids for both past earthquakes and earthquake scenarios are readily available [88, 89]. This detailed regional failure model does come at a cost, as the approach does not scale well to larger network sizes and disaster sets. In particular, computing the end-to-end disconnection probabilities in their evaluation metric is a well-known NP-hard problem for even a single disaster [18].

Eriksson et al. defined the bit-risk miles of a path between s and t as the sum of the geographic distance and the historical and forecast outage risk of each node on the path (weighted by the population weights of s and t, as well as some tuning parameters) [90]. Historical disasters are modeled as points, and a Gaussian kernel is used to convert these to historical outage risks. Forecast outage risks are extracted by parsing hurricane forecasts. This method incorporates the inhomogeneity of disaster risk by extracting disaster

[5] In particular, they extended the work of Wang et al. [62].

frequencies from actual disaster data. However, it does not consider that what makes a regional failure regional: the simultaneous failure of multiple network components. By summarizing disaster risk into bit-risk miles, we lose all information on the correlation between component failures.

Once a disaster strikes the network, it is imperative that traffic is rerouted as soon as possible. By adopting specialized protocols that react quickly to regional failures, one can increase the disaster resilience of a network by improving the ability of the network to adapt to regional failures [36, 55, 91–100]. By integrating these protocols with early warning systems, it is even possible to start rerouting traffic and migrating services before a disaster has struck the network [90, 93, 95, 100–103]. To enable a quick response, protocols like these respond with little to no dependence on human operators. Thus, we consider this process outside the scope of this thesis.

1.4. DATA-DRIVEN DISASTER RESILIENCE

1.4.1. PROBLEM STATEMENT

Overall, there is currently a lack of approaches that allow stakeholders to assess and improve the resilience of communication networks to natural disasters based on actual disaster data. Thus, the main problem statement of this thesis is

> *How to create scalable, data-driven methods for assessing and improving the resilience of communication networks to natural disasters.*

This problem statement can be split into two sub-problems: *assessing* disaster resilience, and *improving* disaster resilience.

The methods we propose in this thesis adhere to three important requirements. They are

- scalable to large disaster datasets;
- demonstrated on publicly available disaster data; and
- applicable to most resilience metrics.

Since the start of this project, several other regional failure models that could incorporate disaster data have been proposed. Interestingly, this research can be categorized into two opposite categories: (1) general approaches that aim to be applicable to almost any disaster, but have not been demonstrated on actual disaster data [104, 105]; and (2) specialized approaches that are only applicable to one type of disaster (or even only one dataset), but have been demonstrated on actual disaster data [106–110].

1.4.2. OUR APPROACH

The main principle behind our approaches - which is described in more detail in Chapter 3 - is to assess the disaster resilience of a network based on a large set of representative disaster scenarios. This relatively simple principle is surprisingly powerful. First, it automatically forces our algorithms to take into account the inhomogeneity of natural disasters. Second, it ensures our algorithms are applicable to currently existing disaster

data[6]; disaster scenarios can be created synthetically, crafted by hand, sampled from a distribution, and/or simply taken from a database of historical disasters.

The drawback of this approach is that the set of representative disasters should be large enough to capture the properties of all potential disasters the analyst wants to consider. This means representative disaster sets can grow very large (up to millions of disasters). Our approaches should reflect this, and should remain tractable for a large number of scenarios. Thus, scalability with the number of disaster scenarios is one of the main challenges for our algorithms.

As assessing disaster resilience is an essential step in effectively improving disaster resilience, the initial focus of this thesis is on resilience assessment. In Chapter 2, we study the risk of earthquakes to the Internet, and show that while a large number of Internet exchange points are threatened by earthquakes, spreading out IXPs and peering links over multiple facilities can significantly reduce the impact of these earthquakes. In Chapter 3, we introduce our framework for assessing the impact of disasters on a network. The basic assumption in these chapters is that the network will only be affected by one disaster at a time. In Chapter 4, we challenge this assumption. We show that the probability of two disasters striking a network in quick succession can be significant, and propose algorithms for assessing the probability of, and resilience to, these successive disasters.

Chapter 5 shifts the focus to resilience enhancement. In this chapter, we show that a targeted attack after a natural disaster can greatly increase the impact of this disaster on network performance. However, we also show that by selecting the right repair strategy, the impact of a potential targeted attack can be severely diminished, at almost no cost to the overall recovery of the network. We continue studying repair strategies in Chapter 6, where we propose a framework to evaluate different node replacement strategies. In Chapter 7, we consider network augmentation. We propose a set of algorithms that allow network operators to find cable routes that enhance disaster resilience in a cost-efficient manner. One of the strengths of these algorithms is the explainability of their solutions; for each potential cable, the algorithm can show which disaster areas the cable aims to avoid, and how much it values avoiding each of these areas.

Our approach is geared towards helping stakeholders anticipate disasters. In Chapters 2 to 5, we assess the ability of the network to absorb disasters. With a change of metric however, these approaches can also assess the ability of the network to adapt and recover from disasters. Chapter 6 aims towards helping operators recover their network more effectively, and the framework we propose in Chapter 7 advises operators on how to improve the ability of their network to absorb disasters.

[6] In fact, representative disaster sets are similar to the stochastic event sets used in catastrophe modeling [111].

2

A Global Study of the Risk of Earthquakes to IXPs

As a demonstration of data-driven resilience assessment, we first conduct an analysis of the resilience of key Internet infrastructure - namely IXPs - to earthquakes. We find that many facilities are at risk of earthquakes. More than 50% of the facilities have at least a 2% probability of experiencing potentially damaging levels of shaking within a period of 50 years. Furthermore, when an earthquake occurs, it is not unlikely that multiple facilities will fail simultaneously. We estimate that there is a 10% probability that at least 20 facilities will simultaneously experience potentially damaging levels of shaking within a period of 50 years. On the positive side, our analysis shows that Internet Exchange Points that host more Autonomous Systems tend to be located in less earthquake-prone areas, and that spreading Internet Exchange Points and peering links out over multiple facilities significantly reduces the impact of earthquakes to Internet Exchange Points and Autonomous Systems. Following this observation, we propose a novel metric and accompanying algorithm to help AS operators select peering facilities, based on the probability of simultaneous facility failures. We show that applying our metric can significantly increase the resilience of individual Autonomous Systems, as well as that of the Internet as a whole.

2.1. Introduction

As the Internet is a vital infrastructure, its resilience has been the focus of many studies. Surprisingly, studies to the resiliency of the Internet as a whole to rare, impactful events, such as natural disasters, are rare themselves. Events such as these can inflict significant, concentrated damage to Internet infrastructure, disrupting local (and sometimes global) connectivity just when people need it the most.

Many of the physical components and facilities making up the Internet are vulnerable to intense levels of shaking. In 2006, an earthquake of the coast of Taiwan damaged 8 submarine cable systems, severely disrupting communications in the region[113]. The

Parts of this chapter have been published in IFIP Networking Conference, 2022 [112].

2011 earthquake and subsequent tsunami of the coast of Japan caused extensive damage to telecommunications buildings and equipment, leading to widespread connectivity problems. The total cost of emergency restoration and reconstruction of the local NTT East network was 80 billion yen (around 1 billion dollars at the time) [14, 114]. In 2015, Nepal was struck by a devastating earthquake, which damaged both cell towers and back-haul infrastructure [15]. The total damage to the telecom industry was estimated at 17.4 million dollars. The damage caused by the 2016 Kaikōura earthquake in New Zealand led to local telecom outages of up to 5 days [115].

In this chapter, we aim to take a global look at the risk of earthquakes to key Internet infrastructure. Our main focus is on Internet Exchange Points (IXPs). An IXP is a physical infrastructure used by Autonomous Systems (ASes) to directly exchange traffic between their networks. Besides potentially reducing costs (by reducing the amount of traffic delivered via transit providers), IXPs have been shown to increase quality of service [116].

Given their critical role in the Internet and the presence of multiple ASes at each of their facilities, the destruction of IXP facilities could have severe consequences for the Internet as a whole. IXPs and network operators do take resiliency measures, such as distributing their services over multiple facilities, and/or rerouting traffic through other IXPs and ASes, in case of failures. But the loss of an IXP facility would certainly cause both temporary issues, as well as reduce quality of service in the local area.

Using publicly available earthquake models (covering 68.9% of our IXP dataset) and hazard computation tools, we estimate both the hazard to individual IXP facilities, as well as the probability of simultaneous facility failures. Our main conclusions are as follows:

- Many IXP facilities are at risk of potentially damaging levels of shaking: 32.4% of facilities have at least a 10% probability of experiencing potentially damaging levels of shaking within 50 years, and 50.9% of facilities at least a 2% probability.

- A minority of facilities host far more ASes than most other facilities. We find that while a number of these facilities are still at risk of earthquakes, overall these more important facilities tend to be located in less earthquake-prone areas.

- There is a real possibility of simultaneous facility failures: In 50 years, there is a 10% probability that at least 20 facilities will experience potentially damaging levels of shaking due to a single earthquake, and a more than 6% probability of 20 IXPs simultaneously experiencing this level of shaking.

- Distributing IXPs over multiple facilities helps. The median probability that an IXP with multiple facilities will simultaneously experience potentially damaging levels of shaking at all its facilities is well below 1%.

To help operators increase the resilience of their ASes to earthquakes, we propose a new metric and algorithm for selecting IXP facilities, based on the probability of simultaneous facility failures. We show that applying our metric can significantly increase the resilience of individual ASes, as well as that of the Internet as a whole.

2.2. RELATED WORK

There have been numerous studies on how to assess the risk of earthquakes and other natural disasters to *single* communication networks. The topic has been addressed by both network scientists (e.g. [62, 105, 117]) and seismologists [110]. Recently, Valentini et al. undertook a multi-disciplinary study, combining insights from both network science and seismology to create a custom-built model for assessing the resilience of Italian communication networks to earthquakes [109]. While crucial to our understanding of disaster risk, such studies focus on the resilience of single communication networks, and their methods and results do not necessarily scale well to the Internet as a whole.

The Internet is an incredibly complex network of millions of devices. A combination of a lack of complete data, as well as its scale, makes undertaking any risk assessment a difficult endeavor. There are few studies on the resilience of the Internet as a whole to disasters. Jyothi studied the risk of solar storms to the Internet by considering a number of risk factors (such as the geographical spread of ASes and datacenters), as well as by looking at the impact of random cable failures (where the probability of failure depends on cable length and latitude) [118]. The author concluded that solar storms have the potential to massively disrupt the Internet.

Anderson et al. analyzed the risk of wildfires to cellular infrastructure in the United States, by studying which cellular transceivers are under threat from wildfires [119]. They computed the number of transceivers that lied within wildfire perimeters in past years, as well as the number of transceivers that lie within higher-threat areas (as identified by the United States Forest Service).

Eriksson et al. proposed RiskRoute, a routing framework that can configure routes based on both historical and forecasted outage threats [90]. The goal of RiskRoute is to minimize bit-risk miles: the sum of geographic distance and expected outage risk encountered along a routing path. The authors also proposed a method for selecting new peering links, with the goal of minimizing overall bit-risk miles. RiskRoute, and this method, do not consider simultaneous component failures, but only consider the individual risk to each point of presence in isolation.

Durairajan et al. and Mayer et al. used data from the Internet Atlas [120] to analyze the risk of global warming [121] and earthquakes [122] to Internet infrastructure in the United States. Both of these works essentially analyze the risk to Internet infrastructure by determining the amount of infrastructure at risk. Mayer et al. do not consider how many, and which, network components could be struck by any individual earthquake.

The true danger of an earthquake to communication networks is not only the damage it can inflict to any individual point of presence, but also its ability to disrupt multiple points of presence at once. Any approach that only considers the risk to individual network components in isolation essentially only paints half the picture. For a more thorough analysis of the risk of earthquakes to the Internet, we need to consider which components may be disrupted simultaneously, and with what probability. This requires a more complex approach that considers individual earthquake scenarios.

While some studies have assessed the impact of disaster scenarios on a smaller scale, to the best of our knowledge, we are the first to analyze the risk of natural disasters to the Internet using a large number of realistic disaster scenarios generated based on actual disaster data, as well as the first to assess the risk of earthquakes to Internet infrastruc-

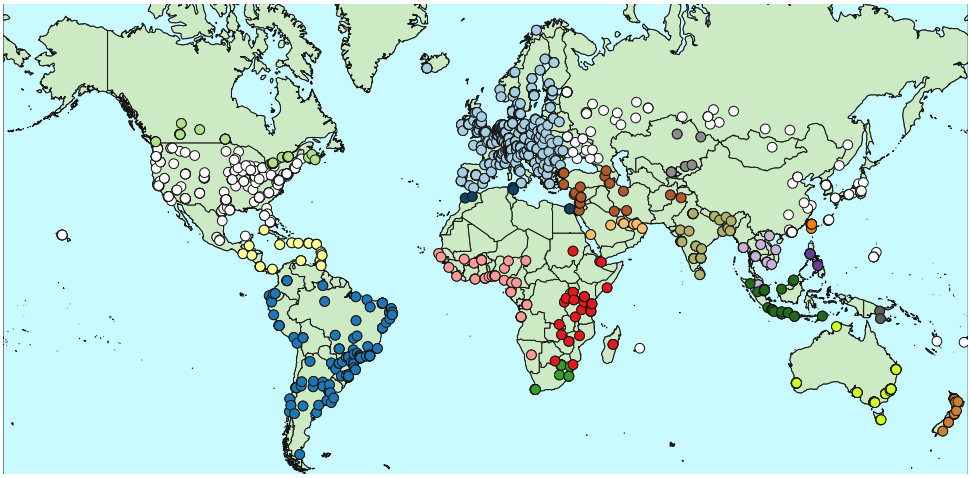


Figure 2.1: Assignment of facility locations to hazard models. All unassigned locations are colored white.

ture globally. We combine a set of 19 earthquake hazard models covering approximately 68.9% of global IXP facilities and generate a total of 902,134,602 earthquake scenarios to estimate the risk to individual facilities, as well as the risk of earthquakes to the Internet as a whole.

2.3. DATASETS

2.3.1. IXPS

We use the CAIDA Internet eXchange Points (IXPs) Dataset [123] as our source of IXP data. This dataset has been constructed by combining information from PeeringDB, Hurricane Electric, and Packet Clearing House. The dataset gives the locations of IXPs (from all sources), the locations of facilities (i.e., datacenters) hosting these IXPs (only from PeeringDB), and the autonomous systems (ASes) peering at each IXP. A single facility can host multiple IXPs, and an IXP can be distributed over multiple facilities.

Our study will be on the level of individual facilities. Thus, as a first step, we create a singly facility for each IXP without assigned facilities. We place these facilities at the location of the IXP itself. IXPs without location information (country + city or lon+lat) are filtered out. We also filter out all facilities that do not host an IXP. The resulting dataset contains 1,887 facilities, hosting a total of 1,162 IXPs.

For our analysis, we need the longitude and latitude of each facility. Most facilities are already assigned precise geographical locations. For the 220 facilities missing coordinates, we assign the coordinates of their city, as given by Geonames[1]. Some facilities were assigned incorrect coordinates by PeeringDB, placing them in the middle of the ocean. In addition, there was a mismatch between the assigned city and country of some IXPs (e.g., Haarlem located in Romania instead of the Netherlands). We manually corrected the locations of these facilities.

[1] http://www.geonames.org/

2.3.2. OPENQUAKE ENGINE

We use the OpenQuake Engine [124] to estimate the earthquake hazard at each facility. The OpenQuake Engine is an open-source software tool for earthquake hazard and risk calculation. One of the key benefits of the OpenQuake Engine is the availability of hazard data for most of the world. This allows us to use largely the same process to determine earthquake risk, independent of the location of a facility.

To calculate earthquake hazards, the OpenQuake engine needs both a seismic source system and a ground motion system. In the remainder of this chapter, we will refer to the combination of seismic source system and ground motion system as a *hazard model*.

SEISMIC SOURCE SYSTEM

The seismic source system consists of multiple seismic source input models, and a seismic source logic tree. A seismic source input model is a list of seismic sources. Each source (e.g., a fault) can generate earthquakes, and the model describes the location, timing, and other properties of these earthquakes. The OpenQuake engine assumes that earthquake occurrences follow a Poisson distribution. The only exception to this are non-parametric sources, which require the user to explicitly provide a set of earthquake ruptures and their occurrence probabilities.

To incorporate epistemic uncertainties, the OpenQuake engine allows users to define multiple seismic source input models, as well as parameter assignments for these models. The potential combinations of input models and parameters are given as a seismic source *logic tree*. Each branch of the logic tree essentially gives a possible modeling choice, and the weight assigned to this choice.

GROUND MOTION SYSTEM

Defining the seismic source system is not sufficient. To compute the potential shaking intensity at each location, we also need one or more ground-motion prediction equations. These allow the engine to compute the expected shaking intensity at each location of interest, for each potential earthquake. As with the seismic source system, epistemic uncertainties are defined in a ground motion logic tree.

Table 2.1: Hazard models used in our calculations.

Model	Version	IXP Facilities
2013 Euro-Mediterranean Seismic Hazard Model (ESHM13) [125]	6.1	629
Hazard Model for South America [126]	2016.0.0	158
2018 National Seismic Hazard Assessment for Australia [127]	2018.032	81
Indian Subcontinent PSHA [128]	2.0.1	61
Hazard Model for Southeast Asia (2018) [129]	2018.0.1	59
Hazard Model for Canada (2015) [130]	2015.1.1	44
Hazard Model for Indonesia [131]	2017.0.0	41
Hazard Model for Western Africa [132]	2018.0.0	40
2014 Earthquake Model of the Middle East (EMME14) [133]	1.5.0-2016-10-31	35
Hazard Model for Eastern Sub-Saharan Africa (2018) [134]	2018.0.0	33
Hazard Model for the Caribbean and Central America [135]	2018.0.0	31
New Zealand 2010 National Seismic Hazard Model[2]	04	27
Hazard Model for South Africa [136]	2018.0.1	14
EMCA Central Asia seismic source model [137]	1.1	10
Hazard Model for the Philippines (2018) [138]	2018.1.1	10
Hazard Model for the Arabian Peninsula [139]	2018.0.0	9
Hazard Model for Taiwan [140]	2015.0.0	9
Hazard Model for Northern Africa (2018) [141]	2018.0.0	7
Papua New Guinea Seismic Hazard Assessment [142]	NSHA_2019	3

HAZARD MODELS

To attain global coverage, we need to combine results from multiple hazard models (see Table 2.1). Versions of these models were also used by the Global Earthquake Model (GEM) foundation to create their Global Seismic Hazard Map[3]. We only make use of publicly available models that are not under any non-disclosure agreement. While this does leave some gaps in our coverage, we are still able to estimate the hazard to 1,301 out of 1,887 facilities (68.9%).

Based on their location, we assign each facility to a single hazard model (see Fig. 2.1). Not unexpectedly, the European ESHM13 model is assigned the largest number of facilities (629, or 48.3% of assigned facilities). The largest gap in our coverage is the United States of America: 394, or 67.2%, of the unassigned facilities are located in the United States.

The public datasets for New Zealand and Central Asia only contained seismic source input models. For these regions, we used the ground motion system specified by GEM[4] instead.

2.4. FACILITIES AT RISK

2.4.1. METHODS

We first study the hazard of earthquakes to individual IXP facilities. Whereas the strength of an earthquake is typically indicated by its magnitude, we are instead interested in the intensity of shaking at each facility. There are a number of intensity measures used to measure the intensity of shaking, each with their own advantages and disadvantages. One of the more common intensity measures in use today is Peak Ground Acceleration (PGA).

[2] The Earthquake Rates – National Seismic Hazard Model is owned by GNS Science and is based on the model explained in [143]. The model is held under licence from GNS Science.

[3] https://www.globalquakemodel.org/gem-maps/global-earthquake-hazard-map

[4] https://hazard.openquake.org/gem/models

As the name implies, PGA measures the peak acceleration of the ground during an earthquake. It is seen as a good indicator of earthquake hazard for short buildings (of up to 7 floors) [144]. We have chosen to focus on PGA in this study as it is one of the more intuitive intensity measures, and because we assume most IXP facilities are located in short buildings.

We are interested in (1) the level of shaking we can expect in a given investigation period and (2) how often we can expect potentially damaging levels of shaking at each facility. Both of these objectives can be achieved by running a classical Probabilistic Seismic Hazard Analysis (PSHA). Simply stated, a classical PSHA considers all potential earthquake ruptures together with the ground motion prediction equations, to compute a hazard curve for each location [145, 146]. A hazard curve gives the probability of exceeding given levels of shaking at a location (or *site*) within a specified investigation time. The hazard curves can be reduced to a hazard map, which shows the level of shaking with a given probability of exceedance (e.g. the PGA with a 2% probability of exceedance) at each site.

As discussed in Section 2.3.2, OpenQuake incorporates epistemic uncertainties within a logic tree. Each path through this logic tree (called a realization in the OpenQuake Engine) constitutes a different combination of ground motion prediction equations and source model. This means that instead of computing a single hazard curve for each site, the engine needs to compute a hazard curve for each realization. Thus, when we mention a probability of exceedance within this chapter, we are actually referring to a *mean* probability of exceedance of all hazard curves.

POTENTIALLY DAMAGING LEVELS OF SHAKING

PGA is an objective measure of the shaking, or ground-motion, due to an earthquake. While this correlates with damage, it is not a direct measure of the damage to buildings and infrastructure. In contrast, a macroseismic intensity scale, such as the Modified Mercalli Intensity scale (MMI), measures the observable (but more subjective) effects of an earthquake. In some papers and hazard maps (e.g., [109, 122, 147]), a macroseismic intensity of 6 (in MMI or the Mercalli-Cancani-Sieberg (MCS) scale) is used as a sort of lower-bound for potentially damaging levels of shaking[5].

It is not straightforward to convert PGA to a macroseismic intensity. The level of damage to a building depends on a variety of factors including construction materials, building codes, and the number of floors. Thus, there are inherent regional differences in the relationship between ground motion and macroseismic intensity. Caprio et al. quantified some of these regional differences, and constructed global ground motion to intensity conversion equations (for a combined MMI/MCS intensity scale) [148]. While one would preferably use regional conversion equations, the global scope of our study makes the global conversion equations a practical, albeit imperfect, alternative.

[5] Note that the building itself does not need to be damaged to disrupt an IXP facility. A facility could also be disrupted if equipment inside the building is damaged or falls down, or if infrastructure in the surrounding area is damaged.

A global macroseismic intensity of 6 roughly corresponds to a PGA of 0.086g. Thus, we will use a PGA of 0.086g as a threshold for potentially damaging levels of shaking. In comparison, using conversion equations for California [149] would result in a threshold of 0.11g (or 0.084g if we round up from an intensity of 5.5), and Mayer et al. assumed infrastructure is potentially damaged if the PGA exceeds 0.092g [122].

CALCULATION SETUP

We run a classical PSHA with an investigation time of 50 years on each hazard model[6]. We compute the probability of exceeding a PGA of 0.086g in 50 years, as well as the PGA with a probability of 10% and 2% of being exceeded in 50 years. The 10% and 2% probabilities of exceedance in 50 years are two common choices for seismic hazard maps. They correspond to a return period of 475 and 2,475 years, respectively.

Configuration To run a PSHA in OpenQuake, we need to provide a configuration file. The configuration file points to the files containing the logic trees, specifies the sites (and importantly, the attributes of these sites), and gives the calculation parameters. The hazard models for Europe, Australia, the Indian subcontinent, the Middle East, and Papua New Guinea include initial configuration files. For these models, we kept most calculation and site parameters, and only changed the sites, investigation time, and the intensity measure type and levels.

For all other hazard models, we set the site attributes to the same reference values used in the ESHM13 (corresponding to a reference rock condition matching Eurocode 8[7] Type A). This closely matches the choice of site attributes used by GEM to construct the Global Earthquake Hazard Map [150], which opted to choose attributes that represent "rock conditions according to the large majority of classification schemes in building codes and normatives".

Most of the calculation parameters are essentially a trade-of between precision and calculation time. We set `rupture_mesh_spacing` to 5[8], `width_of_mfd_bin` to 0.1, and `area_source_discretization` to 10. One can reduce computation times further by setting `complex_fault_mesh_spacing` and `pointsource_distance`. However, this was not necessary for our calculations, as the number of sites we consider is relatively low.

One of the more important parameters is the maximum distance between ruptures and sites at which the OpenQuake engine still considers the rupture when computing the hazard at a site. By lowering this setting, one can reduce computation times by excluding ruptures that are far away and would not significantly impact the computed hazard. We set this distance to an, in our eyes, conservative level of 800km[9]. For Canada, we indicate a maximum distance per tectonic region type, as described in [151].

[6] With the exception of the Caribbean and Central America and the Philippines. These models contain non-parametric seismic sources, and are fixed at an investigation time of 1 year. We convert their results to a 50 year investigation time by assuming Poissonian occurrences.

[7] https://eurocodes.jrc.ec.europa.eu/showpage.php?id=138

[8] For Taiwan, we lowered this value to 1, as the model contains seismic sources with low magnitudes that can not be represented properly with a mesh spacing of 5

[9] For comparison, the maximum distances are set to 200km; 400km and 1000km; 200km; 150km; and 200km and 600km for Europe; Australia; the Indian subcontinent; the Middle East; and Papua New Guinea respectively.

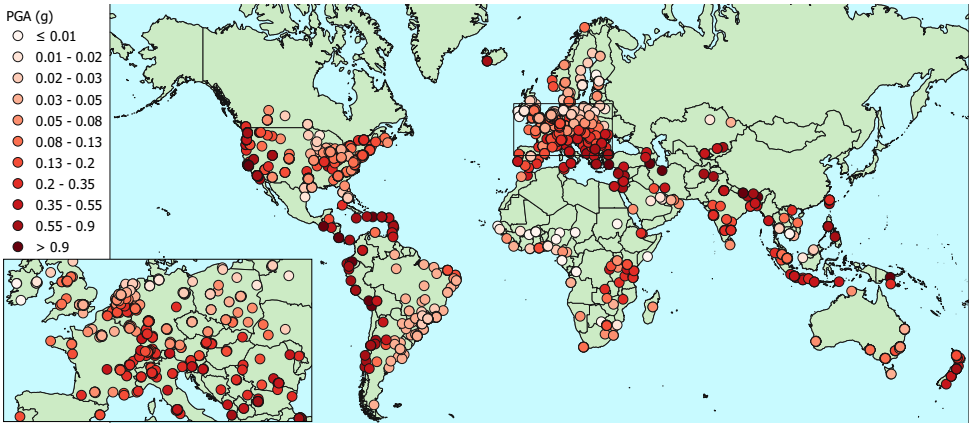


Figure 2.2: Locations of facilities covered by a hazard model from Table 2.1, and the local PGA with a 2% probability of being exceeded in 50 years. Results for the conterminous US were added by extracting PGA values from the 2018 USGS long-term seismic hazard map [147].

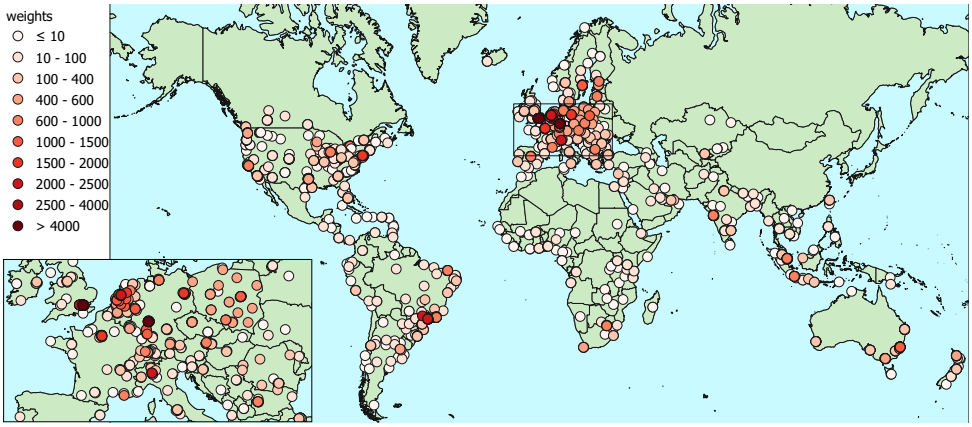


Figure 2.3: Locations of facilities covered by a hazard model from Table 2.1 and facilities in the conterminous United States, and the total weight of these facilities.

To prevent very high, potentially unrealistic estimates of the level of shaking, the tail-end of the ground motion distribution is typically cut off. In the OpenQuake engine, this can be done by setting a truncation level. In our calculations, we use a truncation level of 3.

2.4.2. RESULTS

Fig. 2.2 shows the PGA with a 2% probability of exceedance of each unique location of the facilities covered by one of the hazard models, as well as those in the conterminous United States. In this section, we discuss the facilities covered by the hazard models. For a more complete analysis, we briefly discuss the hazard of US facilities in Section 2.4.4.

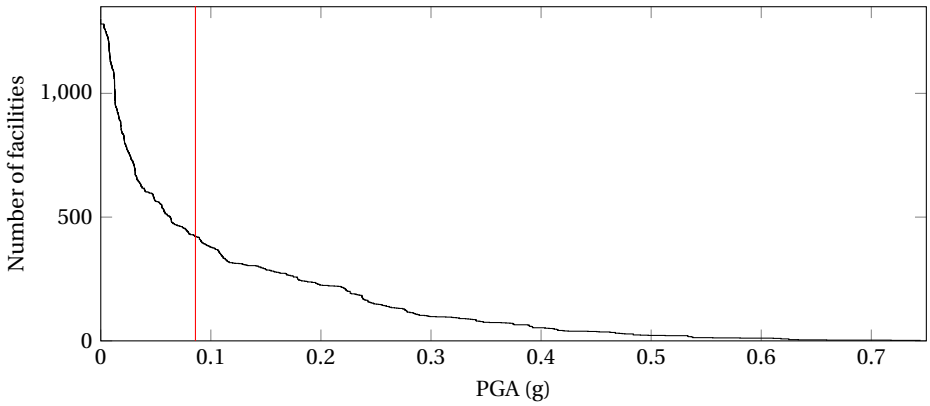


Figure 2.4: The number of facilities with at least a 10% probability of exceeding a given PGA in 50 years. The red line indicates our threshold of potentially damaging levels of shaking.

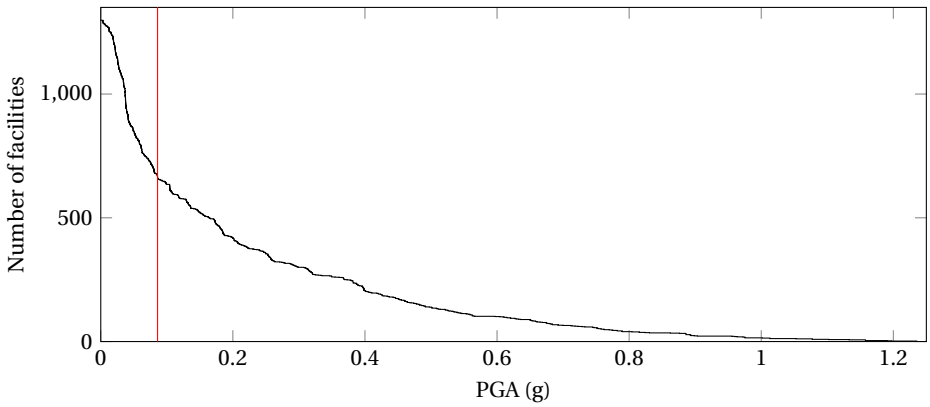


Figure 2.5: The number of facilities with at least a 2% probability of exceeding a given PGA in 50 years. The red line indicates our threshold of potentially damaging levels of shaking.

The hazard models cover a total of 1,301 facilities spread out over 1,135 unique locations. Together, these facilities host 849 unique IXPs. The facilities with the highest shaking hazard are located in Quito, Ecuador. These facilities have a 10% probability of exceeding a PGA of 0.744g and a 2% probability of exceeding a PGA of 1.24g (approximately equivalent to a macroseismic intensity of 9 and 10, respectively).

Figures 2.4 and 2.5 show the PGA versus the number of facilities with at least a 10% (respectively 2%) probability of exceeding this PGA in 50 years. A surprising number of facilities are at risk of potentially damaging levels of shaking. While the median PGA with a 10% probability of exceedance is only 0.0333g, the median PGA with a 2% probability of exceedance is 0.0928g - just above our threshold of 0.086g. Overall, 422 (32.4%) facilities have at least a 10% probability of experiencing potentially damaging levels of shaking within a period of 50 years and 662 (50.9%) facilities at least a 2% probability.

Probability of Exceedance	facilities
≤ 0.01	496
0.01 - 0.02	143
0.02 - 0.1	240
0.1 - 0.2	116
0.2 - 0.5	113
0.5 - 0.8	130
0.8 - 1	63

Table 2.2: The number of facilities with given probabilities of exceeding potentially damaging levels of shaking (PGA of 0.086g) within a period of 50 years.

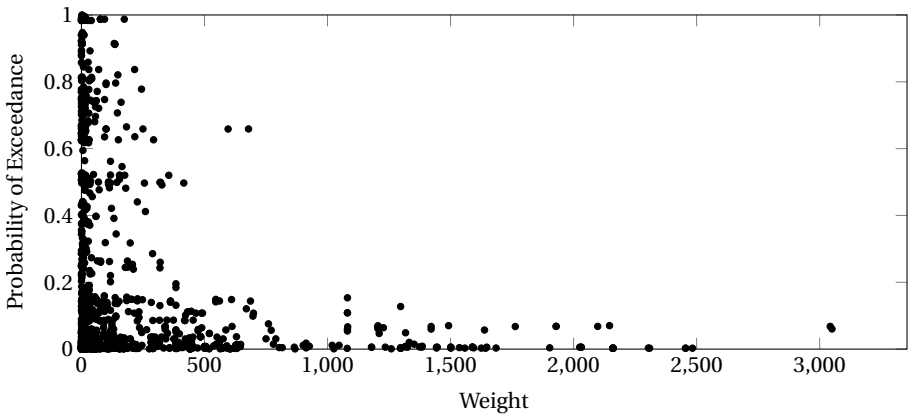


Figure 2.6: Weight of each facility versus the probability of exceeding potentially damaging levels of shaking within a period of 50 years.

Table 2.2 shows the number of facilities with given probabilities of exceeding potentially damaging levels of shaking within a period of 50 years. Interestingly, Quito is not the location with the highest probability of experiencing potentially damaging levels of shaking. The facility with the highest probability, and which thus is most often expected to experience this level of shaking, is located in Changhua, Taiwan. The probability that this facility experiences a PGA of at least 0.086g in a period of 50 years is almost 100%!

Of course, not every facility is equally important. To measure the importance of each IXP, we count the number of ASes at each IXP. Although the dataset does not contain all ASes that peer at every IXP, we expect this number to be proportional to the real number of ASes at an IXP and thus a reasonable measure of its importance. We set the weight of each facility to the sum of the number of ASes of each of the IXPs it hosts, and set the weight of each location to the sum of the weights of all facilities at this location. The results have been plotted in Fig. 2.3.

Most IXPs only host a few ASes: the median number of ASes at an IXP is 11, and there are only 138 IXPs (out of 1162) with at least 100 ASes. The disruption of these larger IXPs would impact the Internet much more than that of other IXPs.

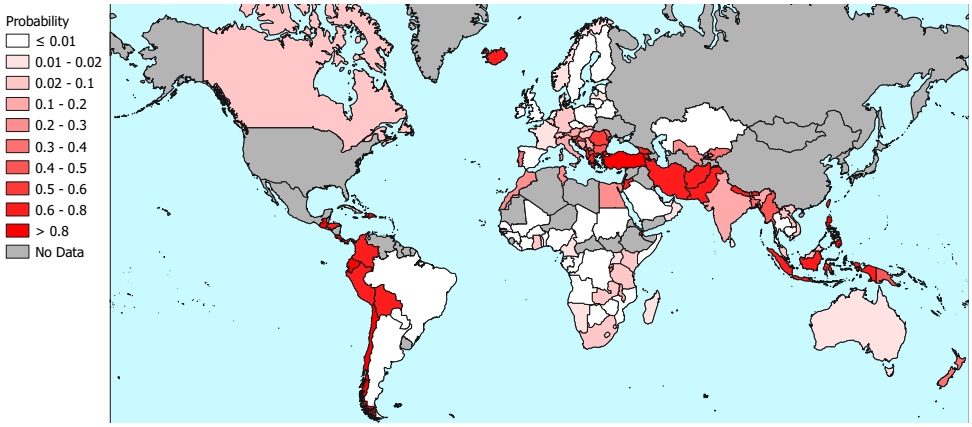


Figure 2.7: The median probability of exceeding potentially damaging levels of shaking within a period of 50 years at each facility of every country. Countries with either (1) no facilities or (2) facilities that were not assigned to a hazard model (see Fig. 2.1) are excluded (No Data).

Fig. 2.6 shows the weight and probability of exceeding potentially damaging levels of shaking of each facility covered by one of the hazard models. Overall, facilities with a larger weight have a lower probability of experiencing damaging levels of shaking: the median probability of exceeding potentially damaging levels of shaking within a period of 50 years is respectively 0.0355, 0.0134, and 0.00576 for facilities with a weight below 100, at least 100, and at least 1,000. However, there are a number of high-weight facilities in higher-risk areas: there are 215 (of 502) facilities with a weight of at least 100 that have at least a 2% probability of exceeding potentially damaging levels of shaking, and 26 (of 84) facilities with a weight of at least 1,000 that have at least a 2% probability of exceeding potentially damaging levels of shaking.

2.4.3. COUNTRY-LEVEL ANALYSIS

In this section, we analyze the risk of earthquakes to IXP facilities on a country-level by mapping each facility to the region denoted by its ISO 3166 two-letter country code [152].

Fig. 2.7 shows the median probability of exceeding potentially damaging levels of shaking for each country. This essentially shows the earthquake hazard that an average facility in each country faces. These values are affected by both the frequency and intensity of earthquakes in each country, as well as the exact placement of facilities within the country.

The median probability does not give the full picture, and even hides the influence of any outliers within a country. Risk is a combination of probability and impact. Thus, what we are more interested in is the number of facilities that could be disrupted by earthquakes in each country.

Fig. 2.8 shows the sum of exceedance probabilities of each country. Or, in other words, the expected number of facilities in each country that will experience potentially damaging levels of shaking at least once within a period of 50 years. We can see that the risk in countries with a low median probability of exceeding potentially damaging levels

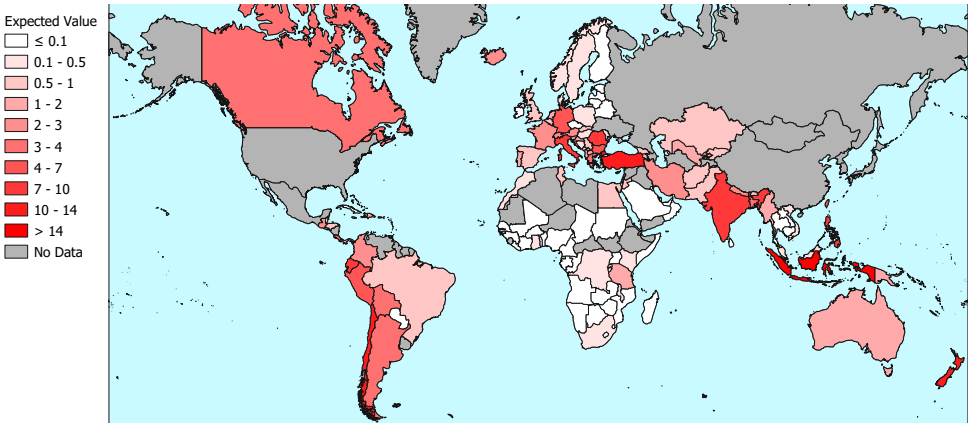


Figure 2.8: The expected number of facilities that will experience potentially damaging levels of shaking within a period of 50 years in each country. Countries with either (1) no facilities or (2) facilities that were not assigned to a hazard model (see Fig. 2.1) are excluded (No Data).

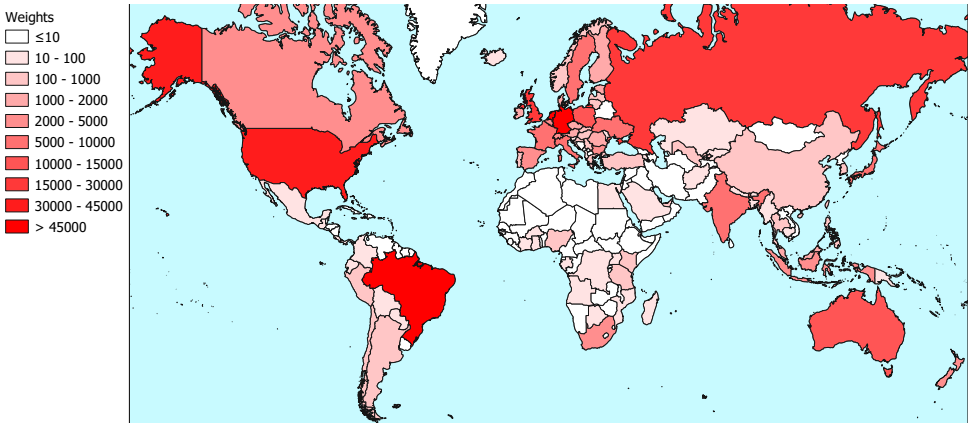


Figure 2.9: The total weight of facilities in each country.

of shaking can still be relatively high, simply due to the number of facilities. Similarly, some countries with a high median probability of exceeding potentially damaging levels of shaking are at lower risk than one might expect, because they do not host many facilities.

Indonesia is both prone to large earthquakes, and hosts a reasonably high number of IXP facilities (38). As such, it is no surprise that it is the country with the highest expected number of facilities that will experience potentially damaging levels of shaking (20.8). Out of all countries covered by our hazard models, Germany hosts the most IXP facilities (101). While it is not the most earthquake-prone country we have studied, it still ranks as the country with the 14th highest expected number of facilities that will experience potentially damaging levels of shaking (4.16).

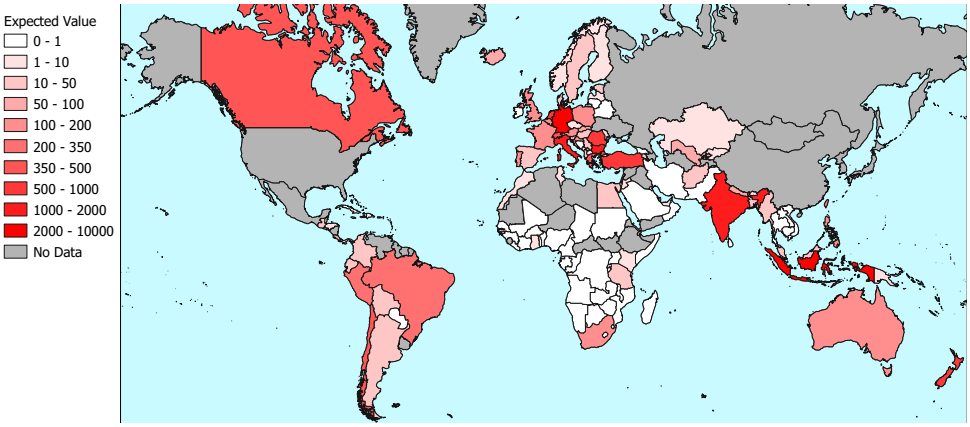


Figure 2.10: The expected total weight of facilities that will experience potentially damaging levels of shaking within a period of 50 years in each country. Countries with either (1) no facilities or (2) facilities that were not assigned to a hazard model (see Fig. 2.1) are excluded (No Data).

As we discussed in the previous section, not every facility is equally important; most facilities host less than 12 ASes, while a small minority host more than 1,000 ASes. Fig. 2.9 shows the sum of the weights of all facilities in each country. Interestingly, if we compare this figure to Fig. 2.7, we can see that the facilities in countries with large total weight (such as the Netherlands and Brazil) tend to have relatively lower probabilities of experiencing potentially damaging levels of shaking.

Fig. 2.10 shows the sum of the product of weight and exceedance probability of each facility for each country. That is, the total expected weight of the facilities in each country that will experience potentially damaging levels of shaking at least once within a period of 50 years. According to this metric, Indonesia only ranks as the second-most country at risk (with an expected weight of 2,021). Due to its concentration of high-weight facilities, Germany has the highest total expected weight of facilities that will experience potentially damaging levels of shaking (2,843[10]). Clearly, different weight functions may lead to a different ranking.

2.4.4. CONTERMINOUS UNITED STATES

Out of the 1,887 facilities in the dataset, 390 are located in the conterminous United States. The US has the 4th highest total weight of all countries. Although we lack a hazard model for the United States, we would be remiss if we completely ignore it. In this section, we give a brief analysis of the seismic hazard to IXPs in the United States based on the 2018 USGS long-term seismic hazard maps [147].

The hazard maps give hazard data for a grid of points spread out over the conterminous United States. To determine the hazard for each facility, we simply map it to its closest grid point. We first extract the PGA with a 2% probability of exceedance (see Fig. 2.2) for site class B/C[11]. It seems the average hazard at US facilities is only slightly

[10] Note that this is only 5.72% of Germany's total weight of 49,710.
[11] Roughly equivalent to the site class used for our own calculations.

higher than that of the rest of the world; the median PGA with a 2% probability of exceedance is 0.103g, compared to 0.0928g in the rest of the world.

The USGS includes a map of the chance of "slight (or greater) damaging earthquake shaking in 100 years" (i.e. the probability of MMI of 6 or higher), which can be easily converted to 50 year probabilities. Note that, while very similar, these probabilities were computed in a different manner than our probability of experiencing potentially damaging levels of shaking, and thus are not perfectly comparable.

As in the rest of the world, a large number of US facilities are at risk of earthquakes; the median probability of experiencing slight (or greater) damaging earthquake shaking in 50 year is 0.0312. Out of the 390 facilities, 89 (22,8%) have at least a 10% probability of experiencing damaging earthquake shaking, and 279 (71,5%) at least a 2% probability.

The expected number of US facilities that will experience damaging earthquake shaking in 50 year is 68.9. While this is indeed more than any other country we analyzed, the United States also contains by far the most facilities of all countries. As before, we assign a weight to each facility equivalent to the sum of the number of ASes at each of the IXPs it hosts. The total weight of facilities in the United States is 40,918 (4th in the world). The expected total weight of US facilities that will experience damaging earthquake shaking in 50 year is 11,384 (much more than any other country!). In the United States, more than in the rest of the world, a large number of facilities with relatively high number of ASes is at high risk of damaging earthquakes.

2.5. COMBINED FAILURES

Depending on its importance, the outage of a single IXP facility could have significant impact. However, the Internet was designed to be resilient, and should be able to reroute traffic around a failed facility. Even a single IXP is often spread out over multiple facilities, allowing their clients to increase redundancy by peering at multiple different locations. The real danger of natural disasters lies in their potential to disrupt multiple facilities simultaneously.

Whereas in the previous section we considered facilities individually, in this section we study the risk of simultaneous facility outages. In other words, we study the potential disruption of multiple IXP facilities due to a single earthquake. To this end, we first run an event-based PSHA in OpenQuake. In contrast to a classical PSHA, an event-based PSHA randomly generates sets of earthquake events, called stochastic event sets, as well as ground motions at each site during each of these events. A single stochastic event set is a realisation of potential earthquakes during the full duration of the investigation time. By generating multiple event sets, and processing the resulting ground motion fields, we can determine which facilities could potentially be disrupted simultaneously (and with which probability).

2.5.1. DISRUPTION

In this section, we say a facility is disrupted by an earthquake if it experiences shaking with a PGA of at least 0.086g. In addition, we say an IXP is disrupted if at least one of its facilities is disrupted, and is *fully* disrupted if all of its facilities are disrupted. Since our threshold of 0.086g is a lower bound on potentially damaging levels of shaking, this gives

us a pessimistic view of the potential impact of an earthquake.

We run an event-based PSHA with almost exactly the same settings as we did for the classical PSHA. To reduce computation time and memory usage, we sample logic trees with more than 200 realizations 200 times (`number_of_logic_tree_samples = 200`). For each sampled realization, we generate 200 seismic event sets (`ses_per_logic_tree_path = 200`)[12]. In total, we generate 902,134,602 events, out of which 8,615,935 disrupt one or more facilities.

Analogously to the probability of exceedance, our goal will be to compute the mean complementary cumulative distribution function (CCDF) of the worst-case impact of an earthquake within a period of 50 years. If we focus on a single hazard model, estimating these probabilities is straightforward. For example, suppose we want to estimate the mean probability that at least 2 facilities are disrupted by a single earthquake within our 50 year investigation time (the CCDF of 1 disrupted facility). We can do so by first counting the number of events that disrupt at least 2 facilities, c_r, for each realization, r. Since we assume earthquake occurrences are Poissonian, given realization r, we can estimate the probability that at least 2 facilities are potentially disrupted by a single earthquake by

$$1 - e^{-\frac{c_r}{200}} \tag{2.1}$$

To estimate the mean probability, we simply average these estimates over each realization.

In our case, the situation is more complex, since we need to combine results from multiple hazard models. Suppose we want to combine the results of two hazard models. For each hazard model, we sample up to 200 realizations. If we were to take the naive approach, we would need to count events for each of the 40,000 combinations of realizations. Clearly, this is not tractable for 19 hazard models.

Fortunately, under some conditions, we can combine mean probabilities instead.

Lemma 1. *Let n be the number of hazard models, and let $X_1, \ldots, X_n$ be random variables measuring the number of events of interest in each hazard model. Furthermore, let R_i be the realizations of hazard model i, and w_r the weight of realization $r \in R_i$.*

We define the mean probability

$$\overline{P}(\sum_{i=1}^{n} X_i \geq 1) = \sum_{r_1 \in R_1} w_{r_1} \cdots \sum_{r_n \in R_n} w_{r_n} P(\sum_{i=1}^{n} X_i \geq 1 | r_1, \ldots, r_n) \tag{2.2}$$

If X_1 to X_n are mutually independent, then

$$\overline{P}(\sum_{i=1}^{n} X_i \geq 1) = 1 - \prod_{i=1}^{n} \overline{P}(X_i = 0) \tag{2.3}$$

where

$$\overline{P}(X_i = 0) = \sum_{r \in R_i} w_r P(X_i = 0 | r) \tag{2.4}$$

[12] For hazard models with an investigation time of 1, we generate 10,000 seismic events sets per realization instead.

Proof. Since

$$P(\sum_{i=1}^{n} X_i \geq 1 | r_1, \ldots, r_n) = 1 - P(\sum_{i=1}^{n} X_i = 0 | r_1, \ldots, r_n) \tag{2.5}$$

and the realization weights of each hazard model sum to 1, we can reformulate Equation 2.2 as

$$\begin{aligned} &\overline{P}(\sum_{i=1}^{n} X_i \geq 1) = \\ &1 - \sum_{r_1 \in R_1} w_{r_1} \cdots \sum_{r_n \in R_n} w_{r_n} P(\sum_{i=1}^{n} X_i = 0 | r_1, \ldots, r_n) \end{aligned} \tag{2.6}$$

Now, since we have mutual independence and the results from hazard model i only depend on realization r_i:

$$\begin{aligned} 1 - \sum_{r_1 \in R_1} w_{r_1} \cdots \sum_{r_n \in R_n} w_{r_n} P(\sum_{i=1}^{n} X_i = 0 | r_1, \ldots, r_n) = \\ 1 - \sum_{r_1 \in R_1} w_{r_1} \cdots \sum_{r_n \in R_n} w_{r_n} \prod_{i=1}^{n} P(X_i = 0 | r_i) = \\ 1 - \sum_{r_1 \in R_1} w_{r_1} P(X_1 = 0 | r_1) \cdots \sum_{r_n \in R_n} w_{r_n} P(X_n = 0 | r_n) = \\ 1 - \prod_{i=1}^{n} \overline{P}(X_i = 0) \end{aligned} \tag{2.7}$$

□

Equation 2.3 allows us to estimate any overall mean CCDF, by separately computing the mean estimated probability of zero events (the weighted average of $e^{-\frac{C_r}{200}}$) for each hazard model.

Our approach ignores the potential overlap between different hazard models. Consider the ESHM13 and EMME14 hazard models for example. We use ESHM13 to estimate the hazard for facilities in Europe, and EMME14 for estimating the hazard in the Middle East. As these areas border each other, it is possible that an earthquake would disrupt facilities in both Europe and the Middle East. Our approach ignores this possibility, and thus potentially overestimates the total number of earthquakes (since multiple hazard models may model the same seismic sources), while underestimating the impact of some of these earthquakes. This problem is an inherent disadvantage of combining multiple hazard models.

RESULTS

We first consider the number of disrupted facilities. As can be seen in Fig. 2.11, the mean probability that at least one facility will be disrupted within 50 years is nearly 1. Worryingly, there are many events that would disrupt multiple facilities at once. There is a 10% probability that at least 20 facilities will be disrupted by a single earthquake. Given the level of facility sharing between IXPs, this could have a significant impact on the Internet.

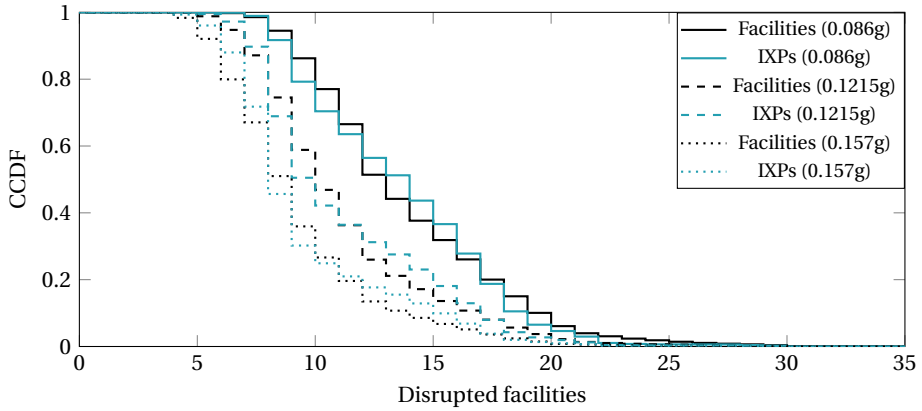
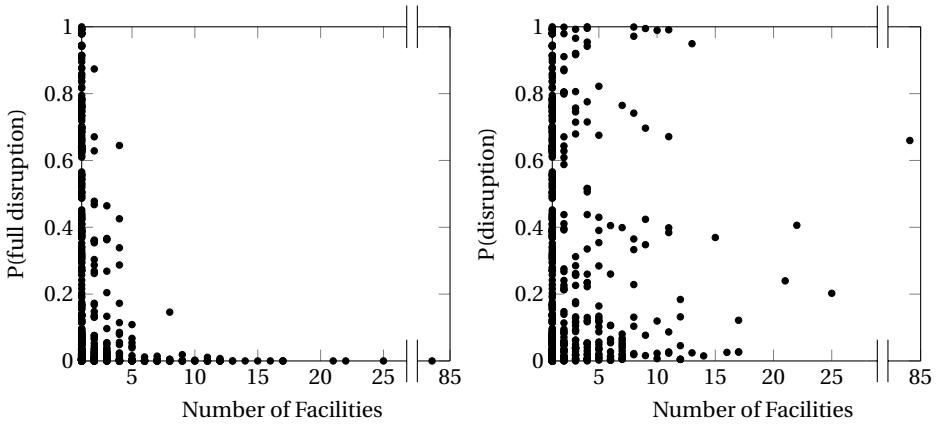


Figure 2.11: The complementary CDF of the maximum number of facilities (IXPs) that are simultaneously disrupted by a single earthquake within a period of 50 years.



(a) Full disruption (b) Disruption

Figure 2.12: The total number of facilities of each IXP, and their probabilities of (full) disruption within a period of 50 years.

Contrary to our expectations, the number of disrupted IXPs is often lower than the number of disrupted facilities. Furthermore, the worst-case number of simultaneously disrupted IXPs is quite a bit lower than the worst-case number of simultaneously disrupted facilities: 72 facilities compared to 46 IXPs. This shows that a number of IXPs are distributed over facilities that can be struck by the same earthquake.

For comparison, we also consider higher PGA thresholds (Fig. 2.11). While there is a clear decrease in earthquake impact if we increase the threshold to 0.157g (roughly corresponding to a macroseismic intensity of 7), the probability of simultaneous facility disruption is still quite high: There is a 3.7% probability that at least 20 facilities will si-

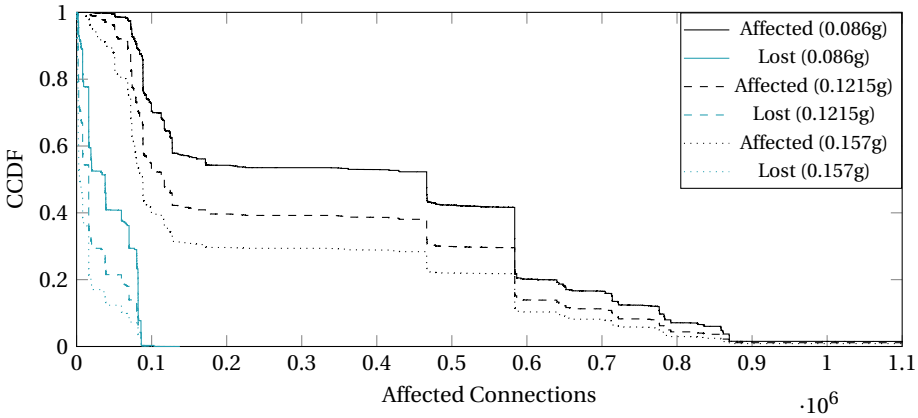


Figure 2.13: The complementary CDF of the maximum number of connections that are affected by a single earthquake within a period of 50 years.

multaneously experience this level of shaking. Nevertheless, the choice of PGA threshold greatly influences our results, and, since we chose a more pessimistic threshold, we are potentially overestimating the impact of earthquakes on IXPs.

Our results raise the question if IXPs spread their facilities over a large enough area. We compute the probability of full disruption of each of the 828 IXPs whose facilities are located in the area covered by the hazard models. It seems like distributing IXPs over multiple facilities helps: the median probability that an IXP is fully disrupted at least once in a 50-year period is 0.0118, while the median probability that an IXP with at least two facilities is fully disrupted is 0.00220. As can be seen in Fig. 2.12, IXPs with more facilities tend to have a lower probability of experiencing full disruption. These results should also translate to individual ASes. By peering with the same neighbors at multiple locations, an AS can significantly reduce the risk of earthquakes to its connectivity.

IMPACT ON CONNECTIVITY

To get a better idea of the impact of these events, we again consider the ASes hosted at each IXP. We assume that, within each IXP, every AS peers with every other AS. We then define a unique (potential) *connection* for every pair of ASes that share at least one IXP. While this is an overestimate of the actual peering density at each IXP, the number of connections should be roughly proportional to the actual number of peering links. Furthermore, the loss in connections due to IXP disruption is equivalent to the loss in available peering links at IXPs.

We assign two impact metrics to each event: (1) the number of affected connections, and (2) the number of lost connections. If two ASes share a disrupted facility, we mark their connection as affected. If the two ASes share no other undisrupted facility, the connection has no remaining backup and we mark it as lost. Note that this does not mean that these two ASes are completely disconnected from each other (packets can potentially still be routed through other ASes or through direct peering outside an IXP), but it does mean that these two ASes can not exchange packets directly at any remaining

IXP. In this manner, the metric is a good indicator of impact on the IXP ecosystem.

We note that a large majority of connections have a backup (Fig. 2.13). This shows the power of peering at multiple IXPs. Even if some facilities are disrupted by an earthquake, there is often another facility available that serves as a suitable backup.

That being said, the number of lost connections is still very high, even at higher probabilities (and at higher PGA thresholds). At best, this means that in case of a strong earthquake, a large number of BGP routes might need to be rerouted. At worst, ASes will be completely disconnected from the rest of the Internet.

2.5.2. INCREASING REDUNDANCY - A NOVEL METRIC

As we discussed in the previous section, operators can reduce the impact of earthquakes on their ASes by peering at multiple facilities. However, selecting a new facility is not trivial. Clearly, peering at a facility with low probability of exceeding damaging levels of shaking helps reduce the risk of earthquakes to the AS. The results from Section 2.4.2 suggest this factor is already taken into account to some degree: facilities that host more ASes tend to have a lower probability of exceeding damaging levels of shaking. However, only peering at low-risk facilities might not always be possible or cost-efficient, and, although less frequently, even a low-risk facility can be struck by an earthquake. Thus, to effectively reduce the risk of earthquakes, an operator would need to consider both the probability that its facilities will be disrupted by the same earthquake, as well as the redundancy of connections at each of its facilities.

In this section, we introduce a novel metric for evaluating sets of peering locations with respect to earthquake risk. Our metric can be applied to both IXP facilities, as well as to private peering. The aim of the metric is to ensure the probability that any of a selection of important connections is disconnected by an earthquake remains below a pre-selected threshold.

Definition 1 (Earthquake-Resistant Peering Metric)**.** *Suppose we are given a set of weights w_i for all ASes, a set of potential facilities F, the cost of peering at each facility $f \in F$, $c(f)$, and a threshold on the disconnection probability, $t \in [0,1]$. Let $h_i \subseteq F$ be the subset of all facilities hosting AS i.*

Given a selection of facilities $s \subseteq F$, the mean probability that the connection with AS i will be disrupted due to an earthquake is equivalent to the mean probability that facilities $h_i \cap s$ will simultaneously be disrupted due to an earthquake. We denote this probability by $p(h_i \cap s)$, and compute it using Equation 2.3.

We define the value of a selection of facilities $s \subseteq F$ as

$$\sum_i w_i I_i(s) - \sum_{f \in s} c(f) \tag{2.8}$$

where

$$I_i(s) = \begin{cases} 1 & \text{if } h_i \cap s \neq \emptyset \text{ and } p(h_i \cap s) \leq t \\ 0 & \text{otherwise} \end{cases} \tag{2.9}$$

Remark 1.1. *Note that one can easily extend this metric to require connectivity with only one out of a set of ASes, or to set individual thresholds per AS.*

EVALUATION

To evaluate our metric, we set a threshold of 0.01 in 50 years, and extract all ASes with at least one connection with a disruption probability above this threshold. We filter out any facilities outside of our hazard models, and any ASes peering at one of these facilities. Our goal will be to increase the resilience of the remaining 4,594 ASes against earthquakes, by connecting each AS to one additional facility.

For the purpose of this experiment, we consider each combination of IXP and facility (hosting the IXP) to be a unique facility. For each AS we aim to protect, we set the cost of each facility to 0, the weight of each of its current peers to 1, and the weight of all other ASes to 0. That is, our goal is to find the IXP-facility pair that protects as many of the current connections as possible.

Out of the 4,594 ASes, we find a new facility for 4,420 ASes. For the other 174 ASes, there is no possible facility that would reduce the disconnection probability of any of its peers to below our threshold of 0.01. The mean number of connections that were previously unsafe but are protected by adding a single facility is 31.8%. However, the mean distance between the closest old facility and this new facility is 2,579km.

If we restrict ourselves to the countries each AS currently peers at, we find a solution for 3,721 ASes. The average distance to the new facility is now 569km, and the facility protects an average of 26.6% of previously unsafe connections.

For 2,280 ASes (almost 50%), we can even find a new facility within 100km of their old facilities. These facilities protect an average of 20.2% of previously unsafe connections, while their average distance to the old facilities is only 24km.

Fig. 2.14 shows the effect of peering at *all* of these facilities on the number of lost connections during an earthquake. Since we chose to protect currently existing connections, we only consider these original connections. We can see that connecting to additional facilities did indeed protect many connections against earthquakes. Interestingly, while peering at additional facilities within the same country increased the resilience of both individual ASes and the Internet as a whole against earthquakes, restricting facilities to a distance of 100km of old facilities has a much smaller impact on the overall resilience of the Internet.

2.5.3. OVERLAP BETWEEN HAZARD MODELS

To study the impact of the overlap of hazard models on our results, we run a new event-based PSHA, where each hazard model is applied to both its own facilities, as well as the facilities of its neighboring models. This gives us an upper bound on our estimates of all mean probabilities. We then compute lower bounds by filtering out all earthquakes that disrupt any neighboring facility.

The difference between the lower and upper bound of the CCDF of the maximum number of disrupted facilities is at most 0.00610, that of the maximum number of disrupted IXPs at most 0.00725, and that of the maximum number of affected connections is 0.00848.

Even when we include neighboring facilities in our calculations, we find that not a single IXP that shares facilities across hazard models would be fully disrupted by any earthquake. Thus, the CCDF of the maximum number of lost connections is completely unaffected by overlap between hazard models.

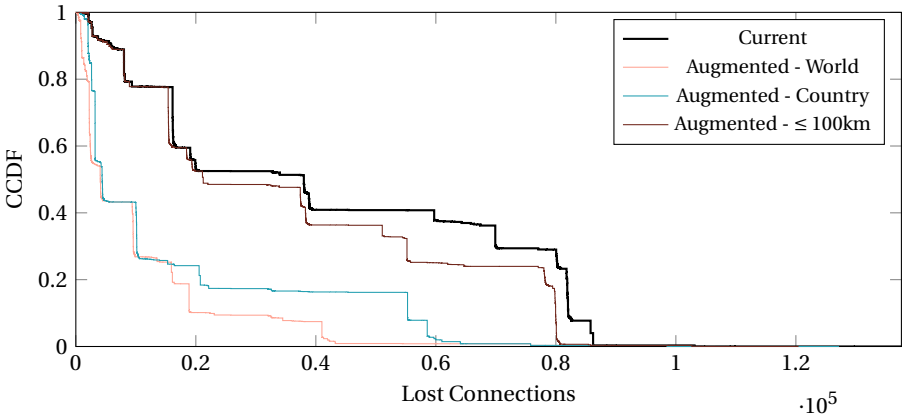


Figure 2.14: The complementary CDF of the maximum number of original connections that are lost due to a single earthquake within a period of 50 years, before and after spreading ASes over more facilities.

2.6. DISCUSSION

Our analysis is a best-effort analysis of the risk of earthquakes to global IXP infrastructure. The maps included in this chapter are not meant to be used to support any important decision involving human life, capital and movable and immovable properties. Due to the scale of our analysis, and our selection of hazard models, a number of concessions were made. The ground motion systems of our hazard models cannot account for the intra-event spatial correlation of ground motions. Furthermore, due to a lack of data, we assume the conditions of each site are equivalent (to reference rock). These conditions affect the level of shaking, and it is possible that some facility locations have been purposely placed in areas that are less susceptible to earthquakes. In addition, since we lack data on the characteristics of each facility as well, we say a facility is disrupted if it experiences potentially damaging levels of shaking. When building characteristics are known, one can use fragility curves to estimate a probability of damage instead.

2.7. CONCLUSION

We have conducted the first global study of the risk of earthquakes to Internet infrastructure. Whereas previous studies only considered the risk to each network component individually, we have combined a set of 19 earthquake hazard models and generated a total of 902,134,602 earthquake scenarios to estimate the hazard to individual facilities, as well as the probability of the disruption of multiple facilities at once, and the impact these disruptions could have on the Internet.

We find that a large number of IXP facilities are at risk of earthquakes: more than 30% of IXP facilities have at least a 10% probability of experiencing potentially damaging levels of shaking within 50 years, and more than 50% at least a 2%. On the positive side, IXP facilities that host large number of ASes tend to be located in less earthquake-prone areas.

Of course, the true danger of natural disasters lies in their ability to damage many facilities at once. There is a 10% probability that at least 20 facilities will experience potentially damaging levels of shaking *simultaneously*, within a period 50 years. This is equivalent to a return period of 475 years. Such an event could significantly disrupt local traffic.

We confirm the effectiveness of spreading out over multiple facilities: IXPs with more facilities tend to have a greatly reduced probability that all their facilities are disrupted simultaneously. However, we find that not all ASes spread out over IXP facilities sufficiently. To this end, we have proposed a novel metric for selecting new peering locations, that takes into account earthquake hazard at both current and new facilities, and the probability of combined facility failures. We have demonstrated that when ASes apply our metric, this both protects their own connections and a great number of existing peering connections of the Internet as a whole.

3

Computing the Impact of Disasters on Networks

In the previous chapter, we studied the risk of earthquakes to Internet Exchange Points. In this chapter, we propose a more general disaster resilience assessment framework that can be applied to any network and disaster set. We give an efficient method to compute the distribution of a network performance metric after a random disaster, based on a finite set of disaster regions and occurrence probabilities. Our approach has been implemented as a tool to help visualize the vulnerability of a network to disasters. With that tool, we demonstrate our methods on an official set of Japanese earthquake scenarios.

Parts of this chapter have been published in SIGMETRICS Perform. Eval. Rev. **45**, 2 (September 2017) [117].

3.1. INTRODUCTION

As discussed in Chapter 1, a large majority of approaches for assessing the disaster resilience of communication networks assume a disaster takes a fixed shape, and can occur anywhere within an area around the network with equal probability. These approaches allow one to assess the impact of regional failures on the network without any additional disaster data, but have a major disadvantage: In reality, as can be observed in Fig. 2.2 of Chapter 2, the probability and properties of disasters do greatly depend on location. To incorporate the inhomogeneity of disasters, we propose a flexible, data-driven framework for assessing the resilience of communication networks to disasters.

Our framework can be applied to any disaster dataset that can be transformed to a finite set of representative disaster regions and probabilities. In contrast to other approaches, which often either compute an expected or worst-case value, we compute the distribution of any given metric after one of the disasters randomly occurs. We show that this distribution can be efficiently calculated, and that it provides more information to a network operator or designer than any single value could.

Our main contributions in this chapter are threefold:

- We propose an efficient method to compute the distribution of any network performance metric, based on a finite set of disaster regions and occurrence probabilities.
- We describe our tool to compute and visualize such distributions for any network topology and disaster set.
- We demonstrate our method on a set of Japanese earthquake scenarios, and show how our approach can provide more insight into the disaster resilience of communication networks.

3.2. MODEL

We model the network as a directed multigraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \psi)$, with nodes $v \in \mathcal{V}$ connected by links $e \in \mathcal{E}$, where $\psi : \mathcal{E} \to \mathcal{V} \times \mathcal{V}$ and $e \in \mathcal{E}$ connects v_1 to v_2 iff $\psi(e) = (v_1, v_2)$. Thus, we permit the same pair of nodes to be connected by multiple links. We assume the network is embedded in a plane, and lies completely in a bounded convex region $R \subseteq \mathbb{R}^2$. Instead of modeling them as straight line segments, each link is modeled as a finite sequence of line segments connecting their nodes.

We model disasters deterministically, i.e., we assume that all links intersecting a disaster region, which we take as the region(s) in which ground motions exceed a specific level, fail. If a node lies within a disaster region, all of its links must have at least one endpoint in the disaster region and therefore would fail. Thus, we do not need to explicitly consider node failures.

Earthquakes typically occur at faults, and thus can not occur everywhere in R. In addition, the ground motion, and thus the disaster region after an earthquake, depends on the earthquake's magnitude, as well as the properties of the rocks and sediments that earthquake waves travel through. Many earthquakes with similar locations affect the

same links of the network, even though their exact disaster regions may differ. We therefore argue that it makes sense to take a finite representative set of earthquakes and use it to calculate the network's vulnerability.

We assume that we are given a finite set of possible disasters $\mathscr{D}$. We further assume that exactly one of these disasters will manifest at a time. The probability of multiple independent earthquakes occurring simultaneously is generally very small and thus is ignored in this chapter. Earthquakes that trigger other disasters (e.g., aftershocks) can still be modeled, by combining their disaster regions. Each disaster $d \in \mathscr{D}$ has a disaster region $A(d) \subseteq \mathbb{R}^2$ and an occurrence probability $P(d)$. Note that $\sum_{d \in \mathscr{D}} P(d) = 1$.

We model a disaster region as either a circle, line segment, simple polygon, or a finite union of these. However, our model and methods can be used with any shape of disaster region, as long as it is possible to calculate if a line segment intersects it. In fact, as we show in Chapter 7, our approach can even be applied to disasters and networks in a spherical model, without first projecting coordinates to the plane.

There are multiple ways to obtain the set $\mathscr{D}$. One can generate potential disaster scenarios in a Monte Carlo approach, as was done in Chapter 2. Another approach is to take a historical set of the last N earthquakes above a certain magnitude. Finally, one can use a given set of scenarios as input. As an example, in the following section, we will convert Japanese J-SHIS earthquake scenarios to our disaster model.

3.3. J-SHIS EARTHQUAKE SCENARIOS

Japan has one of the highest earthquake rates in the world and thus needs to be especially prepared for major earthquakes. The National Research Institute for Earth Science and Disaster Resilience (NIED) provides a large amount of data on Japanese earthquakes through the Japan Seismic Hazard Information Station (J-SHIS) [88]. We use the 2016 version of this dataset. Of particular interest to us are the Seismic Hazard Map and Scenario Earthquake Shaking Maps.

The Seismic Hazard Map gives probabilities for significant ground motion for all of Japan. These probabilities are calculated in a very similar method as our approach: by aggregating over a set of (representative) modeled earthquakes [153]. Unfortunately, as the end result is an aggregation, and the intermediate results are not publicly available, this map was not usable for our purposes.

Instead, we made use of the Scenario Earthquake Shaking Maps. These scenario maps contain, among other data, the JMA seismic intensities for each affected Divided Quarter Grid Square [154] cell in Japan. By converting these to geographical coordinates, and only keeping those grids with an intensity above a specific threshold, a disaster region (of a union of rectangles) can be obtained for every single scenario in the dataset. The resulting disaster regions are not contiguous, as there are gaps where the seismic intensity is below the threshold.

The scenarios do not contain occurrence probabilities. To obtain these probabilities, we take the mean recurrence intervals for each fault from the parameter dataset for the Seismic Hazard Map. If a fault segment has N scenarios and mean recurrence interval i,

the occurrence probability of all its disasters is taken to be

$$\frac{1}{iNT},$$

where T is the sum of the inverses of all recurrence intervals of fault segments with $N > 0$.

3.4. VULNERABILITY DISTRIBUTIONS

Liew et al. proposed characterizing network survivability by a function, rather than by a single value (like the expected value after a random disaster) [155]. In essence, their survivability function is the probability mass function of a given survivability metric after a random disaster. Some interesting values can easily be derived from this function, such as the worst-case survivability, r-percentile survivability, or the probability of zero survivability. Liew et al. did not apply their method to regional failures. In this section, we propose a method to efficiently compute these distributions in our disaster model.

3.4.1. FAILURE STATES

As an intermediate step towards computing metric distributions, we first consider the probability distribution over the state of the network after a random disaster.

Let a failure state s be defined as a set $s \subseteq \mathscr{E}$, where $e \in s$ if and only if e is down.

Let S be the random value indicating the failure state after the disaster and let $S(d)$ be the failure state after disaster $d \in \mathscr{D}$. Thus, $S(d)$ is the set of all links intersecting the disaster region $A(d)$.

Because we assume exactly one disaster occurs, we have

$$P(S = s) = \sum_{d \in \mathscr{D} | S(d) = s} P(d) \tag{3.1}$$

The distribution over S can now be computed as follows:

1. $\forall d \in \mathscr{D}$, compute $S(d)$
2. $\forall s \in S[\mathscr{D}]$ (the image of S), store $S^{-1}(s) = \{d \in \mathscr{D} | S(d) = s\}$
3. $\forall s \in S$, $P(S = s) = \sum_{d \in S^{-1}(s)} P(d)$

Note that $|S[\mathscr{D}]| \leq |\mathscr{D}|$ (trivially), and can be much smaller when many disasters occur in the same small region. The value of a metric only depends on the state of the network, and thus it only needs to be computed once per possible failure state, instead of once for each $d \in \mathscr{D}$. By iterating over possible failure states instead of disasters, we can potentially significantly reduce the computation time of the distribution over a metric.

3.4.2. METRICS

Consider a metric M. Let $M(d)$ be the value of the metric after disaster d, and $M(s)$ be the value of the metric in failure state s. Note that $M(d) = M(S(d))$.

Similarly as in equation 3.1, we have

$$\begin{aligned} P(M=m) &= \sum_{d\in\mathscr{D}|M(d)=m} P(d) \\ &= \sum_{s\in S[\mathscr{D}]|M(s)=m} \left(\sum_{d\in\mathscr{D}|S(d)=s} P(d) \right) \\ &= \sum_{s\in S[\mathscr{D}]|M(s)=m} P(S=s) \end{aligned} \tag{3.2}$$

The distribution over M can now be calculated as follows:

1. $\forall s \in S[\mathscr{D}]$, compute $P(S=s)$ as described in section 3.4.1
2. $\forall s \in S[\mathscr{D}]$, compute $M(s)$
3. $\forall m \in M[S[\mathscr{D}]]$, store $\{s \in S[\mathscr{D}] | M(s) = m\}$
4. $\forall m, P(M=m) = \sum_{s\in S[\mathscr{D}]|M(s)=m} P(S=s)$

Note that this method can be performed in parallel, to further increase performance.

3.5. DISASTER IMPACT VISUALIZATION

The disadvantage of computing a distribution instead of a single value is that one may be overwhelmed by the amount of data. Thus, it is important to properly visualize the results in a useful fashion.

The distributions over a metric can be clearly visualized with a histogram of the cumulative distribution function (CDF), for example, as in figures 3.2 and 3.3.

The intermediate results of the computations in section 3.4.2, such as the distribution over failure states and the coupling of disasters with their resulting state and metric, can also greatly help in preparing the network against disasters.

To this end, we have created the Disaster Impact Visualization Tool (DIVT). This tool can, given any network topology and disaster set, compute and visualize the vulnerability distribution and intermediate results. DIVT maps the network on a world map using the NASA World Wind library (`worldwind.arc.nasa.gov`). By drawing disaster regions over the network, users can clearly see which links are affected by a disaster and why.

The metric distribution, state distribution, and the coupling between these distributions and the disasters themselves, are visualized in a tree structure (see Fig. 3.1).

At the top level one can see and select the values of the metric with their corresponding probability. Their child nodes show the probabilities of the states resulting in these values. Finally, at the lowest level are the individual disasters causing these states. By selecting one or more of these tree nodes, all corresponding disaster regions are drawn in red on the map. Failed links are colored pink.

An example is given in Fig. 3.1. We first expanded all failure states that result in an Average 2-Terminal Reliability – the number of connected node pairs divided by the total amount of node pairs – of 0.87512. Subsequently, we expanded a specific failure state with 4 failed links. This failure state is the result of either disaster scenario “F006104

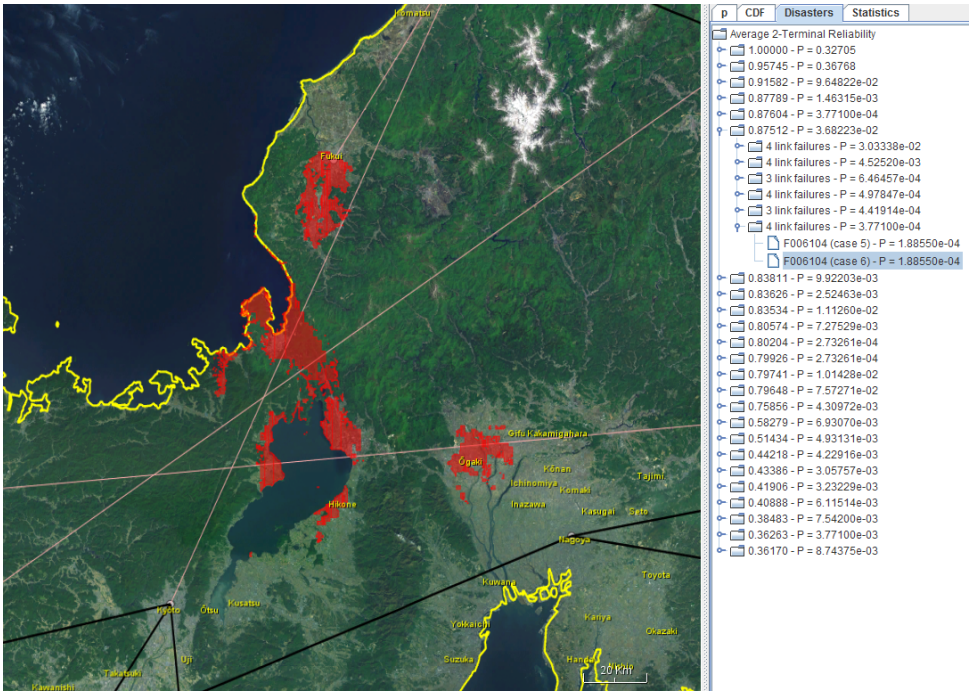


Figure 3.1: Visualization of distributions and disasters. Red: disaster region, pink: affected links.

(case 5)" or "F006104 (case 6)" (names are assigned based on J-SHIS Fault Code and Case Number). "F006104 (case 6)" was selected and is drawn on the map. Some basic statistics, like the expected value, variance, worst case, and all CDF values, are computed and displayed via the Statistics tab.

3.6. EXPERIMENTAL RESULTS

In this section, we demonstrate the use of our methods on two Japanese network topologies: JGN2plus-Japan and Sinet. Both were downloaded from the Topology Zoo [156]. As these files only contain geographical coordinates for the nodes, and not the links, all links are assumed to be straight line segments directly connecting their endpoints. The Mercator projection was used to map all geographical coordinates to the 2-dimensional plane. Nodes without any geographical information were ignored.

JGN2plus-Japan spans almost all of Japan, but only has 11 nodes and 10 links. In contrast, Sinet spans a slightly smaller region, but consists of 47 nodes and 49 links.

As disasters, we took the J-SHIS earthquake scenarios described in section 3.3, specifically those from the 2016 dataset. These comprise 655 scenarios for 189 fault segments. The JMA seismic intensity threshold was set to 5.5. We chose the average 2-terminal reliability (ATTR) as our metric.

In Fig. 3.2 and 3.3, the cumulative distribution functions of the ATTR of both networks, after one of the earthquake scenarios, has been plotted. One may notice a dif-

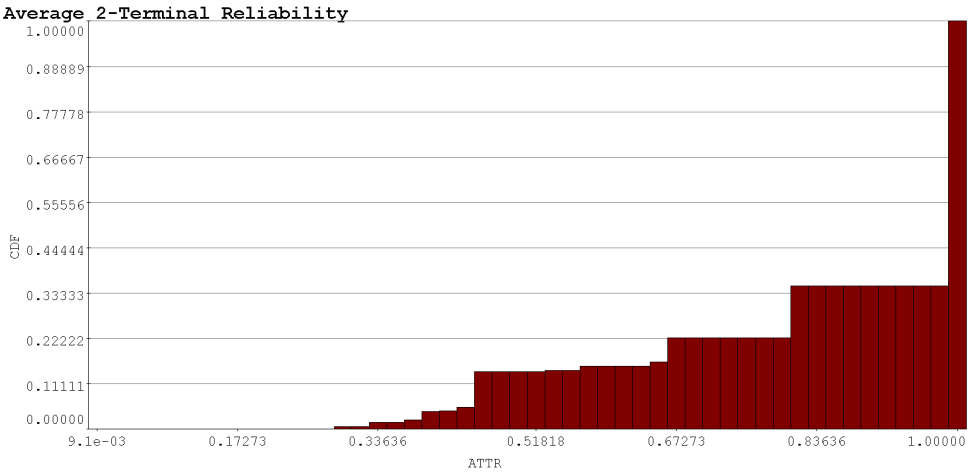


Figure 3.2: ATTR distribution of JGN2plus-Japan.

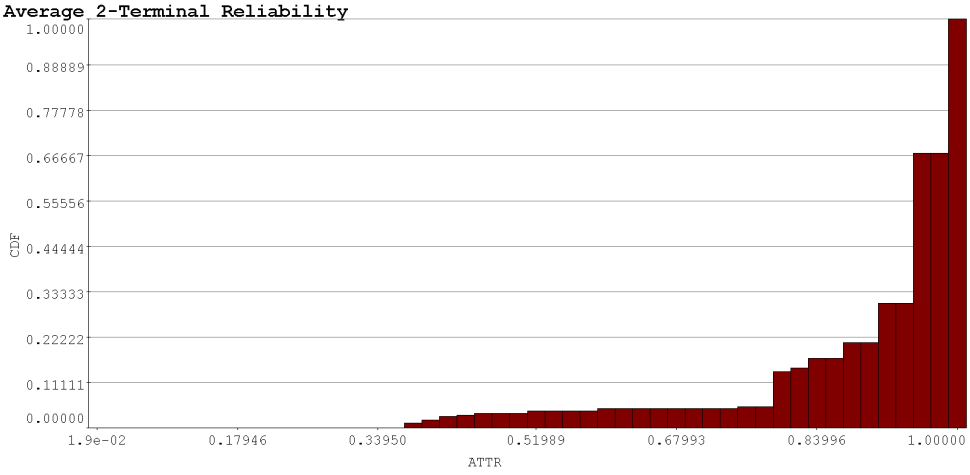


Figure 3.3: ATTR distribution of Sinet.

ference between the two: while JGN2plus-Japan has a much lower probability of becoming disconnected than Sinet (0.352 and 0.673, respectively), its probability of incurring a large ATTR impact is much higher than for Sinet. $P(ATTR \leq 0.7)$ is 0.224 for JGN2plus-Japan and 0.049 for Sinet.

This is probably caused by the large difference in network size between both networks. As JGN2plus-Japan consists of fewer nodes and links, it has a higher probability that it will not be hit by the earthquake at all. However, in the case that the network does get hit, it lacks the backup paths to keep most of its connections. We can confirm this by inspecting P(No Link Failures) in DIVT. Indeed, the probability of all links of JGN2plus-

Japan being unaffected is 0.648, and there are no possible states in which any link fails, but the network stays connected. The comparatively low P(No Link Failures) of Sinet is 0.263.

The worst-case disasters for JGN2plus-Japan all occur around Tokyo, resulting in an ATTR of 0.291 with probability 0.007. The worst-case disasters for Sinet are located around Osaka, and result in an ATTR of 0.362 with probability 0.009. JGN2plus-Japan has an expected ATTR value of 0.866 with a variance of 0.044 and Sinet an expected ATTR value of 0.920 with a variance of 0.016.

For both networks, only computing the ATTR for each possible failure state, instead of for each disaster, had a large effect on performance, reducing the number of times ATTR had to be computed from 655 to 22 and 93 for JGN2plus-Japan and Sinet, respectively.

3.7. CONCLUSION

We have proposed an efficient, data-driven approach for assessing the resilience of a communication network to disasters, by computing the distribution of the impact of a random disaster scenario. One of the key insights of our approach is that, since the number of considered disaster scenarios tends to be much larger than the number of potential outcomes (which we call *failure states*), we can greatly reduce computation times by first computing the distribution of failure states, and only then computing the value of the impact metric for each unique failure state.

We have implemented our approach within a visualization tool that can draw both selected disaster shapes and the network topology itself, and have applied our method and tool to a dataset of Japanese earthquake scenarios to demonstrate how they can give more insight into the disaster resilience of a network.

4

The Risk of Successive Disasters: A Blow-by-Blow Network Vulnerability Analysis

In the majority of this thesis we assume, like many others, that a network will not be struck by multiple disasters in a relatively short period of time; that is, a subsequent disaster will not strike within the recovery phase of a previous disaster. However, recent events have shown that combinations of disasters are plausible. This realization calls for a new perspective on how we assess the vulnerability of our networks and shows a need for a framework to assess the vulnerability of networks to successive independent disasters.

We propose a network and disaster model capable of modeling a sequence of disasters in time, while taking into account recovery operations. Based on that model, we develop both an exact and a Monte Carlo method to compute the vulnerability of a network to successive disasters. By applying our approach to real empirical disaster data, we show that the probability of a second disaster striking the network during recovery can be significant even for short repair times. Our framework enables stakeholders to determine the vulnerability of networks to such successive disasters.

Parts of this chapter have been published in IFIP Networking Conference, 2019 [157].

4.1. INTRODUCTION

The rate at which disasters strike an area is typically very low. Therefore, it is commonly assumed that a network will only be affected by a single (possibly composite[1]) isolated disaster at a time. The probability that two or more independent disasters will occur shortly after one another is seen as negligible and safe to ignore. Recent events have shown that this assumption might not be as rock solid as first thought.

In 2017, the continental United States was hit by 3 hurricanes (Harvey, Irma, and Nate), of which two where categorized as major hurricanes (Harvey and Irma) [158]. Hurricane Irma hit the East Coast only 16 days after Harvey [159, 160]. Out of the top 5 costliest US mainland tropical cyclones on record, 3 occurred in 2017 [161].

In total, there were 16 billion-dollar weather and climate disaster events in the United States in 2017 [162]. The total cost of these events exceeded 300 billion dollars. Over 2013-2017, the United States has had an average of 11.6 major disasters per year with a cost of more than 1 billion dollars.

Also in 2017, Mexico was hit by two major earthquakes in two weeks (where the second quake is not considered an aftershock of the first [163]), leading to a combined economic loss of nearly 6 billion dollars [164, 165].

Recovering a network after a disaster can take several weeks to months, as a large amount of hardware will need to be replaced or repaired in a potentially very inaccessible area [14]. In the context of this chapter, a network is said to be affected by multiple successive disasters if a disaster strikes the network during its recovery from a previous disaster. Depending on the moment in the recovery phase when the next disaster occurs, the total impact and final recovery time will differ significantly.

To increase the resilience of communication networks to disasters, it is essential to be able to compute the vulnerability of networks to these disasters. *While previous work has been instrumental in computing the vulnerability of a network to a single disaster, it has not addressed multiple successive disasters.* In this chapter, we propose a framework to assess the vulnerability of a network to successive disasters. Our main contributions are as follows:

- We compose a network and disaster model capable of modeling a sequence of disasters in time (Section 4.2).
- We develop a method to compute the vulnerability of a network to successive disasters by modeling the network state as a discrete-time Markov chain (Section 4.4). Our methodology allows for arbitrary precision by only computing the effect of at most k successive disasters, with corresponding error bounds. Our results for the Markov chain are subsequently used to derive a faster Monte Carlo method in Section 4.5.
- We apply our methods to empirical disaster data in Section 4.6. These experiments show that the probability of a second disaster striking the network during recovery can be significant, even for short repair times.

[1] Highly correlated disasters such as an earthquake and its aftershocks, can be modeled as a single composite disaster.

To the best of our knowledge, we are the first to propose models and methods for assessing the impact of successive disasters on networks, while taking into account recovery operations.

4.2. NETWORK AND DISASTER MODEL

We model the network as a directed multigraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \psi)$ with nodes $v \in \mathcal{V}$ connected by links $e \in \mathcal{E}$, where $\psi : \mathcal{E} \to \mathcal{V} \times \mathcal{V}$ and $e \in \mathcal{E}$ connects v_1 to v_2 iff $\psi(e) = (v_1, v_2)$. We define a failure set s, where network component $c \in \mathcal{V} \cup \mathcal{E}$ is functioning if and only if $c \notin s$. In the remainder of this chapter, *we refer to the failure set of a network as the state of that network.*

Given such a network, we are interested in three factors: (1) the number of successive disasters we can expect the network to be struck by, (2) the impact of being struck by one or more disasters, and (3) the total time it takes to fully recover from these disasters. To assess these attributes, we need to model the occurrence of disasters over time.

The occurrence of disasters is inherently unpredictable. A common stochastic model for disaster occurrences [104, 166, 167], which we will also employ, is the Poisson process. We model all disaster processes as mutually independent Poisson processes and assume we are given a multiset of disaster processes $d = (a_d, \lambda_d) \in \mathscr{D}^*$, where $a_d \subseteq \mathcal{V} \cup \mathcal{E}$ are the components affected by d and λ_d is the rate of d.

If disaster process d triggers at time t, when the network state is s, the new network state at time t will be $s \cup a_d$. That is, all components in a_d fail. We assume at most one disaster can strike the network at any given time t.

The combination of multiple Poisson processes is again Poissonian, with as rate the sum of its component rates. Thus, we can merge all disaster processes that affect the same set of network components without affecting the outcome of our analysis. Hence, we transform the set $\mathscr{D}^*$ to

$$\mathscr{D} = \{(a_d, \lambda_d) | a_d \neq \emptyset \wedge \lambda_d = \sum_{(a_d, \lambda_{d^*}) \in \mathscr{D}^*} \lambda_{d^*} > 0\} \tag{4.1}$$

Let $(T_n)_{n=1}^{\infty}$ be the ordered sequence such that T_1 is the occurrence time of the first disaster, and for all $n > 1$, T_n is the time between disasters $n-1$ and n. Let $(D_n)_{n=1}^{\infty}$ be the ordered sequence of disasters. In other words, the first disaster $D_1 \in \mathscr{D}$ occurs at time $T_1 \in \mathbb{R}$, the second $D_2 \in \mathscr{D}$ at $T_1 + T_2 \in \mathbb{R}$, etc. Then, for all $n \in \mathbb{N}$, T_n is exponentially distributed:

$$T_n \sim \text{Exp}(\lambda_D) \text{ (where } \lambda_D := \sum_{(a_d, \lambda_d) \in \mathscr{D}} \lambda_d) \tag{4.2}$$

Furthermore, D_n and T_n are independent:

$$P(D_n = d \wedge T_n = t) = P(D_n = d) P(T_n = t) \tag{4.3}$$

4.2.1. EXAMPLE NETWORK AND DISASTERS INSTANCE

To illustrate our network and disaster model, we give an example in Fig. 4.1. We consider a small triangle network of 3 nodes and 3 links. Its representative set of disasters contains four disaster scenarios. As each of these disasters affects a different set of components, $\mathscr{D}^* = \mathscr{D}$. The total disaster rate is $\lambda_D = 1.6$ disasters per year.

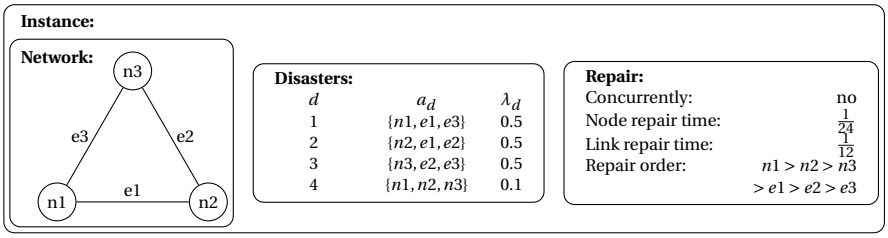


Figure 4.1: Example problem instance.

A network topology and set of disasters are not sufficient to properly compute the vulnerability of the network to successive disasters, as the impact of these disasters significantly depends on how quickly, and in what order, the network can be repaired. Thus, we also need to include some repair properties.

Our framework can include any repair function, but in the example the following repair rules hold: nodes can be repaired in half a month, while links take a full month to repair, and repairs are performed according to a predetermined priority and cannot be performed concurrently.

4.3. PROBLEM STATEMENT

We consider a deterministic repair model. We assume that, given a certain starting state, the recovery of the network is fixed (until a new disaster occurs). For example, if disaster 4 of the example instance occurs, all nodes will be damaged. Afterwards, the nodes will be repaired one by one. Thus, unless another disaster occurs during repair, the state of the network will be

- $\{n1, n2, n3\}$ at time 0
- $\{n2, n3\}$ at time $\frac{1}{24}$
- $\{n3\}$ at time $\frac{2}{24}$
- $\emptyset$ at time $\frac{3}{24}$

Generalizing the above example, we define repair functions $r_{s_0} : \mathbb{R}^+ \to \mathcal{V} \cup \mathcal{E}$ for each $s_0 \in \mathcal{V} \cup \mathcal{E}$. $r(t)_{s_0} \in \mathcal{V} \cup \mathcal{E}$ is the state of the network at time $t+C$, given that the state of the network was s_0 after being struck by a disaster at some time C. We assume the network does not degrade further in the recovery phase:

$$r(b)_{s_0} \subseteq r(a)_{s_0}\ 0 \leq a \leq b, s_0 \in \mathcal{V} \cup \mathcal{E} \tag{4.4}$$

Different repair strategies can be compared by changing the repair functions. Additionally, by increasing the number of components being repaired simultaneously, the benefits of acquiring more personnel can be assessed and compared to the additional cost in salary.

In the following, we elaborate on our research objectives with respect to three properties.

4.3.1. NUMBER OF SUCCESSIVE DISASTERS N

Network operators should decide on how many successive disasters they prepare for. To do so, knowing the probability of at least n successive disasters is essential. In addition, the expected number of successive disasters is also of interest. Hence, our goal is to compute $P(N \geq n)$, as well as $E[N]$.

4.3.2. IMPACT

While knowing the expected number of successive disasters is useful, it is also important to consider their impact. Suppose we have a metric $M : \mathcal{V} \times \mathcal{E} \to [0,1]$ that assigns a value $M(s)$ between 0 (worst case) and 1 (best case) to each state s of the network. We require that $M(a) \leq M(b)$ if $b \subseteq a$.

We analyze the minimum value of M during the disaster-and-recovery process. In the one-disaster case, this would simply be the value of M directly after the disaster. Successive disasters, although rare, can have a significantly higher impact on the network than single disasters. Therefore, given a critical value m, we want to compute the probability that the network reaches a state at least as bad as m during the disaster-and-recovery process, $P(M_{\min} \leq m)$, where $M_{\min}$ is the minimum value of M between T_1 and full recovery.

4.3.3. TOTAL TIME TO FULL RECOVERY

Let T_{total} be the total repair time, from the start of the first disaster to the time when all damage from all previous disasters has been repaired. We aim to compute the expected time to full recovery, $E[T_{\text{total}}]$.

4.4. ANALYSIS

In this section, we describe methods for computing the properties introduced in the previous section by modeling the state of the network as a Discrete-Time Markov Chain (DTMC).

4.4.1. MARKOV CHAIN

Let A_n be the state of network G directly after the nth disaster strikes the network. Now, because the disaster processes are independent and memoryless, and the repair function is deterministic,

$$P(A_n = a_n | A_1 = a_1, A_2 = a_2, \ldots, A_{n-1} = a_{n-1}) = P(A_n = a_n | A_{n-1} = a_{n-1}) \tag{4.5}$$

that is, $(A_n)_{n=1}^{\infty}$ satisfy the Markov property and form a (discrete-time) Markov chain.

The transition probabilities of this Markov chain depend on which disaster strikes next, as well as at which stage of the repair process this disaster strikes. By property (4.3), these two factors are independent. Thus, the transition probabilities can be calculated by summing over all possible disasters $d \in \mathcal{D}$:

$$P(A_n = a_n | A_{n-1} = a_{n-1}) = \sum_{d \in \mathcal{D}} \frac{\lambda_d}{\lambda_D} (\exp(-\lambda_D \mathcal{M}_{a_{n-1},d,a_n}) - \exp(-\lambda_D \mathcal{S}_{a_{n-1},d,a_n})) \tag{4.6}$$

Here, $\frac{\lambda_d}{\lambda_D}$ is the probability that the network will be struck by disaster $d = (a_d, \lambda_d)$. $[\mathscr{M}_{a_{n-1},d,a_n}, \mathscr{S}_{a_{n-1},d,a_n})$ is the period of time during which the occurrence of disaster d will result in network state a_n and $\exp(-\lambda_D \mathscr{M}_{a_{n-1},d,a_n}) - \exp(-\lambda_D \mathscr{S}_{a_{n-1},d,a_n})$ the probability that the next disaster will occur in this period of time[2].

We are specifically interested in the chain of network states until full recovery. Thus, we construct an additional Markov chain $(S_n)_{n=1}^{\infty}$ by adding an absorbing state $\emptyset$ to $(A_n)_{n=1}^{\infty}$ such that $S_n = \emptyset$ if and only if the network has been fully repaired.

Let $R_s := \min\{t \geq 0 | r(t)_s = \emptyset\}$ be the time it takes to fully repair the network (assuming no subsequent disasters occur), starting from network state $s \in \mathscr{V} \cup \mathscr{E}$. The probability that, starting in state s, the network is fully recovered before the next disaster strikes is $\exp(-\lambda_D R_s)$. Therefore, the transition probabilities to the absorbing state $\emptyset$ are

$$P(S_n = \emptyset | S_{n-1} = s_{n-1}) = \begin{cases} 1 & \text{if } s_{n-1} = \emptyset \\ \exp(-\lambda_D R_{s_{n-1}}) & \text{if } s_{n-1} \neq \emptyset \end{cases} \tag{4.7}$$

and the transition probabilities to all other states are

$$P(S_n = s_n \neq \emptyset | S_{n-1} = s_{n-1}) = \begin{cases} 0 & \text{if } s_{n-1} = \emptyset \\ \sum\limits_{d \in \mathscr{D}} \frac{\lambda_d}{\lambda_D} (\exp(-\lambda_D \min(\mathscr{M}_{s_{n-1},d,s_n}, R_{s_{n-1}})) & \\ -\exp(-\lambda_D \min(\mathscr{S}_{s_{n-1},d,s_n}, R_{s_{n-1}}))) & \text{if } s_{n-1} \neq \emptyset \end{cases} \tag{4.8}$$

$S_1 = A_1 = a_{D_1}$, so the initial distribution of the Markov chain $(S_n)_{n=1}^{\infty}$ is

$$P(S_1 = s_1) = \begin{cases} \frac{\lambda_d}{\lambda_D} & \exists d \in \mathscr{D} \text{ s.t. } a_d = s_1 \\ 0 & \text{otherwise} \end{cases} \tag{4.9}$$

4.4.2. NUMBER OF SUCCESSIVE DISASTERS N

We can now compute the probability $P(N \geq n) = 1 - P(S_n = \emptyset)$ of at least n successive disasters without full recovery. This probability decreases exponentially with n.

Lemma 2.

$$P(N \geq n) \leq (1 - \exp(-\lambda_D R))^{n-1} \tag{4.10}$$

where $R := \max\limits_{s \subseteq \mathscr{V} \cup \mathscr{E}} R_s$.

Proof. We give a proof by induction on n. Trivially, $P(N \geq 1) = 1 \leq (1 - \exp(-\lambda_D R))^0$.

Now, suppose

$$\forall k < n \; P(N \geq k) \leq (1 - \exp(-\lambda_D R))^{k-1},$$

then

$$\begin{aligned} P(N \geq n) &= P(N \geq n-1) P(N \geq n | N \geq n-1) \\ &\leq (1 - \exp(-\lambda_D R))^{n-2} P(N \geq n | N \geq n-1) \end{aligned}$$

[2] $\mathscr{M}_{a_{n-1},d,a_n}$ is the first time at which $r_{a_{n-1}} \cup a_d = a_n$ (or ∞ if no such time exists), and $\mathscr{S}_{a_{n-1},d,a_n}$ is the first time after $\mathscr{M}_{a_{n-1},d,a_n}$ at which $r_{a_{n-1}} \cup a_d \neq a_n$ (or ∞).

By direct application of (4.7):

$$\begin{aligned} P(N \geq n | N \geq n-1) &= 1 - P(S_n = \emptyset | S_{n-1} \neq \emptyset) \\ &= 1 - \frac{1}{P(S_{n-1} \neq \emptyset)} \sum_{s \neq \emptyset} P(S_{n-1} = s) P(S_n = \emptyset | S_{n-1} = s) \\ &= 1 - \frac{1}{P(S_{n-1} \neq \emptyset)} \sum_{s \neq \emptyset} P(S_{n-1} = s) \exp(-\lambda_D R_s) \\ &\leq 1 - \frac{1}{P(S_{n-1} \neq \emptyset)} \sum_{s \neq \emptyset} P(S_{n-1} = s) \exp(-\lambda_D R) \\ &= (1 - \exp(-\lambda_D R)) \end{aligned}$$

So,

$$P(N \geq n) \leq (1 - \exp(-\lambda_D R))^{n-1}$$

□

Remark 2.1. *If $R_s = R \; \forall s \neq \emptyset \in \mathcal{V} \cup \mathcal{E}$, then*

$$P(N \geq n) = (1 - \exp(-\lambda_D R))^{n-1}$$

Typically, $R = \max_{s \subseteq \mathcal{V} \cup \mathcal{E}} R_s$ will be the amount of time it takes to repair all network components ($R_{\mathcal{V} \cup \mathcal{E}}$).

Unfortunately, computing $E[N]$ directly is intractable in most cases, as the number of possible states can be as high as $2^{|\mathcal{V}|+|\mathcal{E}|}$. However, we can approximate (from below) the expected number of successive disasters by only constructing the Markov model for k successive disasters and computing the distribution of S_1 to S_k. The choice of k depends on the required accuracy.

Theorem 3 (Stopping conditions 1)**.** *Let $\hat{E}[N] = \sum_{n=1}^{k} P(N \geq n)$, then*

$$0 \leq E[N] - \hat{E}[N] \leq \frac{(1 - \exp(-\lambda_D R))^k}{\exp(-\lambda_D R)} \tag{4.11}$$

In addition, if $P(N \geq k) \leq \epsilon \frac{\exp(-\lambda_D R)}{1 - \exp(-\lambda_D R)}$, then

$$E[N] - \hat{E}[N] \leq \epsilon \tag{4.12}$$

Proof. We start by proving (4.11).

$$\begin{aligned} E[N] - \hat{E}[N] &= \sum_{n=k+1}^{\infty} P(N \geq n) \\ &\leq \sum_{n=k+1}^{\infty} (1 - \exp(-\lambda_D R))^{n-1} \text{ (Lemma 2)} \\ &= \frac{(1 - \exp(-\lambda_D R))^k}{\exp(-\lambda_D R)} \end{aligned}$$

If $P(N \geq k) \leq \epsilon \frac{\exp(-\lambda_D R)}{1-\exp(-\lambda_D R)}$, then (for $n \geq k$):

$$P(N \geq n) \leq \epsilon \exp(-\lambda_D R)(1-\exp(-\lambda_D R))^{n-k-1}$$

This can be proved analogously to Lemma 2. But this means that the absolute error

$$\begin{aligned} & E[N] - \hat{E}[N] \\ & \leq \sum_{n=k+1}^{\infty} \epsilon \exp(-\lambda_D R)(1-\exp(-\lambda_D R))^{n-k-1} \\ & = \sum_{n=0}^{\infty} \epsilon \exp(-\lambda_D R)(1-\exp(-\lambda_D R))^{n} \\ & = \epsilon \end{aligned}$$

4

□

Thus, to guarantee an upper bound on the absolute error, we can either choose the number of steps k beforehand, or test if $P(N \geq k)$ is below the threshold after every iteration, where the latter requires fewer iterations than the former.

4.4.3. IMPACT

As M is minimal directly after a disaster, $M_{\min} = \min_n M(S_n)$. The cumulative distribution function $P(M_{\min} \leq m)$ is the hitting probability of $M^{\leq m} := \{s \in \mathcal{V} \times \mathcal{E} | M(s) \leq m\}$. We can take a similar approach as before and approximate these probabilities as

$$\hat{P}(M_{\min} \leq m) := P(M^k_{\min} \leq m) \tag{4.13}$$

where $M^k_{\min} = \min_{n \leq k} M(S_n)$.

Suppose we have computed the first k states and corresponding transition probabilities of the Markov chain $(S_n)_{n=1}^{\infty}$. To compute $P(M^k_{\min} \leq m)$, we construct a new Markov chain $(S^{\leq m}_n)_{n=1}^{\infty}$ by replacing all $s \in M^{\leq m}$ with a single absorbing state $\mathscr{A}^{\leq m}$. Now,

$$P(M^k_{\min} \leq m) = P(S^{\leq m}_k = \mathscr{A}^{\leq m}) \tag{4.14}$$

Theorem 4 (Stopping conditions 2)**.** *Let*

$$\hat{P}(M_{\min} \leq m) = P(M^k_{\min} \leq m) = P(S^{\leq m}_k = \mathscr{A}^{\leq m})$$

Then

$$\begin{aligned} 0 \leq P(M_{\min} \leq m) - \hat{P}(M_{\min} \leq m) \leq \\ 1 - \hat{P}(M_{\min} \leq m) - P(S^{\leq m}_k = \emptyset) \leq P(N \geq k) \\ \leq (1-\exp(-\lambda_D R))^{k-1} \end{aligned} \tag{4.15}$$

Proof. If $m \geq 1$, then
$P(M_{\min} \leq m) = \hat{P}(M_{\min} \leq m) = 1$, so we assume that $m < 1$.

In this case

$$\begin{aligned} & P(M_{\min} \leq m) - \hat{P}(M_{\min} \leq m) \\ = & P(M_{\min} \leq m) - P(M^k_{\min} \leq m) \\ = & P(M_{\min} \leq m \wedge M^k_{\min} > m) \\ \leq & 1 - P(M^k_{\min} \leq m) - P(S_k^{\leq m} = \emptyset) \\ \leq & P(N \geq k) \end{aligned}$$

□

4.4.4. TOTAL TIME TO FULL RECOVERY

The total time to full recovery, or the total repair time, T_{total}, is equivalent to the sum of the time spent on repair in all states of $(S_n)_{n=1}^{\infty}$:

$$T_{\text{total}} = \sum_{n=1}^{\infty} \mathscr{R}_n \tag{4.16}$$

where $\mathscr{R}_n$ is the time spent on repairs between the nth and $(n+1)$th disaster. Thus, $\mathscr{R}_n$ is 0 if $S_n = \emptyset$ and $\mathscr{R}_n$ is the minimum between the total repair time of failures S_n and the time till the next disaster otherwise:

$$\mathscr{R}_n = \begin{cases} 0 & \text{if } S_n = \emptyset \\ \min(R_{S_n}, T_{n+1}) & \text{if } S_n \neq \emptyset \end{cases} \tag{4.17}$$

The expected value of $\mathscr{R}_n$ is

$$\begin{aligned} & E[\mathscr{R}_n] = \\ & \sum_{s \neq \emptyset} P(S_n = s)\left(\int_0^{R_s} \lambda_D \exp(-\lambda_D t)\, t\, dt + \exp(-\lambda_D R_s) R_s\right) \\ & = \sum_{s \neq \emptyset} P(S_n = s)\left(\frac{1}{\lambda_D}(1 - \exp(-\lambda_D R_s))\right) \\ & = \frac{1}{\lambda_D} \sum_{s \neq \emptyset} P(S_n = s)(1 - \exp(-\lambda_D R_s)) \end{aligned} \tag{4.18}$$

As before, we propose approximating $E[T_{\text{total}}]$ by truncating (4.16). That is, we approximate $E[T_{\text{total}}]$ by summing the expected values of $\mathscr{R}_1$ to $\mathscr{R}_k$, which only requires the distributions of S_1 to S_k.

Theorem 5 (Stopping conditions 3)**.** *Let* $\hat{E}[T_{total}] := \sum_{n=1}^{k} E[\mathscr{R}_n]$*, then*

$$0 \leq E[T_{total}] - \hat{E}[T_{total}] \leq \frac{(1 - \exp(-\lambda_D R))^k}{\lambda_D \exp(-\lambda_D R)} \tag{4.19}$$

In addition, if $P(N \geq k) \leq \epsilon \lambda_D \frac{\exp(-\lambda_D R)}{1 - \exp(-\lambda_D R)}$*, then*

$$E[T_{total}] - \hat{E}[T_{total}] \leq \epsilon \tag{4.20}$$

Proof. By the monotone convergence theorem,

$$E[T_{\text{total}}] = E[\sum_{n=1}^{\infty} \mathscr{R}_n] = \sum_{n=1}^{\infty} E[\mathscr{R}_n]$$

In addition, by (4.18), $E[\mathscr{R}_n] \leq \frac{1}{\lambda_D} P(N \geq n)$.
Now, the proof follows analogously to that of Theorem 3. □

4.5. MONTE CARLO

The Markov chain in Section 4.4 has a large number of states. Most of these states have a very small probability of ever being reached. However, we can not simply ignore these states, as the aggregate of their probabilities is relatively high. This is a perfect use case for Monte Carlo simulations.

We propose an efficient Monte Carlo method, based on the results from Section 4.4, for estimating $P(N \geq n)$, $E[N]$, $E[M_{\text{min}}]$, and $E[T_{\text{total}}]$. The method is given in detail in Fig. 4.2. The main idea is to simulate many sequences of successive disasters simultaneously, and cut off these sequences when the error bounds on the values of interest are small enough. As all sequences are cut off after the same number n of successive disasters, we only allow transitions to subsequent disaster states and keep track of the probability of reaching the absorbing state separately. This allows us to closer estimate the values of interest.

In essence, we approximate the lower bounds described in Section 4.4. By Theorems 3 to 5, these lower bounds, combined with $P(N \leq n)$, give us the upper bounds as well. The method can be tuned with respect to two values: Stopping condition β determines the maximum difference between the approximated bounds, while the number of simulations η can be adjusted to affect the accuracy of the approximation of the bounds themselves. When the probability of successive disasters is too high, lowering β can keep computation times manageable by reducing the number of successive disasters taken into account.

4.6. EXPERIMENTS

To demonstrate our methods, we apply them to a version of the Sinet topology (Fig. 4.3) from the Topology Zoo [156], where all nodes without geographical information have been removed. This backbone network of 47 nodes connected by 49 bidirectional links is located in Japan, and hence is vulnerable to a variety of different disasters such as earthquakes, landslides, and typhoons. All experiments are performed on an Intel Xeon Processor E5-2620 v3.

4.6.1. DATASET

We create a set of disasters $\mathscr{D}^*$ by combining datasets from two sources: (1) the Japan Seismic Hazard Information Station (J-SHIS) [88] and (2) the International Best Track Archive for Climate Stewardship (IBTrACS) [168].

input: Number of simulations η, and bound β
output: $\hat{P}(N \geq n)$, $\hat{E}[N]$, $\hat{P}(M_{\min} \leq m)$, $\hat{E}[M_{\min}]$, and $\hat{E}[T_{\text{total}}]$
Let $\text{State}_{i,j}$ be the network state in simulation i after the jth disaster
$\hat{P}(N \geq 1) \leftarrow 1$
$\hat{P}(M_{\min} \leq m) \leftarrow 0$
for $i = 1$ to $i = \eta$ **do**
    Sample starting state $\text{State}_{i,1}$ from S_1
    $P_{i,1} \leftarrow 1$
    $M_{i,1} \leftarrow M(\text{State}_{i,1})$
    if $M_{i,i} \leq m$ **then**
        $\hat{P}(M_{\min} \leq m) \leftarrow \hat{P}(M_{\min} \leq m) + \frac{1}{\eta}$
    end if
end for
$n \leftarrow 1$
while $\hat{P}(N \geq n) > \beta$ **do**
    $n \leftarrow n + 1$
    for $i = 1$ to $i = \eta$ **do**
        $P(S_n = \emptyset) \leftarrow \exp(-\lambda_D R_{\text{State}_{i,n-1}})$
        $P_{i,n} \leftarrow P_{i,n-1}(1 - P(S_n = \emptyset))$
        Sample next disaster occurrence time T_n, conditioned on $T_n < R_{\text{State}_{i,n-1}}$
        Sample next disaster
        Compute $\text{State}_{i,n}$, given occurrence time T_n
        $M_{i,n} \leftarrow \min(M_{i,n-1}, M(\text{State}_{i,n}))$
        if $M_{i,n-1} > m$ and $M_{i,n} \leq m$ **then**
            $\hat{P}(M_{\min} \leq m) \leftarrow \hat{P}(M_{\min} \leq m) + \frac{1}{\eta} P_{i,n}$
        end if
    end for
    $\hat{P}(N \geq n) \leftarrow \frac{1}{\eta} \sum_{i=1}^{\eta} P_{i,n}$
end while
$\hat{E}[N] \leftarrow \sum_{j=1}^{n} \hat{P}(N \geq n)$
$\hat{E}[M_{\min}] \leftarrow \sum_{i=1}^{\eta} \sum_{j=1}^{n-1} P_{i,j} \exp(-\lambda_D R_{\text{State}_{i,j}}) M_{i,j}$
$\hat{E}[M_{\min}] \leftarrow \frac{1}{\eta} \hat{E}[M_{\min}] + \frac{1}{\eta} \sum_{i=1}^{\eta} P_{i,n} M_{i,n}$
$\hat{E}[T_{\text{total}}] \leftarrow \frac{1}{\eta \lambda_D} \sum_{i=1}^{\eta} \sum_{j=1}^{n} P_{i,j}(1 - \exp(-\lambda_D R_{\text{State}_{i,j}}))$

Figure 4.2: Monte Carlo method for estimating $P(N \geq n)$, $E[N]$, $P(M_{\min} \leq m)$, $E[M_{\min}]$, and $E[T_{\text{total}}]$.

EARTHQUAKE DATA (J-SHIS)

We create a disaster process d for each earthquake scenario. The affected components a_d of each scenario are the set of network components that intersect (or lie within) one or more grid cells with a seismic intensity larger than or equal to 5.5. The disaster rates

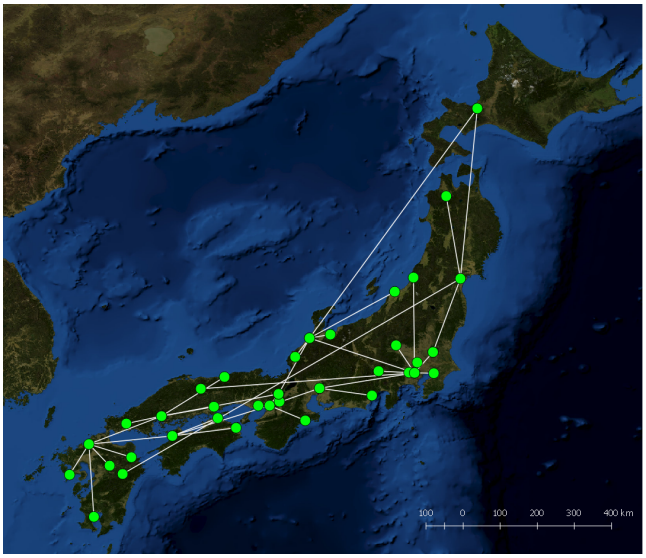


Figure 4.3: Sinet Topology.

λ_d are the inverse of the mean recurrence intervals of each fault, divided by the total number of scenarios of the fault.

TROPICAL CYCLONE DATA (IBTRACS)

IBTrACS is a collection of tropical cyclone data from numerous agencies maintained by the National Centers for Environmental Information (NCEI) of the (U.S.) National Oceanic and Atmospheric Administration (NOAA) [168]. In our experiments, we use IBTrACS beta version 4 and limit ourselves to cyclones from 1980 to 2017. We filter out any storms that never reached wind speeds of 74 mph, leaving us with a set of 1649 historical storms. As disaster area, we would prefer to use the regions that reached 74 mph winds. Unfortunately, this information is only available for some storms (in the form of the radius maximum extent per quadrant). Therefore, we apply the concept of the hurricane strike circle instead.

A strike circle is a circle with diameter 231.5 km, centered 23.15 km to the right of the hurricane center (based on its direction of motion). It is meant to depict the typical extent of hurricane force winds [169].

For each typhoon-level storm, we find the first registered center point p_a where the storm had a maximum sustained wind speed of at least 74 mph, as well as the last center point p_b with at least 74 mph maximum sustained wind speed. Then, we select the range of center points from p_a up to and including the first registered center point after p_b. Connecting these points forms a track. a_d is the selection of all components within or intersecting a strike circle of any point (including points on the line segment between registered center points) on this track. The resulting set of disasters includes many storms that do not affect any components of Sinet (e.g., hurricanes striking the U.S.). However, this is not an issue, as empty a_d are filtered out when generating $\mathscr{D}$.

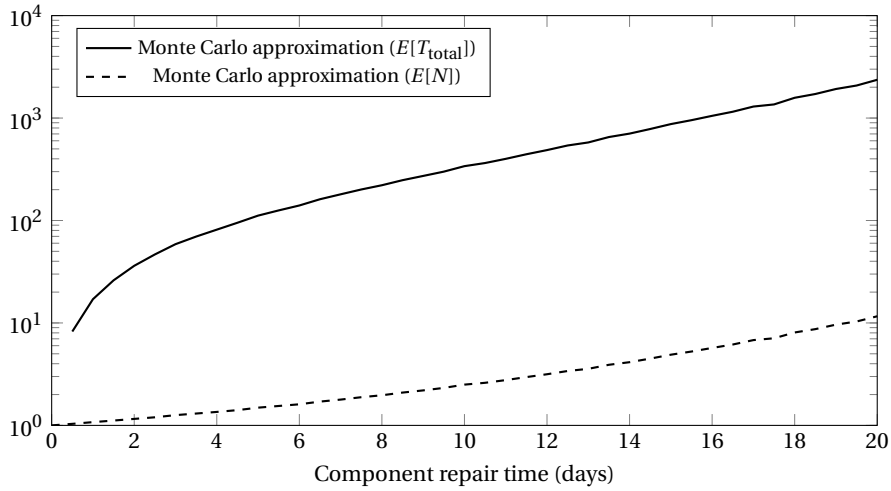


Figure 4.4: Approximations of the expected number of successive disasters, $E[N]$, and the expected time to full recovery, $E[T_{\text{total}}]$, against the component repair time.

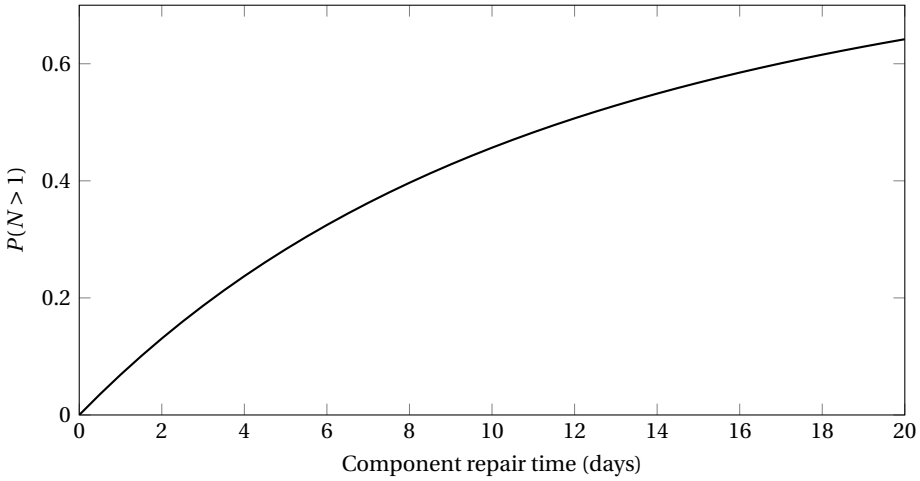


Figure 4.5: The probability of a successive disaster during recovery of the first disaster, $P(N > 1)$, against the component repair time. Exact.

The final set $\mathscr{D}^*$ is the union of the earthquake scenarios and historical tropical cyclones. This set of 2304 potential disasters can be reduced to a set $\mathscr{D}$ of 160 unique scenarios affecting Sinet. The total rate λ_D of these scenarios is 1.648 per year.

4.6.2. THE EFFECT OF COMPONENT REPAIR TIME

We first examine the effect of repair time. In a one-disaster scenario, the relation between component repair time and total repair time is simple: Ignoring start-up time, if

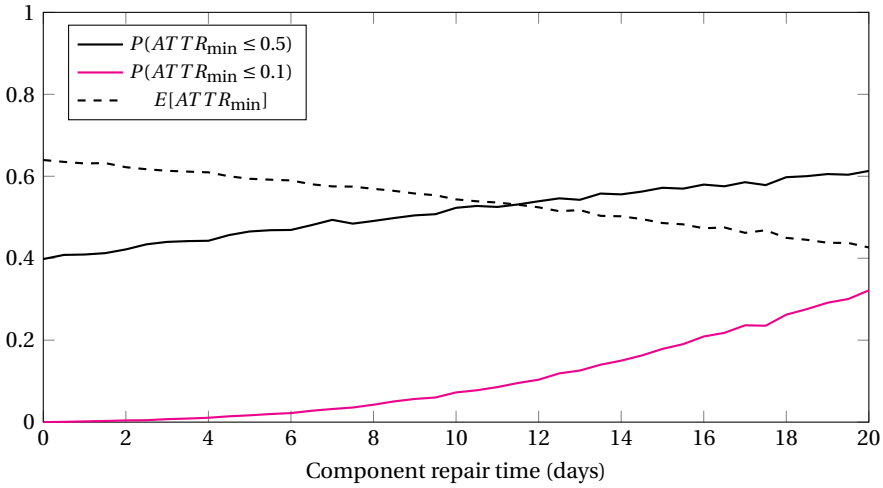


Figure 4.6: Approximations of $P(\text{ATTR}_{\min} \leq 0.5)$, $P(\text{ATTR}_{\min} \leq 0.1)$, and $E[\text{ATTR}_{\min}]$ against the component repair time. Computed by Monte Carlo simulations.

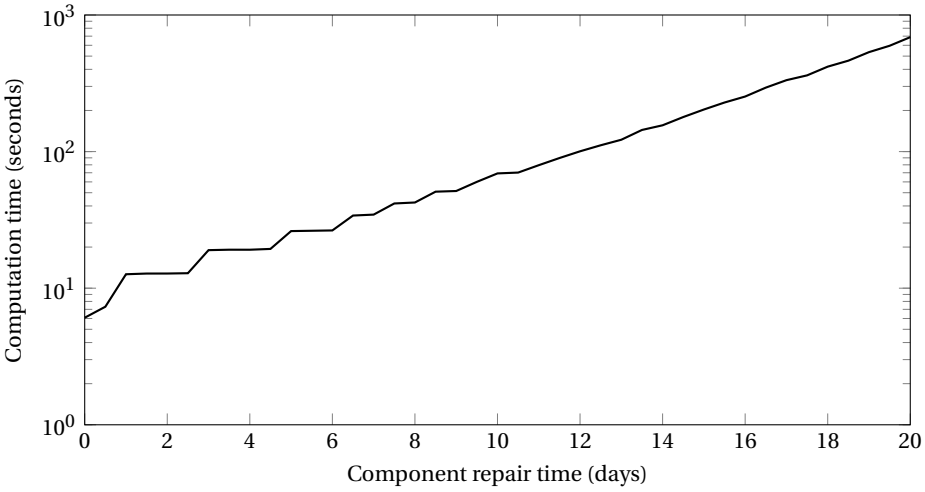


Figure 4.7: Computation time of the Monte Carlo approximations against the component repair time.

repairing components takes twice as long, the total time to full recovery will also take twice as long. However, if we take the possibility of multiple disasters into account, we encounter another effect of repair time: When the time to repair the network increases, so does the probability that the network will be struck by a subsequent disaster during recovery. These successive disasters further increase the expected total recovery time on top of the increase in component repair time itself. Our experiments show this effect can be significant.

We consider a situation where components are repaired one-by-one, using a greedy strategy that tries to maximize the number of connected node-pairs. We vary the time it takes to repair a component between 0 and 20 days. As we would need to compute a large number of steps of the DTMC to get precise results for higher repair times, we approximate all results. We use $\eta = 10,000$ simulations for each Monte Carlo approximation and set $\beta = 0.05$.

The expected number of successive disasters and the expected time to full network recovery are plotted in Fig. 4.4. $E[N]$ rapidly increases with the (component) repair time. Although, as could be expected, for more reasonable repair times[3] $E[N]$ remains below 2. Due to the influence of successive disasters, $E[T_{\text{total}}]$ grows exponentially in the component repair time.

Fig. 4.5 shows the probability of a subsequent disaster during recovery of the first disaster, $P(N > 1)$. This value can be computed exactly by computing one step of the DTMC. Interestingly, even with a component repair time of less than 5 days, the probability of facing more than 1 disaster is relatively high. Probabilities of around 0.2, or even 0.1, are significant enough to stop ignoring the possibility of successive disasters.

Next, we consider the connection between repair time and network performance. To do so, we analyze the minimum value of the Average Two-Terminal Reliability (ATTR) survivability metric in the period after the first disaster strikes and before all damage has been repaired. In Fig. 4.6, we have plotted $E[\text{ATTR}_{\text{min}}]$ against the component repair time. While the repair time does affect the expected minimum ATTR, this effect is much smaller than that on the expected time to full recovery.

A similar outcome can be observed when computing the probability that at most half of all node pairs remain connected (Fig. 4.6). However, while $P(\text{ATTR}_{\text{min}} \leq 0.5)$ increases relatively slowly with the repair time, $P(\text{ATTR}_{\text{min}} \leq 0.1)$ increases much faster.

Fig. 4.7 shows the computation time of the Monte Carlo method (parallelized to 11 threads) against the component repair time. The computation time grows exponentially in the component repair time, as the method has to simulate longer sequences of disasters to satisfy the stopping condition. Nevertheless, even for unrealistically large $E[N]$ and $P(N \geq 1)$, the computation time is more than manageable.

The repair time has a significant effect on both the total recovery time and ATTR during the recovery process. Thus, reducing it, by repairing more components at once or by decreasing the time it takes to repair individual components, should be a high priority.

4.6.3. CONCURRENT REPAIR

To evaluate our methods, we consider a use-case in which multiple components can be repaired simultaneously. In addition, we assume only nodes are damaged by the disasters. As Sinet is a backbone network and individual nodes are connected to many additional network components (which will also be affected by the disaster) that are not included in our topology, we assume repairing a single node takes half a month. However, by sending out multiple repair crews, 10 nodes can be repaired simultaneously.

[3] When components are repaired non-concurrently and the component repair time is 20 days, it can take more than 5 years to fully repair Sinet.

Table 4.1: Comparison of the exact results from Section 4.4 and the results of the Monte Carlo method from Section 4.5. The exact column shows the lower and upper bounds of the value. The runtime of the exact computation only includes the time to compute S_1 to S_k. The Monte Carlo approximation is obtained by performing 50,000 Monte Carlo simulations with stopping condition $\hat{P}(N \geq n) \leq 0.0001 \frac{exp(-\lambda_D R)}{1-exp(-\lambda_D R)}$.

	Exact	**Monte Carlo**
$E[N]$	1.0850 - 1.0851	1.0851
$P(N > 1)$	0.0763	0.0763
$P(\text{ATTR}_{\text{min}} \leq 0.5)$	0.3834 - 0.3834	0.3825
$P(\text{ATTR}_{\text{min}} \leq 0.1)$	0.0021 - 0.0022	0.0021
$E[T_{\text{total}}]$ (days)	19.4576 - 19.4674	19.4816
1-Threaded Computation Time (s)	1,556.7398	120.1504

To compute exact lower and upper bounds of the properties of interest, we construct the DTMC up to 5 successive disasters. By applying the methods from Section 4.4 and limiting ourselves to 5 successive disasters, we obtain lower bounds of $E[N]$, $P(\text{ATTR}_{\text{min}} \leq m)$, and $E[T_{\text{total}}]$. By computing the upper bound on the error, applying Theorems 3, 4, and 5, we can obtain the upper bounds on these values as well.

We approximate the lower bounds of these properties with our Monte Carlo method from Section 4.5. We set the number of simulations η to 50,000, and choose β such that the difference between the approximation of the lower and upper bounds of $E[N]$ is smaller or equal to $\epsilon = 0.0001$. That is, the method stops if $\hat{P}(N \leq n) \leq \beta = 0.0001 \frac{exp(-\lambda_D R)}{1-exp(-\lambda_D R)}$. The resulting values can be found in Table 4.1.

The computation time of the Monte Carlo method is much lower than that of the exact bounds. In addition, the Monte Carlo approximations are quite accurate. Thus, this method can be a good alternative to the exact approach, especially when the network or repair times are very large.

The probability of a second disaster striking the network during repair of a previous disaster has a low, yet still significant, probability (0.0763), but a very high impact. It is disastrous to the network if more than 90% of all node pairs lose their connection. While this outcome is not even considered feasible when only considering a single disaster, our successive disaster model shows that it is possible, although with low probability.

4.7. RELATED WORK

The amount of research into assessing the impact of multiple regional failures on communication networks is rather sparse. In [104], disaster occurrences were characterized by independent Poisson processes. However, in contrast to our framework, the methods of [104] did not consider the difference between single or multiple disaster occurrences in a short period of time.

When the possibility of more than one regional failure is considered, it is often in the form of deliberate, simultaneous attacks. In this case, the goal is to find a set of attack locations where the damage to the network is maximized [22, 23, 170] or to compute the

minimum number of regional failures required to disconnect two nodes [19, 20, 171].

In [60], Neumayer and Modiano showed how to compute the average two-terminal reliability after a randomly located disk or line cut. They briefly discussed how to extend their approach to multiple simultaneous events.

Regional failures can be modeled as Shared-Risk Link Groups (SRLG). SLRGs reflect possible combinations of links that can fail simultaneously, for example due to disasters or cable cuts. Yang et al. considered the problem of finding a set of at most k paths with an availability of at least δ under, potentially multiple simultaneous, single link failures and SRLG failures [75]. As this problem is NP-hard, they provided both a heuristic and an integer non-linear program formulation to find these paths.

Rahnamay-Naeini et al. proposed a model for multiple correlated random disasters, based on spatial point processes [69]. Using their model, Monte Carlo simulations can be performed by randomly generating a fixed number of disaster events and their effects. The model from [69] does not account for network repair and does not consider disaster processes over time.

Heegaard and Trivedi considered the recovery of a network after a single pre-selected disaster [82]. They proposed a detailed model of the performance of a network directly after the failure event and during subsequent recovery operations.

To the best of our knowledge, none of the work on multiple regional failure events considers time or network repair.

4.8. CONCLUSION

Recently, natural disasters have struck the same area shortly after one another on a number of occasions. Successive disasters like these are rare, but can inflict a massive amount of damage on the network. Consequently, the risk of successive disasters is significant and should be considered when assessing the resilience of a network.

To this end, we have composed a network and disaster model capable of modeling a sequence of disasters in time and applied this model to construct a discrete-time Markov chain of the network state after one or more successive disasters. We have shown how to adopt this Markov chain to compute with arbitrary precision (1) the probability of more than one successive disaster, (2) the expected number of successive disasters, and (3) the expected time to fully recover from these disasters. Analogously to the survivability metrics in single-disaster models, we considered the minimum value of a metric during the disaster-and-recovery process.

Building upon these results, we have developed a Monte Carlo method that can compute the vulnerability of networks to successive disasters in a matter of minutes. Since these types of analyses only need to be conducted sporadically and can be done well in advance, this computation time can be considered to be very fast.

We have applied our model to empirical disaster data. Our experiments show that when considering successive disasters, the expected time to complete recovery grows exponentially in the time it takes to repair a network component. Additionally, the probability of a second disaster striking the network during recovery can be significant, even for short repair times. Combined, our single-disaster and multi-disaster frameworks give stakeholders the ability to efficiently conduct extensive, data-driven assessments of the resilience of their networks to natural disasters.

5

A Moment of Weakness: Protecting Against Targeted Attacks Following a Natural Disaster

By targeting communication and power networks, malicious actors can considerably disrupt our society. As networks are more vulnerable after a natural disaster, this moment of weakness may be exploited to disrupt the network even further. The potential impact and mitigation of such a follow-up attack have yet to be studied.

In this chapter, we extend our multi-disaster framework to analyze the impact of a combination of a natural disaster followed by a targeted single node failure. We apply this framework on empirical disaster data and two network topologies. Our experiments show that even small targeted attacks can significantly augment the already grave network disruption caused by a natural disaster. We further show that this effect can be greatly mitigated by adopting a repair strategy that actively takes the possibility of targeted attacks into account.

Parts of this chapter have been published in SIGMETRICS Perform. Eval. Rev. **47**, 4 (March 2020) [172].

5.1. INTRODUCTION

Communication and power networks are critical to our society. This makes them a prime target for malicious actors trying to destabilize or terrorize a country. In fact, in 2016, the director of the NSA, warned, "It's only a matter of the when, not the if, you are going to see a nation state, a group or an actor engage in destructive behavior against critical infrastructure of the United States" [173].

For many of these actors, the exact timing of their attack may not be essential. Their focus is to inflict as much damage as possible, preferably using only a small amount of resources. A strategy they might adopt is to delay their attack until the network is most vulnerable, such as after a natural disaster. By attacking the network at its weakest moment, a bad actor can multiply the damage he, or the disaster by itself, could otherwise inflict.

As it takes time to prepare and execute an attack, a network operator has a limited window to try to reduce the impact of any potential attack. However, to the best of our knowledge, while there is a large body of research on the resilience of networks to natural disasters and targeted attacks (e.g. [174–176]), the potential combination of a disaster followed by a targeted attack has yet to be studied.

In this chapter, we propose a framework to analyze the impact of a combination of a natural disaster and targeted single node failure[1]. Our main contributions are as follows:

- We extend our successive disaster framework from Chapter 4 to incorporate *targeted* attacks.
- We apply our framework to empirical disaster data and show that a small follow-up attack can significantly increase the impact of a natural disaster.
- We study the effect of changing the repair strategy to prepare for potential follow-up attacks, and demonstrate that the impact of follow-up attacks on networks can be greatly mitigated, at almost no cost to network performance, by making calculated modifications to the order in which network components are repaired.

5.2. FRAMEWORK

In Chapter 4, we introduced a model and framework for assessing the resilience of networks to successive disasters, taking into account network recovery. In this chapter, we extend this framework to include the risk of targeted attacks.

As a reminder, we model the network as a directed multigraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \psi)$ with nodes $v \in \mathcal{V}$ connected by links $e \in \mathcal{E}$, where $\psi : \mathcal{E} \to \mathcal{V} \times \mathcal{V}$ and $e \in \mathcal{E}$ connects v_1 to v_2 if and only if $\psi(e) = (v_1, v_2)$. We define the *network state* s of this network by its failures: network component $c \in \mathcal{V} \cup \mathcal{E}$ is functioning if and only if $c \notin s$.

Now, to be able to assess the resilience of this network to follow-up attacks, we first need to model the impact of a disaster itself. We assume disaster occurrences are Poissonian, and we are given a multiset $\mathcal{D}^*$ of disaster processes $d = (a_d, \lambda_d)$, where $a_d \subseteq \mathcal{V} \cup \mathcal{E}$ are the components affected by d and λ_d is the rate of d. Thus, if disaster $d \in \mathcal{D}^*$ occurs at time t, when the network state is s, the new network state at time t will be

[1] Be it through physical means or cyber attacks.

$s \cup a_d$. As the combination of multiple Poisson processes is itself Poissonian, the disaster processes in $\mathscr{D}^*$ can be combined as follows:

$$\mathscr{D} = \{(a_d, \lambda_d) | a_d \neq \emptyset \wedge \lambda_d = \sum_{(a_d, \lambda_{d^*}) \in \mathscr{D}^*} \lambda_{d^*} > 0\} \tag{5.1}$$

In this chapter, we only consider attacks after a single disaster. In other words, we assume a single disaster occurs and is then followed by an attack on the network. However, by applying the techniques of Chapter 4, our model can be easily generalized to capture an attack after an arbitrary number of successive disasters or any other mix of natural disasters and attacks. We fix the time of the initial disaster (D_1) at $T_1 = 0$. Now, we can compute the distribution of the network state at T_1, S_1, by

$$P(S_1 = s) = \sum_{d \in \mathscr{D} | a_d = s} P(D_1 = d) = \sum_{d \in \mathscr{D} | a_d = s} \frac{\lambda_d}{\lambda_D} \tag{5.2}$$

where λ_D is $\sum_{d \in \mathscr{D}} \lambda_d$.

A follow-up attack after a disaster can be pre-planned or opportunistic. In either case, it will take some time to react to the disaster and execute the attack. We consider two different attack models: (1) the attack occurs after a fixed amount of time t_{attack}, or (2) the time between the disaster and attack is exponentially distributed with rate λ_{attack}. In both cases, if the network has been fully repaired before the attack has been executed, we assume the attack will be canceled and the network will not suffer any further damage.

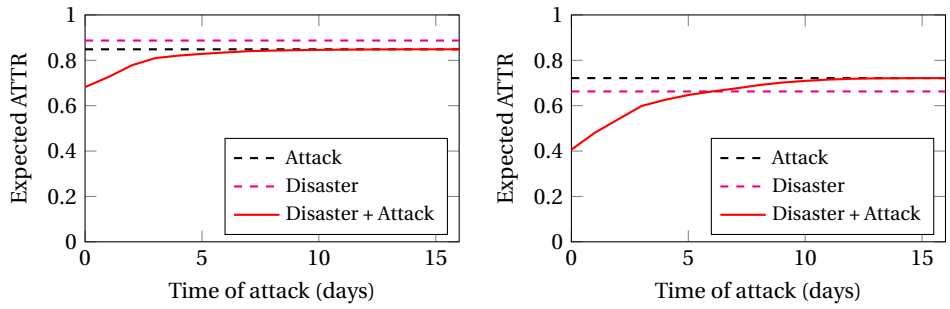
Let T_{attack} be the time of the attack, and let $M(s)$ be a measure over network state s. We assume the attacker has perfect knowledge of the network at all times, and will always take down the node that minimizes M at T_{attack}. In other words, an attack is modeled as a worst-case node failure.

The target and impact of this attack greatly depend on the progress of network repair at T_{attack}. We consider a deterministic repair model. That is, we assume that, given a certain starting state, the recovery of the network is fixed (until the attack occurs). For each possible starting state s, we define a repair function $r_s : \mathbb{R}^+ \to \mathscr{V} \cup \mathscr{E}$. $r(t)_s \in \mathscr{V} \cup \mathscr{E}$ is the state of the network at time $t \leq T_{\text{attack}}$, given that the state of the network after being struck by the initial disaster was $S_1 = s$. Thus, the state of the network just before the attack is $r_{S_1}(T_{\text{attack}})$.

Let $R_s := \min\{t \geq 0 | r(t)_s = \emptyset\}$ be the time it takes to fully repair the network (assuming no attack occurred beforehand). Given M, we can consider the follow-up attack as a function $att : \mathscr{V} \cup \mathscr{E} \to \mathscr{V} \cup \mathscr{E}$ from the network state just before the attack to the network state just after the attack:

$$att(s) = \begin{cases} \emptyset & \text{if } s = \emptyset \\ s \cup \operatorname{argmin}_{v \in \mathscr{V}} M(s \cup v) & \text{otherwise} \end{cases} \tag{5.3}$$

By combining our disaster, repair, and attack models, we can now directly compute the distribution of the state S_{attack} of the network just after the attack. In the fixed attack



(a) Deltacom (b) Sinet

Figure 5.1: Impact of a follow-up attack. The disaster occurs at t=0, and is followed by a targeted attack after 0 to 16 days (even if the network is already repaired). One node is repaired every day.

time case, the distribution of S_{attack} is given by

$$P(S_{\text{attack}} = s) = \sum_{d \in \mathscr{D} | s = att(r_{a_d}(t_{\text{attack}}))} \frac{\lambda_d}{\lambda_D} \tag{5.4}$$

while in the random attack time case the distribution of S_{attack} is given by

$$\begin{aligned} P(S_{\text{attack}} = s) &= \sum_{d \in \mathscr{D}} \frac{\lambda_d}{\lambda_D} P(S_{\text{attack}} = s | D_1 = d) \\ &= \sum_{d \in \mathscr{D}} \frac{\lambda_d}{\lambda_D} (\exp(-\lambda_{\text{attack}} \min(\mathscr{M}_{a_d,s}, R_{a_d})) \\ &\qquad - \exp(-\lambda_{\text{attack}} \min(\mathscr{S}_{a_d,s}, R_{a_d}))) \end{aligned} \tag{5.5}$$

where $[\mathscr{M}_{a_d,s}, \mathscr{S}_{a_d,s})$ is the period of time during which an attack would result in network state s[2].

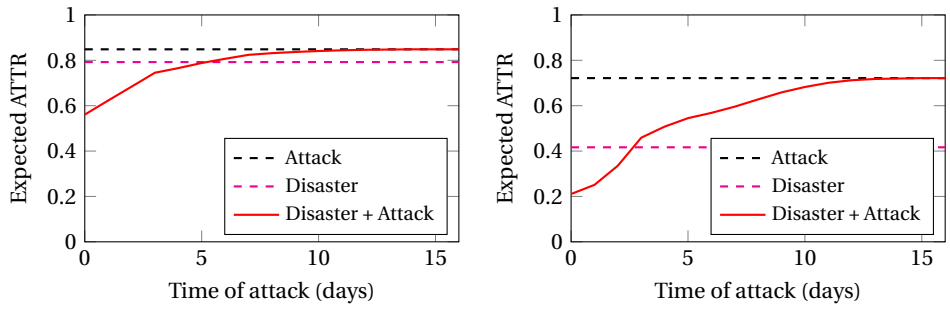
Given the distribution of S_{attack}, we can directly compute the distribution of any performance metric after the follow-up attack, such as the number of remaining connections, $M(S_{\text{attack}})$. In addition, our framework allows network operators to assess the impact of different repair strategies or network configurations by simply exchanging repair functions or modifying the initial network.

5.3. EXPERIMENTS

In this section, we apply our framework to two slightly modified[3] versions of undirected networks from the topology zoo [156]: Sinet and Deltacom. Sinet is a Japanese network of 47 nodes connected by 49 links, and Deltacom is a US network of 99 nodes connected by 151 links.

[2] $\mathscr{M}_{a_d,s}$ is the first time t at which $att(r_{a_d}(t)) = s$ (or ∞ if no such time exists), and $\mathscr{S}_{a_d,s}$ is the first time t after $\mathscr{M}_{a_d,s}$ at which $att(r_{a_d}(t)) \neq s$ (or ∞).

[3] We have removed all nodes without a geographical location or with degree 0.



(a) Deltacom (b) Sinet

Figure 5.2: Impact of a follow-up attack on Deltacom (Sinet) if the attacker waits for a disaster that damages at least 5 (10) nodes. The disaster occurs at t=0, and is followed by a targeted attack after 0 to 16 days (even if the network is already repaired). One node is repaired every day.

We make use of the same disaster set as was used in Chapter 4: a set of earthquake scenarios and historical tropical cyclones. For Sinet, we consider both types of disasters, while for Deltacom, we only consider tropical cyclones. We assume only network nodes are affected by these disasters and all network links remain functioning. This gives us a yearly disaster rate λ_D of 1.597[4] for Sinet and 1.342 for Deltacom.

For ease of reading, we make the assumption that one node is repaired every day. However, by scaling both the attack and repair time, our results can easily be transformed to any other repair time.

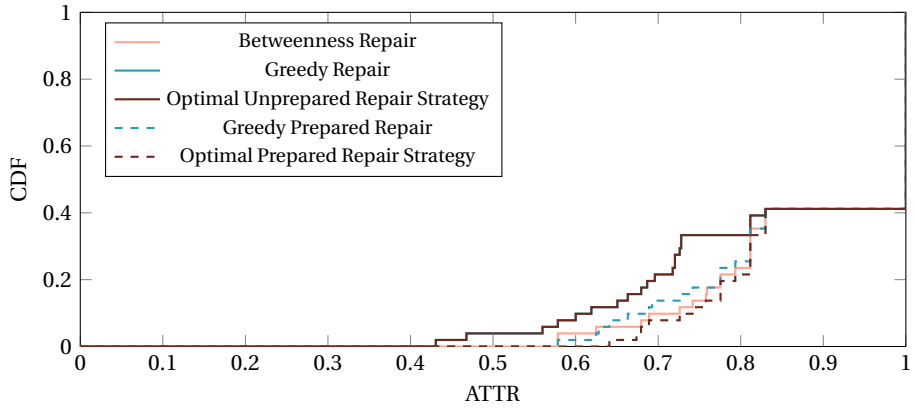
5.3.1. IMPACT

We first compare the impact of follow-up attacks to those of a disaster or attack by itself. We assume both networks use a greedy repair function that continuously chooses the node with the largest impact on ATTR to repair. To make a fair comparison, we modify the attack function att by continuing the follow-up attack even if the network has been fully repaired.

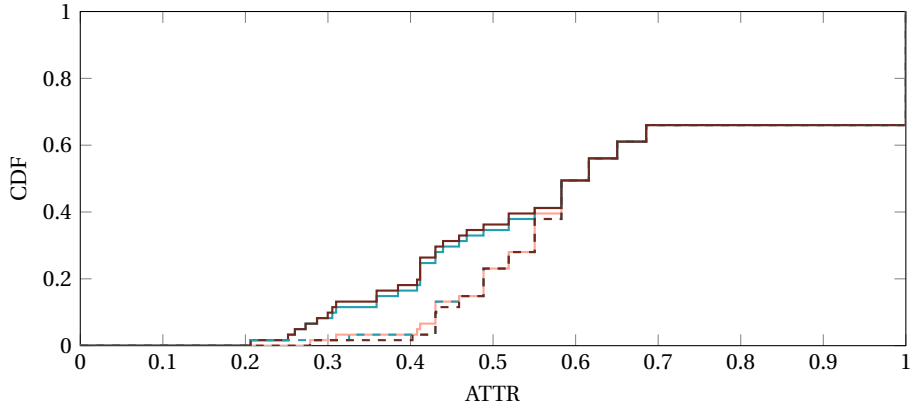
Fig. 5.1 shows the expected ATTR after a targeted attack, disaster, or disaster and follow-up attack. Sinet is clearly more vulnerable to both targeted attacks and disasters than Deltacom. However, for both networks, a follow-up attack can significantly increase the impact of a disaster. For Sinet, the combination of disaster and follow-up attack disconnects more than half of all node pairs on average.

We have assumed that a malicious actor strikes after the *first* natural disaster that hits the network. On average, such an opportunity occurs more than once per year. However, he could also decide to wait for a larger disaster, which would allow him to inflict even more damage to the network. We consider an attacker that waits for a disaster that damages at least 5 (10) of Deltacom's (Sinet's) nodes[5]. Fig. 5.2 shows the impact of this more patient follow-up attack. Waiting for larger events allows the attacker to inflict

[4] Lower than in Chapter 4, since we now exclude disasters that only strike links.
[5] An opportunity that occurs around once every 2 years on average.



(a) Deltacom



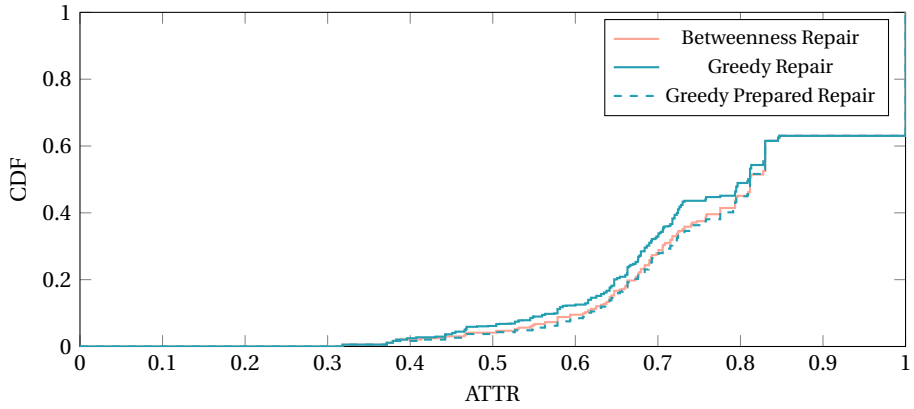
(b) Sinet

Figure 5.3: Cumulative Distribution Function of the ATTR after the follow-up attack for different repair strategies. The disaster occurs at t=0, and is followed by a targeted attack after 3 days. If the network is repaired before the attack, the attack is canceled (and ATTR = 1).

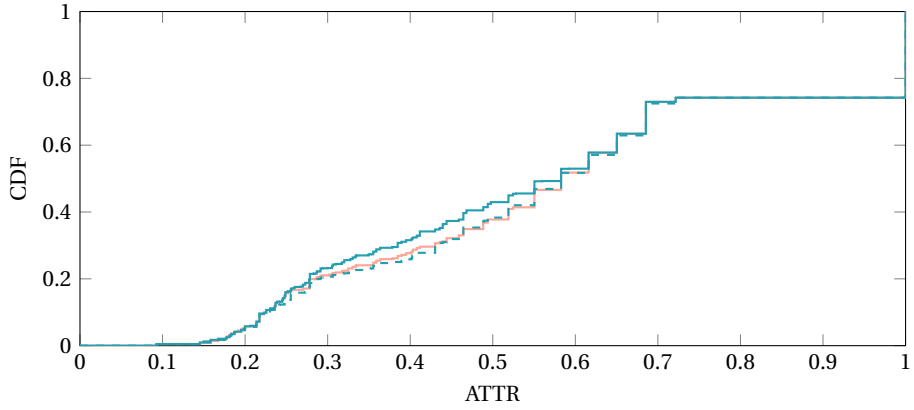
much more damage. In the case of Sinet, the expected impact of the follow-up attack is the disconnection of around half of the *remaining* node pairs.

5.3.2. REPAIR STRATEGIES

After a disaster, the network operator will typically try to restore as much functionality as quickly as possible. The impact of the follow-up attack greatly depends on the progress of these repair operations at the time of the attack. Although speeding up repair would have the largest effect, the network operator can also change the order in which components are repaired to try to minimize the impact of any attacks. While this might lower the speed at which network functionality is restored, it could be a worthy trade-off if the network is under threat.



(a) Deltacom

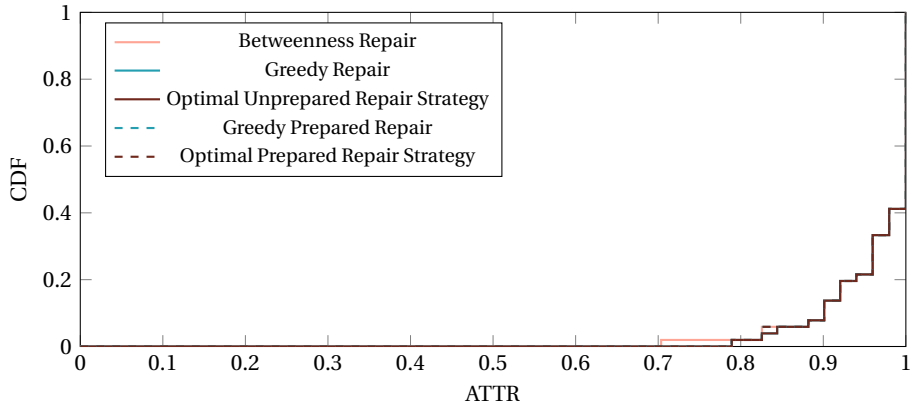


(b) Sinet

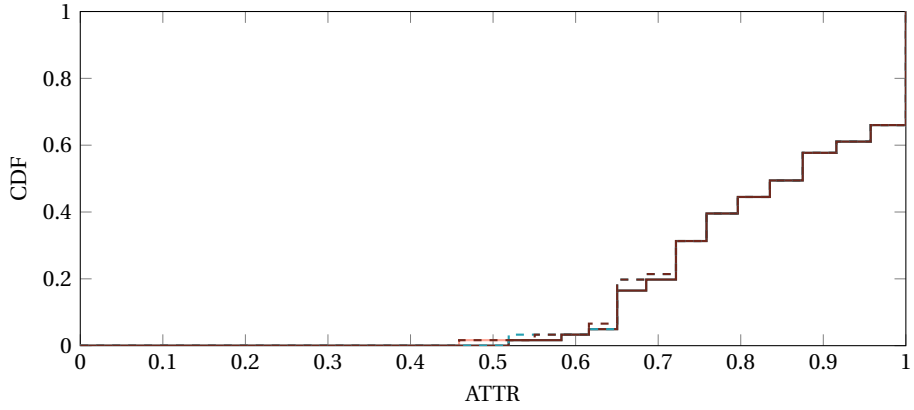
Figure 5.4: Cumulative Distribution Function of the ATTR after the follow-up attack for different repair strategies. The disaster occurs at t=0, and is followed by a targeted attack after a random exponentially distributed delay of on average 3 days. If the network is repaired before the attack, the attack is canceled (and ATTR = 1).

In this section, we consider the effect of changing the node repair order on the impact of the follow-up attack. We compare 5 different repair strategies:

- *Betweenness Repair*: Repair nodes in the order of their betweenness centrality [177].
- *Greedy Repair*: Every day, repair the node with the highest impact on the ATTR.
- *Optimal Unprepared Repair Strategy*: Maximizes the ATTR at the time of attack.
- *Greedy Prepared Repair*: Every day, repair the node that would increase the ATTR the most if the network would be attacked immediately afterwards.
- *Optimal Prepared Repair Strategy*: Maximizes the ATTR immediately after the attack.



(a) Deltacom



(b) Sinet

Figure 5.5: Cumulative Distribution Function of the ATTR 3 days after the initial disaster (without follow-up attack) for different repair strategies.

Fig. 5.3 shows the Cumulative Distribution Function (CDF) of the ATTR after a follow-up attack with a fixed t_{attack} of 3 days. The large spikes at ATTR = 1 show the probability of completely repairing the network within 3 days (0.340 for Sinet and 0.588 for Deltacom). In these cases, the attack is canceled and the repair strategy has no impact. However, in most other cases the repair strategy does significantly impact the ATTR after a follow-up attack. In particular, there is a large gap between the strategies that try to maximize the ATTR by itself compared to those that try to maximize the ATTR after the attack. Experiments on randomly delayed follow-up attacks show similar, albeit less pronounced, results to those of fixed-time attacks (see Fig. 5.4).

Since changing the order of repair can reduce the impact of follow-up attacks, one might wonder what the impact of these prepared repair strategies is on the performance of the network during repair without an attack. Or in other words, what does preparing

for a follow-up attack cost us if no such attack occurs? Fig. 5.5 shows the CDF of the ATTR 3 days after the initial disaster (without any follow-up attack). For the considered networks, the difference between the different repair strategies is extremely small and even the prepared strategies perform close to optimally.

5.4. CONCLUSION

Critical infrastructure networks are prime targets for malicious actors trying to destabilize or terrorize a country. As part of their attack strategy, they might delay their attack until critical infrastructure is significantly more vulnerable, for example right after a natural disaster has struck. However, current disaster vulnerability frameworks do not consider the potential risk of these follow-up attacks to a network.

We have extended our successive-disaster framework from Chapter 4 with the ability to analyze the impact of follow-up attacks. The extended framework can take into account a variety of natural disasters and two kinds of attacks: a worst-case node failure after (1) a fixed amount of time or (2) an exponentially distributed random delay after an initial disaster.

In our experiments, we have shown that small targeted attacks can significantly augment the impact that a natural disaster has on the network. Fortunately, our results also reveal that the right choice of repair strategy allows network operators to reduce the threat of follow-up attacks at almost no cost to network performance compared to other repair strategies. Our framework aids in determining the efficacy of repair strategies.

6

Evaluating Local Disaster Recovery Strategies

Whereas in the previous chapters we focused on assessing disaster resilience, in the remainder of this thesis we study how to improve the resilience of networks to disasters. Our methods for improving disaster resilience build upon our single-disaster framework. In this chapter, we extend our framework with the ability to evaluate various repair strategies, with the goal of improving the ability of the network operator to quickly restore network connectivity after a disaster.

We specifically focus on the possibility of temporarily replacing damaged nodes by emergency nodes. We prove that computing the optimal choice of nodes to replace is an NP-hard problem and propose several simple strategies. We evaluate these strategies on two U.S. topologies and show that a simple greedy strategy can perform close to optimal.

Parts of this chapter have been published in SIGMETRICS Perform. Eval. Rev. **46**, 2 (September 2018) [178].

6.1. INTRODUCTION

Repairing a network can take days to months, during which functionality is only slowly restored. Thus, there is a need for a simultaneous quick response to recover a bare amount of network functionality in the affected areas as quickly as possible.

In this chapter, we consider the possibility of temporarily replacing some of the failed network components by emergency equipment, such as Movable and Deployable Resource Units (MDRUs) [179]. We extend our framework from Chapter 3 with the ability to evaluate recovery strategies. The evaluation only considers the effect of the recovery on the network area enclosing the disaster region, as the focus of these recovery efforts is to restore vital network functionality to the affected area.

Using our framework, network operators can decide beforehand which strategy they want to employ, such that after a disaster the strategy can be implemented immediately.

Our main contributions are as follows:

- We propose a model (Section 6.2) and algorithm (Section 6.4) for evaluating the effectiveness of a recovery strategy.
- We describe an optimal strategy as an optimization problem (Section 6.3.1), and prove that it is NP-hard.
- As the time to determine a repair strategy is limited, and communication within and from a disaster region is disrupted, we propose alternative, simple strategies that facilitate quick, local decision-making (Section 6.3.2).
- We demonstrate our framework on two topologies, and evaluate our strategies (Section 6.5). In our example use case, a simple greedy strategy gives close to optimal results.

While there has been other work on network recovery strategies after a large-scale disaster, to the best of our knowledge we are the first to propose an evaluation framework for different strategies, as well as the first to focus on a local area enclosing the disaster region.

6.2. EVALUATION MODEL

We model a telecommunications network as an undirected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ of nodes $\mathcal{V}$ connected by links $\mathcal{E}$. The nodes of the network are the routing and computing nodes of the network, as well as its base stations, while the edges are the cables (or radio links) connecting them.

To evaluate different strategies to a wide range of possible situations, we work with a representative set of disaster scenarios $\mathscr{D}$, as was introduced in Chapter 3. These can for example be historical disasters, randomly sampled disasters, or specific scenarios created by experts.

Each disaster $d \in \mathscr{D}$ affects a region of the network, called the disaster region. We assume all nodes in the affected region fail. Links remain unaffected in our model, but our methodology is easily extended to other damage patterns as well.

We assume exactly one disasters occurs and that we are given the occurrence probability $P(d)$ of each disaster $d \in \mathscr{D}$. We use these probabilities to weigh the relative importance of each disaster to the overall evaluation of a recovery strategy.

To quickly recover the functionality of the network, damaged nodes can be replaced by temporary emergency nodes, such as MDRUs [179]. The exact functionality (e.g. base station, router, edge computing) of these nodes would depend on the node it replaces. To connect the emergency node to the rest of the network, the cable to the old node will be digged up, spliced, and connected to the emergency node. The emergency nodes can have a smaller capacity as the node they replace, as long as it at least takes over some bare minimum of its functionality.

In the case of a disaster, large amounts of manpower will be made available to recover the network. However, the number of other resources available might be more limited. As such, we assume that only K temporary nodes can be placed, but the process of placing these K nodes can be worked on simultaneously.

The time it takes to place and connect an emergency node depends on both the reachability of its intended location, as well as the properties of the area and soil around it. For example, it could take much more time to place a device on top of a mountain than on farmland. We assume we are given a cost(v) for each $v \in V$, where cost(v) is the time it takes to replace node v.

Let $S(d)$ be all nodes affected by disaster $d \in \mathscr{D}$. The choice to be made after a disaster, using a recovery strategy, is the set of at most K nodes out of $|S(d)|$ to replace.

Given such a choice of actions, the state of the network after a disaster d can be described by a vector

$$[S_k(d)]_{k=1}^{K+1} = [(\mathscr{G}_1, 0), (\mathscr{G}_2, t_2), \ldots, (\mathscr{G}_{K+1}, t_{K+1})] \tag{6.1}$$

of length $K+1$. Where $\mathscr{G}_1$ is the topology of the network directly after the disaster, i.e. the graph $\mathscr{G}$ minus the affected nodes. $\mathscr{G}_2$ is the topology of the network at time t_2, directly after the first recovery action has been completed, $\mathscr{G}_3$ is the topology of the network at time t_3, directly after the second recovery action has been completed, etc.

6.2.1. LOCAL AREA

The focus of recovery efforts is to restore vital network functionality to the local affected area. However, it is also important to consider those nodes that are disconnected by the disaster, but are not in the disaster region, and are thus still functioning. The most effective method to reconnect these nodes will be through the disaster region.

As such, we only consider the placement of emergency equipment and the effect of this equipment in a local area around the disaster region. By limiting ourselves to a smaller area, we also limit the size of the graph we need to consider when determining where to place the emergency nodes and when evaluating the effectiveness of the approach, thus reducing the amount of processing time required, and increasing the level of network details that can be considered.

Specifically, we define the local nodes $\mathscr{V}_L \subseteq \mathscr{V}$ after a disaster as the nodes of the network that are struck by the disaster ($S(d)$), or are distanced only 1 hop from such a node. Thus the local network of interest is $\{\mathscr{V}_L, \mathscr{E}_L\}$, where $\mathscr{E}_L = \{(v,x) \in \mathscr{E} \,|\, v,x \in \mathscr{V}_L\}$.

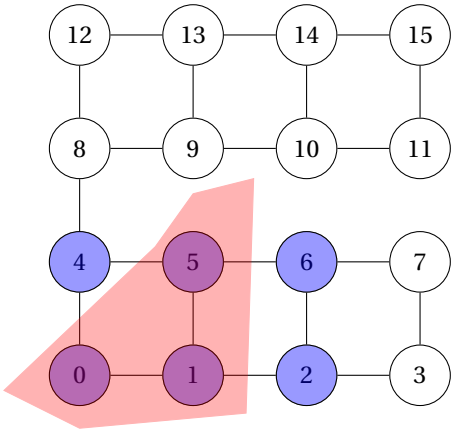


Figure 6.1: A disaster (in red) takes down nodes 0, 1 and 5. The local nodes V_L consists of the nodes in blue

As an example, consider the network and disaster region in Fig. 6.1. The disaster damages nodes 0, 1 and 5. The local nodes are $\mathcal{V}_L = \{0,1,2,4,5,6\}$. Nodes 2, 3, 6 and 7 have been cut of from the giant connected component.

6.2.2. EVALUATION METRICS

Nodes that are cut off by the disaster, but are not part of the local area, still need to be reconnected to the rest of the network. This is taken into account by increasing the weights of nodes on the border in proportion to the portion of the network they connect to the local area.

Let $p(x)$ be the weight of node x in $\mathcal{G}$. For any local node $v \in \mathcal{V}_L$, define

$$\mathscr{C}(v) := \{x \in \mathcal{V} | h(v,x) \leq h(y,x), \forall\, y \in \mathcal{V}_L\} \tag{6.2}$$

as the nodes closest to v, where $h(v,x)$ is the smallest number of hops from v to x in $\mathcal{G}$.

Now, the weight of node $v \in \mathcal{V}_L$ is set to

$$w(v) = \sum_{x \in \mathscr{C}(v)} n(x)p(x) \tag{6.3}$$

where

$$n(x) = \frac{1}{|\{v \in \mathcal{V}_L | x \in \mathscr{C}(v)|} \tag{6.4}$$

Note that $w(v) = p(v)$ for all nodes in the disaster region itself. These weights can be seen as representative for the amount of traffic demand we expect to/from the nodes.

Functioning nodes in the giant connected component will have a much higher weight than other functioning nodes, which in turn generally have a higher weight than the nodes in the disaster area. Thus, by setting these weights, we prioritize connecting areas to the core network and connecting the smaller components to the giant connected component.

For example, if all nodes v in Fig. 6.1 have $p(v) = 1$, then the weight of nodes 0, 1 and 5 would be 1. However, nodes 2 and 6 would have weight 2, and node 4 weight 9. Thus, one of the first priorities would be reconnecting nodes 2 and 6 to node 4.

Our framework can be used with any network metric. In this chapter we consider a weighted version of the Average Two-Terminal Reliability (ATTR).

Definition 2. *Weighted Average 2-Terminal Reliability (WATTR)*
Let

$$I(v,x) = \begin{cases} 1 & \text{if node } v \text{ is connected to node } x \\ 0 & \text{otherwise} \end{cases}$$

The weighted average 2-terminal reliability (WATTR) is defined as

$$WATTR := \frac{1}{W} \sum_{v \in \mathcal{V}_L} \sum_{x \in \mathcal{V}_L - \{v\}} w(v)\, w(x)\, I(v,x) \tag{6.5}$$

where $W := \sum_{v \in \mathcal{V}_L} \sum_{x \in \mathcal{V}_L - \{v_i\}} w(v)\, w(x)$.

WATTR can be seen as a measure of the proportion of potential connections in a network that are still functioning.

If we let $C \subseteq \mathcal{V}_L$ be the set of all connected components of the network in $\mathcal{V}_L$, and define $\text{sum}(c) := \sum_{v \in c} w(v)$ for all $c \subseteq \mathcal{V}_L$. Then

$$W = \sum_{v \in \mathcal{V}_L} w(v) * (\text{sum}(\mathcal{V}_L) - w(v)) \tag{6.6}$$

and

$$\text{WATTR} = \frac{1}{W} \sum_{c \in C} \sum_{v \in c} w(v) * (\text{sum}(c) - w(v)) \tag{6.7}$$

The metric evaluates the network at a specific state. To evaluate the complete emergency recovery process, we use a weight function $\mathcal{W} : \mathbb{R}^+ \to \mathbb{R}^+$ such that $\int_0^\infty \mathcal{W}(t)\,dt = 1$

We then evaluate the vector $[S_k(d)]_{k=1}^{K+1}$ after the disaster as

$$M(d) = \sum_{k=1}^{K+1} \frac{M(\mathcal{G}_k) - M(\mathcal{G}_1)}{1 - M(\mathcal{G}_1)} \int_{t_k}^{t_{k+1}} \mathcal{W}(t)\,dt, \tag{6.8}$$

where $M(\mathcal{G}_k)$ is the value of the metric (in our case WATTR) on the graph $\mathcal{G}_k$, $t_1 := 0$ and $t_{K+2} := \infty$. The value $\frac{M(\mathcal{G}_k) - M(\mathcal{G}_1)}{1 - M(\mathcal{G}_1)}$ measures the effect of the recovery operations in the local network and ranges from 0 (no effect) to 1 (full recovery). In case $s(d) = \emptyset$, i.e., the disaster does not affect the network, we define $M(d) = 1$.

6.3. RECOVERY STRATEGIES

6.3.1. OPTIMAL STRATEGY

If we let $\mathcal{V} = \{v_1, v_2, \ldots, v_{|\mathcal{V}|}\}$, and describe the choice of nodes as a vector of binary values $\mathbf{x}$ such that $x_i = 1$ if and only if v_i is replaced, then an optimal strategy is the solution to the problem

$$\max M(d|\mathbf{x}) \tag{6.9}$$

$$\text{s.t.} \sum_{i=1}^{|\mathcal{V}|} x_i \leq K \tag{6.10}$$

$$x_i = 0 \qquad \forall\ v_i \notin S(d) \tag{6.11}$$

$$x_i \in \{0,1\} \qquad \forall\ i \tag{6.12}$$

where $M(d|\mathbf{x})$ is the value of $M(d)$ given the choice $\mathbf{x}$ of nodes to replace.

Theorem 6. *When using WATTR as the evaluation metric, computing the optimal strategy is strongly NP-hard even for the 0 cost case. i.e., when repair time is not considered.*

Proof. Our proof is inspired by the proof of theorem 1 in [180].

We prove theorem 6 by giving a reduction from the well-known NP-complete SET COVER problem to the decision version of the optimization problem (with costs 0).

Note that the weight function $\mathcal{W}$ is irrelevant if all replacement costs are 0, thus, we will not include further mentions of the weight function in the proof.

The SET COVER problem can be described as follows: given a set $\mathcal{U} = \{u_1, u_{2,3}, \ldots, u_n\}$, a family $\mathcal{F} = \{F_1, F_2, F_3, \ldots, F_m\}$ of subsets of $\mathcal{U}$ s.t. $\cup_{i=1}^{m} F_i = \mathcal{U}$ and an integer $k \leq m$, is there a cover $C \subseteq \mathcal{F}$ such that $\cup_{c \in C} c = \mathcal{U}$ and $|C| \leq k$?

Given an instance of the SET COVER problem, we construct a (local) graph with nodes $\mathcal{V}_L = \{b\} \cup \mathcal{U} \cup \mathcal{F}$. That is, $\mathcal{V}_L$ consists of a base node b, a node for each element in $\mathcal{U}$ and a node for each set in $\mathcal{F}$.

We directly connect b to all nodes in $\mathcal{F}$. In addition, for all nodes $F_i \in \mathcal{F}$ we add the links $\{\{F_i, u_j\} | u_j \in F_i\}$ to $\mathcal{E}_L$. More formally, $\mathcal{E}_L = (\{b\} \times \mathcal{F}) \cup \{\{F_i, u_j\} \in \mathcal{F} \times \mathcal{U} | u_j \in F_i\}$.

The weight of all nodes in $\mathcal{F}$ is set to 0, and the weight of all other nodes to 1. We let $S(d) = \mathcal{F}$, i.e., a node F_i is in the disaster region of the disaster iff $F_i \in \mathcal{F}$.

Note that this is a valid local selection of nodes and links, as all nodes in $\mathcal{V}_L$ are within 1 hop of the failed nodes. Now, let $K = k$, the decision problem will be to determine if there exists a choice of at most K nodes of $S(d)$ to be replaced such that $M(d) = \text{WATTR}(\mathcal{G}_K)$ will be greater or equal than 1.

Suppose there is a solution to the problem instance of SET COVER. That is, there exists a $C \subseteq \mathcal{F}$ such that $\cup_{c \in C} c = \mathcal{U}$ and $|C| \leq k$. By replacing all corresponding nodes $F_i \in C$, all nodes with a weight greater than zero will be connected to each other (through b). Thus, C is also a solution to the optimal strategy instance.

Conversely, suppose there is a solution to the optimal-strategy instance. That is, we have a set C of at most K nodes in $S(d)$, such that when these nodes are replaced, the WATTR of the local network will be 1. So every node $U_i \in \mathcal{U}$ must be connected to b through at least 1 node $F_j \in C$. That is, $\forall u_i \in \mathcal{U}\ \exists F_j \in C$ s.t. $u_i \in F_j$. Or alternatively, $\cup_{c \in C} c = \mathcal{U}$. So C is also a solution to the SET COVER instance.

We have provided a (polynomial) reduction from the strongly NP-complete SET COVER problem to the decision variant of the optimal-strategy problem with costs 0. As a result, we can conclude that the optimal strategy problem for the 0 cost case is strongly NP-hard. □

As computing the optimal strategy is an NP-hard problem and there might only be a limited amount of resources available after a disaster due to the destruction and chaos, computing the optimal choice of nodes might take too much time. In addition, the choice of which nodes to replace has to be made as quickly as possible after a disaster, at which point the complete state of the network might not be known. As such, it might be preferable to make some quick decisions based on a simple rule of thumb instead.

These rules of thumb, or simple strategies, might be suboptimal for the specific situation, but give good results in general, whatever state the network might be in. In the following section, we propose several simple strategies.

6.3.2. SIMPLE STRATEGIES

We use $R \subseteq S(d)$ to indicate the nodes that will be replaced.

The basic idea of these strategies is as follows. Choose some node-metric $\mathscr{M}$, then iteratively select nodes to replace with the highest value of $\mathscr{M}$:

1. $R \leftarrow \emptyset$
2. Let $\mathscr{B} \subseteq S(d)$ be all nodes $v \in S(d)$ such that v is at most 1 hop away from (i.e., directly connected to) at least one node in $\mathscr{V}_L - S(d)$. That is, $\mathscr{B}$ is the intersection of the neighborhood of $\mathscr{V}_L - S(d)$ and $S(d)$. We want to limit ourselves to only replacing nodes in $\mathscr{B}$, as otherwise we would replace nodes without connecting them to a connected component.
3. Pick a $v \in \mathscr{B} - R$ such that $\mathscr{M}(v) \geq \mathscr{M}(y) \forall y \in \mathscr{B} - R$.
4. $R \leftarrow R \cup \{v\}$
5. $\mathscr{B} \leftarrow \mathscr{B} \cup \{y \in S(d) | \{v, y\} \in \mathscr{E}_L\}$
6. If $|R| < K$ and $|R| < |S(d)|$, repeat steps 3-6

We consider 4 node-selection strategies:

- Greedy, that is, pick the node that has the largest effect on M: $\mathscr{M}(v) := M(d|R \cup \{v\}) - M(d|R)$.
- Pick the node with the highest weight-to-cost ratio: $\mathscr{M}(v) := \frac{w(v)}{\text{cost}(v)}$
- Pick the node with the highest neighbors-to-cost ratio: $\mathscr{M}(v) := \frac{|y \in \mathscr{V}_L | \{v,y\} \in \mathscr{E}_L|}{\text{cost}(v)}$
- Pick a node randomly. This strategy might not perform very well, but is very easy to execute after a disaster.

If $M(d)$ can be computed in polynomial time, the node-metrics can also be computed in polynomial time. As such, the simple node-selection strategies are all of polynomial complexity.

input: undirected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, disaster set $\mathcal{D}$, recovery strategy function $R : \mathcal{V} \rightarrow \mathcal{V}$
output: $P(M = m)\ \forall m \in \mathbb{R}$
$O \leftarrow \emptyset$
for all $d \in \mathcal{D}$ **do**
    Determine $S(d) \subseteq \mathcal{V}$
    if $S(d) \in O$ **then**
        $P(S(d)) \leftarrow P(S(d)) + P(d)$
    else
        $P(S(d)) \leftarrow P(d)$
        $O \leftarrow O \cup \{S(d)\}$
    end if
end for
for all $o \in O$ **do**
    $\mathcal{G}_1 \leftarrow \mathcal{G} - o$ ▷ $o \subseteq \mathcal{V}$
    $\mathcal{V}_L \leftarrow \{v \in \mathcal{V} | \exists x \in o\ h(x, v) \leq 1\}$
    Compute $R(o)$
    Order $[v_1, v_2, \ldots] = R(\mathcal{G}_1)$ such that
$\text{cost}(v_1) \leq \text{cost}(v_2) \leq \text{cost}(v_3) \leq \cdots$
    $t_1 \leftarrow 0$
    for $i \leftarrow 1, |R(\mathcal{G}_1)|$ **do**
        $\mathcal{G}_{i+1} \leftarrow \mathcal{G}_i + v_i$ ▷ Where $\{\mathcal{V}_i, \mathcal{E}_i\} + v_i = \{\mathcal{V}_i \cup \{v_i\}, \mathcal{E}_i \cup \{(x, y) \in \mathcal{E} | x, y \in \mathcal{V}_i \cup \{v_i\}\}$
        $t_{i+1} \leftarrow \text{cost}(v_i)$
    end for
    $s \leftarrow [(\mathcal{G}_1, t_1), (\mathcal{G}_2, t_2), \ldots]$
    Compute $M(s)$
    $M(o) \leftarrow M(s)$
end for
$\forall m \in \mathbb{R}\ P(M = m) = \sum_{o \in O | M(o) = m} P(o)$

Figure 6.2: Recovery strategy evaluation algorithm.

6.4. ALGORITHM

Let M be the random value of the evaluation metric after one of the disasters in $\mathcal{D}$ randomly occurs. Given a general recovery strategy, we want to compute the distribution over all possible values of M. Then, by comparing these distributions and the comparative effort to implement each strategy, a general recovery strategy can be chosen by the network operator and other involved parties. When a disaster actually occurs, this strategy can then be implemented immediately, thus wasting no time on deciding on how to best recover the network.

For the purpose of our evaluation algorithm, we consider each possible recovery strategy as a function $R : \mathcal{V} \rightarrow \mathcal{V}$ from the damaged nodes $S(d)$ to a choice of nodes to replace with emergency nodes. Our algorithm is given in Fig. 6.2. We start by computing the set of affected nodes (the outcome) $S(d)$ for each disaster. As the state vector

$[S_k(d)]_{k=1}^{K+1}$ will be the same for each disaster affecting the same nodes:

$$S(d_1) = S(d_2) \Rightarrow [S_k(d_1)]_{k=1}^{K+1} = [S_k(d_2)]_{k=1}^{K+1} \; \forall d_1, d_2 \in \mathscr{D} \tag{6.13}$$

we can compute these states, and M, for each possible outcome instead of for each possible disaster to reduce the computation time.

Next, we go over each possible set of affected nodes and compute the corresponding local network, choose the nodes to recover, create the final state vector $[S_k(d)]_{k=1}^{K+1}$ and compute the value of M.

Using these properties, we can easily compute $P(M = m)$ for each $m \in \mathbb{R}$ by taking the sum of the probabilities of all disasters/outcomes resulting in this value of M. Computing all possible outcomes requires us to iterate over each disaster and each node, which takes $\mathscr{O}(|\mathscr{D}||\mathscr{V}|)$ time (assuming we can determine if a node is in the disaster region in constant time).

Creating the local network takes $\mathscr{O}(|\mathscr{V}| + |\mathscr{E}|)$ time. However, computing the weights of the local nodes takes more time, as we need to find the closest nodes in $\mathscr{V}_L$ of each node in $\mathscr{V}$. This can be accomplished by doing $|\mathscr{V}_L|$ breadth-first searches, and thus takes $\mathscr{O}(|\mathscr{V}_L||\mathscr{V}| + |\mathscr{V}_L||\mathscr{E}|)$ time.

The time it takes to compute the choice of nodes to recover depends on the strategy that is used. For example, the weight-to-cost ratio strategy takes $\mathscr{O}(|K||\mathscr{V}_L| + |K||\mathscr{E}_L|)$ time to compute $R(\mathscr{G}_1)$.

Finally, assuming the weight function can be integrated in constant-time, and the metric used is the WATTR, computing M takes $\mathscr{O}(|K||\mathscr{V}_L| + |K||\mathscr{E}_L|)$ time.

Thus, the time complexity of the algorithm is

$$\mathscr{O}(|\mathscr{D}||\mathscr{V}|^2 + |\mathscr{D}||\mathscr{V}||\mathscr{E}| + |\mathscr{D}|F(|\mathscr{V}|, |\mathscr{E}|, |K|)) \tag{6.14}$$

where $F(|\mathscr{V}|, |\mathscr{E}|, |K|)$ is the time-complexity of the strategy.

6.5. EXPERIMENTS

We apply the framework to two U.S. topologies from the Topology Zoo [156]: ITC Deltacom and Kentucky Datalink. We ignore all nodes without any geographical coordinates.

ITC Deltacom consists of 101 nodes connected by 151 links, while Kentucky Datalink consists of 726 nodes connected by 822 links. Both networks are concentrated in the eastern half of the United States.

For each node v of these networks, we set $p(v)$ to the population of the county containing this node, according to the 2010 US Census [181].

The replacement costs $\text{cost}(v)$ of each node are set randomly to a value between 6 hours and 120 hours (5 days). We use a weight function that decreases linearly to 0 at $t =$ 120 hours, and is constant from then on. After 5 days the emergency recovery operations should be over, and repair operations should be in full swing.

As a use case, we consider a scenario where the network operator knows a hurricane will make landfall in a few days, but not the exact path it will take. Thus, his goal will be to decide on both a strategy and the number of emergency nodes to prepare. We generate a disaster set based on the *5 AM EDT THU AUG 25 2005* hurricane Katrina track prediction of the National Hurricane Center (NHC) [182].

	K=1	K=2	K=3	K=6	K=10
Optimal	0.067	0.098	0.127	-	-
Greedy	0.067	0.095	0.122	0.180	0.215
Weight/Cost	0.022	0.043	0.063	0.134	0.191
Neighbors/Cost	0.035	0.064	0.083	0.128	0.189
Random	0.006	0.013	0.020	0.041	0.082

(a) ITC Deltacom

	K=1	K=2	K=3	K=6	K=10
Optimal	0.076	0.101	0.123	-	-
Greedy	0.076	0.099	0.118	0.149	0.161
Weight/Cost	0.064	0.088	0.105	0.138	0.153
Neighbors/Cost	0.063	0.082	0.096	0.130	0.151
Random	0.030	0.038	0.046	0.074	0.101

(b) Kentucky Datalink

Table 6.1: Expected value of M for different strategies, and different number K of temporary nodes, after hurricane Katrina, based on the *5 AM EDT THU AUG 25 2005* hurricane Katrina track prediction.

To predict potential storm surge flooding, and to assess the probability of wind surface probabilities, the NHC performs Monte Carlo simulations based on the predicted hurricane track and historical errors in their predictions. We propose using these Monte Carlo simulations as representative disaster set. As we do not have access to these simulations, and to demonstrate our approach, we use a simpler hurricane model, based on the NHC Track Forecast Cone. The "Tropical Cyclone Track Forecast Cone" shows the probable path of the center of a tropical cyclone. The cone is formed by simply placing a circle around each predicted track position and connecting them. The size of each circle is set so that two-thirds of historical official forecast errors over a 5-year sample fall within the circle.

We assume the actual track positions (in 2D projected coordinates) are distributed around the predicted positions according to a bivariate Normal distribution. This distribution is composed of normal distributions for the horizontal and vertical positions, each with a standard deviation of $\sqrt{(\frac{r^2}{\ln(10000/1225)})}$, where r is the radius of the corresponding circle, to ensure 65% of samples lie inside the cone.

We can randomly sample hurricane tracks for our own Monte Carlo approach by sampling the track positions and then connecting them with a straight line segment. This only leaves us with the problem of computing the disaster region based on a hurricane track. The strike circle of a hurricane, based on the typical extent of hurricane force winds, is a circle with diameter 231.5 km, centered 23.15 km to the right of the hurricane center (based on its motion) [169]. In our approach we take this circle as the disaster region. Because the hurricane moves through the network area, the complete disaster region of each sampled track takes the form of a union of hippodromes.

Thus the complete approach to generating $\mathscr{D}$ is as follows:

1. Sample N sets of track positions.

2. For each track: compute the resulting disaster region.

3. Set all occurrence probabilities to $\frac{1}{N}$.

The potential hurricane realizations affect between 5 and 38 nodes of the ITC Deltacom network, and between 0 and 89 nodes of the Kentucky Datalink network, depending on their track through the network. On average, around 16 ITC Deltacom nodes and around 18 Kentucky Datalink nodes fail.

Table 6.1 shows the expected values of M utilizing each strategy for different values of K. Due to its high computation cost, we did not compute the expected values of the optimal strategy for $K > 3$. The randomized node selection was evaluated by taking the average of 20 random recovery choices for each possible disaster outcome.

Selecting nodes at random performs very badly compared to the other strategies, especially on the ITC Deltacom topology. This shows how much of a difference it can make to recover nodes according to a suitable strategy.

In this use case, and for these topologies, the greedy strategy performs very close to optimal (at least for $K \leq 3$). As this strategy has polynomial complexity, it seems like a suitable choice.

6.6. RELATED WORK

In this chapter, we aim to restore connectivity in an area by rapidly replacing a selection of network nodes. This approach is essentially equivalent to repairing these nodes (albeit much quicker, but with reduced capacity). Besides replacing or repairing existing network components, connectivity in an area can also be restored by setting up an ad-hoc emergency network. For an overview of this approach, we refer the reader to a survey by Miranda et al. [183].

Want et al. were the first to consider the repair of a network after a disaster, and proposed the progressive network recovery model [180]. Under progressive network recovery, the time spend repairing the network is divided into several stages. In each stage, all available repair resources (e.g., repair crews) are divided among broken network components. This process continues until the network is fully repaired, or we have run out of stages. Want et al. studied the problem of optimizing the assignment of repair resources in each stage, as to maximize the weighted sum of maximum flow over the entire recovery process. They proved this problem is NP-hard, and provided heuristics as well as an MIP formulation for solving it.

A number of variants of the progressive network recovery model have been proposed. Genda and Kamamura modified the objective function to provide a balance between total network flow and the division of bandwidth between different logical flows [184]. Al Sabeh et al. applied the progressive network recovery model to opaque, transparent, and elastic optical networks [185]. Pourvali et al. proposed several progressive network recovery schemes for network virtualization services [186]. Mazumder et al. studied the progressive recovery of interdependent, multi-layer networks [187]. Ferdousi et al. proposed coordinating datacenter and network repair [188]. Ciavarella et al. considered progressive network recovery under uncertainty of the state of the network [189]. In their model, the state of the network is gradually uncovered as more and more network components are repaired.

Ishigaki et al. applied machine learning to progressive network recovery [190]. They formulated the progressive network recovery of an interdependent network of infrastructure and virtual network function nodes, and proposed a deep reinforcement learning algorithm for solving it. Note that the algorithm is trained after the network has been damaged, using the state of the network as its initial state.

The progressive network recovery model does not take into account the travel time of repair crews (instead, it considers repair crews as abstract resources it needs to assign). To help operators create repair schedules that do take travel time into account, Ma et al. proposed the multiple traveling repairmen problem and gave two heuristics to solve it [191]. The objective of this problem is to find an optimal schedule for teams of repairmen, taking into account network virtualization, as well as the travel time between locations.

Bartolini et al. did not study the creation of an optimal schedule, but instead introduced the problem of finding a minimum-cost selection of components to repair that satisfied all required demands [192]. This more closely resembles the optimization problem discussed in this chapter, with the main difference being that we aim to maximize performance, while Bartolini et al. aim to minimize cost.

Zad Tootaghaj et al. extended the work of Bartolini et al. by introducing uncertainty about the state of the network [193]. Since in their problem formulation, the state of some network components is unknown, they propose an iterative approach: components are selected for repair, their repair provides the operator with more knowledge, which in turns helps in selecting more components to repair. While this approach is progressive, since components are essentially repaired in stages; the objective is still to minimize total repair cost.

Xu et al. proposed an SDN-based system that allows different carriers to collaborate by connecting their networks together after a disaster [194]. Setting up this interconnection-based emergency network still requires carriers to repair network components. Thus, Xu et al. formulated optimization problems for selecting which components to repair [195]. Xu et al. also proposed another scheme, in which carriers can sell lightpaths to one another [196]. By collaborating, carriers can reduce the number of repair operations (and repair cost) required to restore network connectivity.

The commonality among all approaches described thus far is that they are reactive: Once a disaster has damaged the network, repair schedules are setup and the network is slowly restored. In contrast, we propose a more pro-active approach, where the network operator already evaluates and decides on a repair strategy before a disaster has even struck the network.

While our repair model is simpler than the progressive network recovery model (due to our focus on the quick replacement of nodes), our approach for systematically evaluating repair strategies and algorithms based on a large set of representative disasters can be applied to any repair model and strategy.

6.7. CONCLUSION

In the period shortly after a natural disaster, the need for communication networks only increases. Unfortunately, repairing the network and restoring network functionality to the affected area can be a long process. Thus, in this chapter, we have studied strategies

for strategically replacing a number of nodes with emergency equipment. This emergency equipment can temporarily take over some of the functionality of these nodes, and restore connectivity within the local area.

We have proposed an extension to our single-disaster framework for evaluating any potential node replacement strategy. One of the main ideas behind our approach is to reduce computation time by only considering the local area that has been directly affected by the disaster, as this is the area we want to reconnect to the network. Since finding the optimal repair strategy is an NP-hard problem, we have proposed that network operators pro-actively select a simple repair strategy, which they can then immediately implement after a disaster has struck the network. We have demonstrated this concept by applying our evaluation framework to two U.S. topologies and a selection of strategies. By evaluating, selecting, and preparing a repair strategy before a disaster has even struck the network, operators can increase their ability to adapt to and recover the network from a disaster, which improves its disaster resilience.

7

GOING THE EXTRA MILE WITH DISASTER-AWARE NETWORK AUGMENTATION

Network outages have significant economic and societal costs. While network operators have become adept at managing smaller failures, this is not the case for larger, regional failures such as natural disasters. Although it is not possible, and certainly not economic, to prevent all potential disaster damage and impact, network operators can reduce their impact by adding cost-efficient, geographically redundant, cable connections to the network.

In the previous chapter, we considered strategies for restoring network performance after the network has been struck by a disaster. In contrast, adding new, geographically redundant connections to the network improves the ability of the network to absorb the impact of a disaster, by reducing the initial loss in network performance.

In this chapter, we provide algorithms for finding cost-efficient, disaster-aware cable routes based on empirical hazard data. In contrast to previous work, our approach finds disaster-aware routes by considering the impact of a large set of input disasters on the network as a whole, as well as on the individual cable. For this, we propose the Disaster-Aware Network Augmentation Problem of finding a new cable connection that minimizes a function of disaster impact and cable cost. We prove that this problem is NP-hard and give an exact algorithm, as well as a heuristic, for solving it. Our algorithms are applicable to both planar and geographical coordinates. Using actual seismic hazard data, we demonstrate that by applying our algorithms, network operators can effectively raise the resilience of their network and future cable connections.

Parts of this chapter have been published in IEEE INFOCOM 2021 [197].

7.1. INTRODUCTION

In 2006, an earthquake of the coast of Taiwan damaged 8 submarine cable systems, severely disrupting communications in the region [113]. In 2008 and 2009, new cable systems were installed that deliberately avoided this earthquake region. Thus, when a similar event damaged the same 8 cable systems again in 2009, network operators were able to restore service much quicker [198]. Numerical simulations suggest disaster-aware submarine cable deployments could potentially save society billions of dollars [73].

By installing a new cable connection, a network operator can introduce *geographic redundancy*. In case of a disaster, connections can be routed through the new cable instead of through the disaster region. As a simple example, consider the new cable connection depicted in Fig. 7.1. By avoiding D1, the new link ensures that nodes 2 and 3 remain connected. Note that avoiding D1 forces the new cable to either go through disaster D2 or make a large detour. Designing disaster-resilient topologies requires operators to make these kinds of compromises for hundreds of disaster regions, taking into account cable laying costs, disaster probabilities, the impact of disasters on the network as a whole, as well as the impact of a disaster on the new cable itself. Taking into account all possible combinations of failures and disaster regions manually would be too time-consuming. Thus, to create truly disaster-resilient network topologies, we need an automated system that can suggest potential cable routes based on actual hazard data.

Although there is a large body of work on finding disaster-resilient cable connections, none of the previous work considers, simultaneously, the impact of a large class of disasters on the cable route itself, as well as on network connectivity as a whole. To fill this gap, we propose a set of algorithms for finding cost-efficient, disaster-aware cable connections based on a large set of representative disasters. The main idea behind our algorithms is to separate the decision of which disaster regions to avoid from the design of the route itself. This allows us to develop exact and heuristic algorithms that search through the problem space and are able to incorporate any pathfinding algorithm for computing the actual routes.

Of course, the final decision on the design of a network or cable must be made by the stakeholders, and not by an automated system. By varying an input parameter, our algorithms can quickly generate multiple routes that are Pareto-optimal in cable cost and expected disaster impact. In addition, because our algorithms assign a specific cost to each disaster, and specifically select a set of regions to avoid, they can provide detailed information on *why* a proposed cable connection takes a certain route. Armed with this data, network operators, governments, and other stakeholders can make an informed, disaster-aware decision on any new cable connection.

Our main contributions are as follows:

- We define the Disaster-Aware Network Augmentation Problem of finding a new cable connection that minimizes a cost-function of expected disaster impact and cable cost (Section 7.2). By varying a parameter, α, in the objective function we make it possible for network operators to find various different Pareto-optimal connections for cable cost and expected impact.

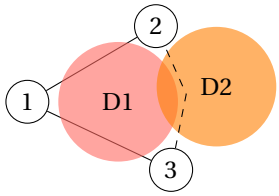


Figure 7.1: Example disaster-aware network augmentation. Node 3 is connected to node 2 through node 1. However, if disaster D1 were to occur, this connection would fail. To increase the resilience of this network, we can add a cable connecting nodes 3 and 2 (dotted line), avoiding disaster D1.

- Since the Disaster-Aware Network Augmentation Problem is NP-Hard (Section 7.3), we propose both an exact branch & bound algorithm (Section 7.4), as well as a heuristic (Section 7.5). Our algorithms are applicable to both the plane, as well as geographical coordinates.

- We demonstrate our approach by augmenting a real network topology based on actual seismic hazard data (Section 7.6). Given a representative disaster set of 100,000 disasters, our algorithms are able to compute cost-efficient network augmentations within 3 minutes.

7.2. PROBLEM STATEMENT

We model the network as a directed multigraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \psi)$, with nodes $v \in \mathcal{V}$ connected by links $e \in \mathcal{E}$, where $\psi : \mathcal{E} \to \mathcal{V} \times \mathcal{V}$ and $e \in \mathcal{E}$ connects v_1 to v_2 iff $\psi(e) = (v_1, v_2)$. We model the physical structure and location of network nodes as points in $\mathbb{R}^2$ and denote these points as $p(v)$ for all $v \in \mathcal{V}$. Any link $e \in \mathcal{E}$ between v_1 and v_2 is modelled as a finite sequence of line segments or geodesics connecting v_1 to v_2, $seg(e) = (s_1, s_2), (s_2, s_3), \ldots, (s_{l-1}, s_l)$ where $s_1, \ldots, s_l \in \mathbb{R}^2$, $s_1 = p(v_1)$, and $s_l = p(v_2)$.

To determine how to optimally augment this network, we assume we are given a finite set of representative disasters[1] to protect against, $\mathcal{D}$: We model each potential disaster $d \in \mathcal{D}$ as a disaster region $A(d)$ in the plane or on the globe, with associated probability $P(d)$, and assume exactly one of these disasters occurs (i.e. $\sum_{d \in \mathcal{D}} P(d) = 1$) and destroys all network components intersecting its disaster region. We denote this random disaster by D.

Representative sets of disasters are similar to the stochastic event sets used in catastrophe modeling [111] and should be easily obtainable by network operators. In contrast to many approaches that make use of stochastic event sets, we do not assign structural failure probabilities to network components, but instead assume a more pessimistic outcome where every component inside an affected region fails. For network augmentation, such a worst-case perspective should result in more resilient and less overfitted cable connections among the many uncertainties involved in disaster modeling.

If a node lies in a disaster region, all its incident links also intersect this region. Thus, we do not need to explicitly consider node failures, as the failure of all incident links

[1] Our disaster model is based on the model introduced in Chapter 3.

would disconnect nodes from the network as well. We define the *failure state*, $S(d) \subseteq \mathscr{E}$, of a disaster d to be the set of links intersecting the disaster region $A(d)$, where we say a link $e \in \mathscr{E}$ intersects $A(d)$ if and only if one or more of its line segments, $seg(e)$, intersects $A(d)$.

Before augmenting the network, we first need an impact metric over these failure states to optimize towards. For this purpose, we construct the set $\mathscr{E}^+_{\mathscr{G}}(d) \subseteq \mathscr{V} \times \mathscr{V}$ of all node pairs that are still directly connected by a functioning link:

$$\mathscr{E}^+_{\mathscr{G}}(d) := \psi[\mathscr{E} \setminus S(d)] \tag{7.1}$$

We allow any function $M : \mathscr{P}(\mathscr{V} \times \mathscr{V}) \to \mathbb{R}$ over these sets of node pairs as an impact metric, as long as

$$\forall B \subseteq C \subseteq \mathscr{V} \times \mathscr{V}, M(B) \geq M(C) \tag{7.2}$$

7.2.1. CABLE COSTS

When suggesting the addition of new links to the network, it is imperative to take the costs of installing these links into account. A simple measure of this cost is cable length. However, costs can vary greatly depending on the specific path of the cable; e.g., if it crosses less accessible areas. To take these factors into account, we divide a rectangular area encompassing the network into a grid of $w \times h$ cells and assume we are given the costs of laying a cable from the center of each cell to the centers of all 8 of its neighbors.

We formulate the route of a new link e from node v_1 to v_2 as a sequence of grid cells, $r(e) = c_{x_1,y_1}, \ldots c_{x_l,y_l}$, where any successive cell c_{x_i,y_i} is a neighbor of the previous cell $c_{x_{i-1},y_{i-1}}$, $v_1 \in c_{x_1,y_1}$, and $v_2 \in c_{x_l,y_l}$. The cost of this route is

$$\mathscr{C}(r(e)) = \sum_{i=1}^{l-1} \mathscr{C}(c_{x_i,y_i}, c_{x_{i+1},y_{i+1}}), \tag{7.3}$$

where $\mathscr{C}(c_{x_i,y_i}, c_{x_{i+1},y_{i+1}})$ is the cost of laying a cable between cells c_{x_i,y_i} and $c_{x_{i+1},y_{i+1}}$. The exact path of the fiber, $seg(e)$, can now be constructed by connecting the centers of the grid cells in $r(e)$:

$$seg(e) = (p(v_1), ctr(c_{x_2,y_2})), \ldots, (ctr(c_{x_{l-1},y_{l-1}}), p(v_2)), \tag{7.4}$$

where $ctr(c_{x_i,y_i})$ is the center of cell c_{x_i,y_i}.

Note that the existing links $\mathscr{E}$ of $\mathscr{G}$ do not need to adhere to this grid system, and we solely use the grid and $\mathscr{C}$ as a means of computing the path and cable cost of *new* links. Furthermore, our algorithms are also applicable to any other system for computing cable costs, as long as it allows us to compute a shortest path avoiding a given set of disaster regions.

7.2.2. DISASTER-AWARE NETWORK AUGMENTATION PROBLEM

Any augmentation can now be defined by three properties: (1) the source node, v_1; (2) the destination node, v_2; and (3) the route of grid cells connecting these two nodes, r. Given such a *link triple*, we define a network augmentation as follows:

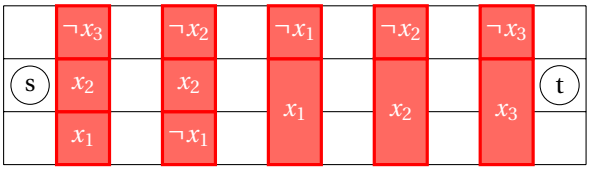


Figure 7.2: Example of a reduction from a 3-SAT instance $((x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2 \vee \neg x_2))$ to a Disaster-Aware Network Augmentation Problem instance.

Definition 3 (Network Augmentation)**.** *Given a link triple* (v_1, v_2, r)*, where* $v_1, v_2 \in \mathcal{V}$ *and* r *is a valid route of cells connecting these nodes,*

$$(\mathcal{V}, \mathcal{E}, \psi) + (v_1, v_2, r) = (\mathcal{V}, \mathcal{E} \cup \{e\}, \psi') \tag{7.5}$$

where

$$\psi'(e') = \begin{cases} (v_1, v_2) & \text{if } e' = e \\ \psi(e') & \text{otherwise} \end{cases} \tag{7.6}$$

and $seg(e)$ *is given by Eq. 7.4.*

Any impact metric M for $\mathcal{G}$ is also applicable to $\mathcal{G} + (v_1, v_2, r)$, giving us a straightforward way of computing the benefit of augmenting the network with any link triple.

Definition 4 (Disaster-Aware Network Augmentation Problem)**.** *Given a directed multigraph* $\mathcal{G}$*, node locations* p*, link segments* seg*, cable costs* $\mathscr{C}$*, metric* M*, and* $\alpha > 0$*, find a link triple* (v_1, v_2, r) *that minimizes*

$$\begin{aligned} &cost(v_1, v_2, r) := \\ &\alpha E[M(\mathscr{E}^+_{\mathcal{G}+(v_1,v_2,r)}(D))] + \mathscr{C}(r) \end{aligned} \tag{7.7}$$

Remark 6.1. *By varying* α*, we can find different Pareto-optimal link triples for expected impact and cable cost. If possible, one should choose* α *such that* $\alpha E[M(\mathscr{E}^+_{\mathcal{G}}(D))]$ *roughly represents the expected future cost of the class of disasters taken into consideration.*

This problem can be divided into two sub-problems:

1. Given two nodes $v_1, v_2 \in \mathcal{V}$, find a route, r, that minimizes $\text{cost}(v_1, v_2, r)$;
2. Find the optimal source and destination nodes $v_1, v_2 \in \mathcal{V}$.

7.3. NP-HARDNESS

Theorem 7. *The Disaster-Aware Network Augmentation Problem is NP-hard, even if we restrict ourselves to a single node pair.*

Proof. We will provide a polynomial-time reduction from the NP-complete 3-SAT problem [199] to the decision variant of the Disaster-Aware Network Augmentation Problem.

Suppose we are given a Boolean formula f in conjunctive normal form, where each clause contains exactly three literals:

$$f = C_1 \wedge C_2 \wedge \cdots \wedge C_k \tag{7.8}$$

with

$$C_1 = l_{1,1} \vee l_{1,2} \vee l_{1,3}, C_2 = l_{2,1} \vee l_{2,2} \vee l_{2,3}, \ldots \tag{7.9}$$

The 3-SAT problem is to determine if this formula is satisfiable. Let V be the set of all variables in the formula. To reduce f to an instance of the Disaster-Aware Network Augmentation Problem, we first create a grid of $(1+2k+2|V|) \times 3$ cells and assign the same cost of $\frac{1}{3(1+2k+2|V|)}$ to each of the possible connections from a cell to its neighbors. We then create a graph $\mathscr{G} = \{\mathscr{V}, \mathscr{E}, \psi\}$ of two nodes ($\mathscr{V} = \{s,t\}$) and no links ($\mathscr{E} = \emptyset$) and place s in the middle cell of the left-most column, and t in the middle row of the right-most column.

We form a disaster set $\mathscr{D}$ by combining all possible literals, i.e., $\mathscr{D} = \bigcup_{x \in V} \{x, \neg x\}$. Fig. 7.2 demonstrates how we construct the disaster regions:

1. We create a column of 3 disaster sub-regions for each clause C_i of f that fills up exactly one column of cells: the disaster region of $l_{i,1}$ fills up the first cell of this column, the region of $l_{i,2}$ the second cell, and the region of $l_{i,3}$ the third cell.

2. We also create a column of 2 disaster sub-regions for each variable $x \in V$, where x fills up the first two cells of the column, and $\neg x$ the third cell.

3. We then place all these columns directly to the right of the column containing s, and put a spacing of 1 cell between each successive column.

We set $P(d) = \frac{1}{\mathscr{D}}$ for all $d \in \mathscr{D}$ and choose $\alpha = |\mathscr{D}|$. Finally, we assign an impact of 0 if s is connected to t, and an impact of 1 otherwise:

$$\begin{aligned} M(\{(s,t),(t,s)\}) = 0, M(\{(s,t)\}) = 0, \\ M(\emptyset) = 1, M(\{(t,s)\}) = 1 \end{aligned} \tag{7.10}$$

Note that this means that a solution that connects t to s instead of s to t will always have a cost of at least $|\mathscr{D}|$.

Now, suppose there is a route r from s to t with cost

$$\alpha E[M(\mathscr{E}^+_{\mathscr{G}+(s,t,r)}(D))] + \mathscr{C}(r) < \frac{1}{2}\mathscr{D} + 1 \tag{7.11}$$

This would mean that r intersects at most $\frac{1}{2}\mathscr{D} = |V|$ disasters, as otherwise $\alpha E[M(E^+_{G+(s,t,r)}(D))] \geq \frac{1}{2}\mathscr{D} + 1$. To reach t from s, the route must intersect x or $\neg x$ for each variable $x \in V$. Because the number of intersected disasters is at most $|V|$, this means that for all $x \in V$, the route can not intersect both x and $\neg x$. In addition, the route must also intersect at least one disaster region of the literals of each clause. Thus, the selection of literals intersected by r form a satisfying assignment for f.

Vice versa, suppose f is satisfiable. In other words, there is an assignment of TRUE and FALSE to each variable $v \in V$ such that each of the clauses of f (and thus f itself) is satisfied. We will construct a route r from s to v such that

$$\alpha E[M(\mathscr{E}^{+}_{\mathscr{G}+(s,t,r)}(D))] + \mathscr{C}(r) < \frac{1}{2}\mathscr{D} + 1 \tag{7.12}$$

First, our route will need to cross all the columns of clauses. Because each clause is satisfied by the assignment, at least one literal of each clause must evaluate to TRUE. Thus, we construct the route in such a way that we only intersect this literal of the clause. Note that due to the spacing between successive clauses, this is always possible. Next, the route will need to cross the columns of variables as well. Here, as with the clauses, we intersect the literal that evaluates to TRUE. This way, we can construct a route r from s to t that only intersects the $|V|$ literals that evaluate to TRUE. Furthermore, because we pass through each cell at most once,

$$\begin{aligned} &\alpha E[M(\mathscr{E}^{+}_{\mathscr{G}+(s,t,r)}(D))] + \mathscr{C}(r) < \\ &\frac{1}{2}\mathscr{D} + \frac{3(1+2k+2|V|)}{3(1+2k+2|V|)} = \\ &\frac{1}{2}\mathscr{D} + 1 \end{aligned} \tag{7.13}$$

We thus have a polynomial-time reduction from the 3-SAT problem to the decision variant of the Disaster-Aware Network Augmentation Problem, and can conclude that the Disaster-Aware Network Augmentation Problem is NP-hard. □

Remark 7.1. *This proof also applies to sub-problem 1 by itself. Thus, determining the optimal route between a given source and destination node is already NP-hard.*

7.4. BRANCH AND BOUND

In this section we describe an exact algorithm for sub-problem 1 based on the branch and bound paradigm. Suppose we are given two nodes $v_1, v_2 \in \mathcal{V}$, our goal is to find a route r of cells from v_1 to v_2 that minimizes $\text{cost}(v_1, v_2, r)$. This might seem similar to the shortest path problem. However, the difficulty lies in that, unlike for the shortest path problem, where the sub-path of a shortest path is itself a shortest path, in our case a sub-route of a minimum-cost route is not necessarily a minimum-cost route itself.

The key insight behind our approach is that if we decide on a specific set of disasters to avoid, $R \subseteq \mathscr{D}$, the problem of finding a route with minimum cable cost between v_1 and v_2 that does not intersect any disaster in R *is* a shortest path problem. In the rest of this chapter, we call these subsets of representative disasters *restrictions*, and the minimum-cable-cost route avoiding a restriction a *restricted shortest route*. As we show in this section, we can quickly compute the cost of any given route between v_1 and v_2 as a sum of pre-computed disaster *penalties* and cable costs. Thus, to determine the cost associated to a given restriction R, we simply find a restricted shortest path for R, and compute the cost of this path. This allows our algorithm to search for the optimal restriction R instead of the optimal route, greatly simplifying the problem.

1: **input:** $\mathcal{G}, v_1, v_2, M, \mathscr{C}, \mathscr{D}, \alpha$
2: **output:** optimal route from v_1 to v_2, r
3: compute $M(d)^+ \; \forall d \in \mathscr{D}$ ▷ Equation 7.15
4: $W \leftarrow \alpha \sum_{d \in \mathscr{D}} P(d) I(d,r) M(d)^+ + \mathscr{C}(r)$ ▷ Equation 7.17
5: $W(\emptyset) \leftarrow \infty$
6: $\mathscr{D} \leftarrow \{d \in \mathscr{D} | M(d)^+ > 0, p(v_1) \notin A(d), p(v_2) \notin A(d)\}$
7: $r \leftarrow$ SEARCH$(\mathscr{D}, \emptyset, \emptyset, \emptyset)$

8: **function** SEARCH$(\mathscr{D}, R, \mathscr{D}^-, r_*)$
9: cutoff $\leftarrow W(r_*) - \alpha \sum_{d \in \mathscr{D}^-} P(d) M(d)^+$
10: try to find a restricted shortest route $sp(R)$ from v_1 to v_2 with a cutoff cost of cutoff
11: **if** $sp(R)$ not found **then**
12: $r \leftarrow r_*$
13: **else**
14: **if** $W(sp(R)) < W(r_*)$ **then**
15: $r \leftarrow sp(R)$
16: **else**
17: $r \leftarrow r_*$
18: **end if**
19: **for all** $d \in \mathscr{D} \setminus (\mathscr{D}^- \cup R)$ intersected by $sp(R)$ **do**
20: $r \leftarrow$ SEARCH$(\mathscr{D}, R \cup \{d\}, \mathscr{D}^-, r)$
21: $\mathscr{D}^- \leftarrow \mathscr{D}^- \cup \{d\}$
22: **end for**
23: **end if**
24: **return** r
25: **end function**

Figure 7.3: Pseudocode for the exact depth-first branch and bound algorithm for finding the minimum-cost route from node v_1 to node v_2.

We first introduce an indicator value $I(d,r)$:

$$I(d,r) = \begin{cases} 1 & \text{if } seg(r) \text{ intersects } A(d) \\ 0 & \text{otherwise} \end{cases} \tag{7.14}$$

Regardless of our choice of route r, the impact of any disaster $d \in \mathscr{D}$ is $M(\mathscr{E}^+_{\mathcal{G}}(d))$ if r intersects it, and $M(\mathscr{E}^+_{\mathcal{G}}(d) \cup \{(v_1, v_2)\})$ otherwise. If we take the difference of these values, we get a measure of the benefit of adding a connection from v_1 to v_2 in case of a disaster d:

$$M(d)^+ := M(\mathscr{E}^+_{\mathcal{G}}(d)) - M(\mathscr{E}^+_{\mathcal{G}}(d) \cup \{(v_1, v_2)\}) \tag{7.15}$$

Now, for any route r,

$$\begin{aligned} E[M(\mathscr{E}^+_{\mathscr{G}+(v_1,v_2),r)}(D))] = \\ \sum_{d\in\mathscr{D}} P(d)(M(\mathscr{E}^+_{\mathscr{G}}(d)\cup\{(v_1,v_2)\}) + I(d,r)M(d)^+) = \\ E[M(\mathscr{E}^+_{\mathscr{G}}(D)\cup\{(v_1,v_2)\})] + \sum_{d\in\mathscr{D}} P(d)I(d,r)M(d)^+ \end{aligned} \tag{7.16}$$

If we subtract any constant from our objective function the resulting optimization problem is equivalent to our old one. Thus, we subtract $\alpha E[M(\mathscr{E}^+_G(D)\cup\{(v_1,v_2)\})]$ to obtain the new objective function

$$W(r) := \alpha \sum_{d\in\mathscr{D}} P(d)I(d,r)M(d)^+ + \mathscr{C}(r) \tag{7.17}$$

As $P(d)M(d)^+$ does not depend on the route of the link itself and can be pre-calculated, $W(r)$ can be seen as the sum of the cable cost, $\mathscr{C}(r)$, and a pre-computed penalty, $\alpha P(d)M(d)^+$, for every intersected disaster region $A(d)$.

We denote a restricted shortest route by $sp(R)$, where $R\subseteq D$ is the set of disasters this route should avoid. Our exact algorithm is a depth-first search for the optimal restriction R. The algorithm starts at $R=\emptyset$, and tries to find the optimal restriction and route from there. The algorithm is provided in pseudocode in Fig. 7.3. For readability, the algorithm is formulated as a recursive function call. However, our implementation uses an iterative approach.

7.4.1. BRANCHING

The number of possible restrictions ($2^{|\mathscr{D}|}$) grows exponentially in $|\mathscr{D}|$. Thus, to keep computation times manageable, reducing the number of considered disasters is essential. Fortunately, it is likely that for many of the representative disasters $d\in\mathscr{D}$, connecting v_1 to v_2 will not bring any benefit and $M(d)^+=0$. In addition, disasters that intersect $p(v_1)$ or $p(v_2)$ might have a positive benefit $M(d)^+)>0$, but can not be avoided. These two sets of disasters do not need to be considered by our algorithm and are excluded.

After computing $sp(R)$, we can limit the number of potential branches even more. If $sp(R)$ does not intersect a disaster $d\in\mathscr{D}$, adding d to the restriction will not change the restricted shortest route. Thus, for any restriction R, we only consider extending R with disasters intersected by $sp(R)$, where we say a route $sp(R)$ intersects a disaster $d\in\mathscr{D}$ if and only if any of the line segments $seg(sp(R))$ intersect $A(d)$.

We choose to branch on individual disasters: If $sp(R)$ intersects k disasters with positive benefit, $d_1,\ldots,d_k\in\mathscr{D}$ from large to small benefit, we create k branches, $R\cup\{d_1\},\ldots,R\cup\{d_k\}$. If the optimal solution avoids $R\cup\{d_i\}$, our approach will find this solution in branch $R\cup\{d_i\}$. Thus, after having visited branch $R\cup\{d_i\}$, we remove d_i from consideration in branches $R\cup\{d_{i+1}\},\ldots,R\cup\{d_k\}$. This both prevents the algorithm from visiting the same restriction twice and further limits the number of considered restrictions.

7.4.2. BOUNDING

Throughout our algorithm, we keep track of the best route encountered so far, r_*, and its objective value, $W(r_*)$. Once we know that all further restrictions on R would lead

to a worse solution than r_*, we can stop exploring branch R. For every restriction R, $\mathscr{C}(sp(R))$ is a lower bound on the objective value W for any further restrictions. However, by taking into account the disasters we explicitly removed from consideration in our branching approach, we can improve upon this bound.

Let $\mathscr{D}^-$ be the set of all disasters removed from consideration. If the optimal route avoids a disaster $d \in \mathscr{D}^-$, it has already been found previously and r_* is the optimal route. Thus, we can stop exploring a branch when

$$\mathscr{C}(sp(R)) \geq W(r_*) - \alpha \sum_{d \in \mathscr{D}^-} P(d) M(d)^+ \tag{7.18}$$

7.4.3. SHORTEST ROUTE COMPUTATIONS

Our algorithm needs to compute a new route $sp(R)$ for every considered restriction R. To speed up computations, we pre-compute the minimum distance between each cell and v_2 using Dijkstra's Algorithm, and then use A^* to compute restricted shortest routes from v_1 to v_2. Furthermore, we immediately stop computing a route once the cost to reach the current cell and the minimum distance between the current cell and v_2 exceeds the cutoff value given in Eq. 7.18.

While computing restricted shortest routes, we constantly need to check if a line segment between two adjacent cells does not intersect any disaster $d \in R$. To reduce the computation time spent on these checks, we use caches to keep track of which line segments intersect which disasters.

7.4.4. GLOBAL OPTIMIZATION

To find the optimal link triplet (v_1, v_2, r), we apply our branch & bound algorithm to every pair of nodes in the network. In this context, we make some small adjustments to the algorithm to further reduce computation times. First, we propose pre-computing the failure state $S(d)$ and impact $M(\mathscr{E}^+_{\mathscr{G}}(d))$ of all disasters $d \in \mathscr{D}$. This allows us to compute penalties for each failure state instead of for each disaster. As the number of failure states tends to be much smaller than the number of disasters, this significantly speeds up the pre-computation phase of each node pair.

Second, we keep track of the minimum-cost route r_* and the corresponding upper bound across all node pairs and pass this global upper bound to the branch and bound algorithm. This requires us to transform the global upper bound (on cost) to a local upper bound (on W). Let u_{global} be a global upper bound, we can transform u_{global} to a local upper bound by

$$u_{\text{local}} = u_{\text{global}} - \alpha(E[M(\mathscr{E}^+_G(D) \cup \{(v_1, v_2)\})]) \tag{7.19}$$

We apply this bound as an initial upper bound for our search function, as well as a limit on the cells we pre-compute A^* heuristics for. Note that the transformed local upper bound might be negative. In this case no possible route from v_1 to v_2 could improve upon our current global best route, and we can skip node pair (v_1, v_2).

The global upper bound is essential in reducing computation times. However, depending on the order we traverse node pairs in, it might take a long time before we have obtained a low upper bound. Thus, to obtain a reasonable upper bound a priori, we

initially compute the shortest route for each node pair. We then select the route with minimum cost and use it as the initial value for r_*.

7.5. HEURISTIC

Although our branch & bound algorithm is fast enough for many practical use cases, its runtime is still exponential in the number of representative disasters, $\mathscr{D}$. Since the Disaster-Aware Network Augmentation Problem is NP-hard, we propose a heuristic for sub-problem 1 that can find near-optimal solutions for larger disaster sets in a fraction of the runtime of the branch & bound approach.

In the previous section, we have reduced the problem of finding the optimal route between two nodes to the problem of determining which regions to avoid. We use this same concept to create a heuristic and apply simulated annealing to find an approximate solution for the following optimization problem:

$$\min_{R \subseteq \mathscr{D}} W(sp(R)) = \alpha \sum_{d \in \mathscr{D}} P(d) I(d, sp(R)) M(d)^+ + \mathscr{C}(sp(R)), \tag{7.20}$$

where $sp(R)$ is a restricted shortest route from v_1 to v_2 avoiding R. To find the global link triple, we use the same approach as described in Section 7.4.4, and simply replace the branch & bound approach with our heuristic.

7.5.1. SIMULATED ANNEALING

A simulated annealing algorithm searches for an optimal solution by randomly selecting and evaluating neighboring solutions. If the neighbor has a lower cost, the algorithm directly switches to this solution. If not, it does so with a probability depending on the current *temperature* as well as the difference in costs [200]. By starting with a high temperature and slowly decreasing it over time, simulated annealing initially searches throughout the problem space and then gradually "locks in" to a local minimum.

An important consideration when applying simulated annealing to any problem is the selection of neighbors. We take the same approach as described in Section 7.4.1 (but do not exclude $\mathscr{D}^-$). The neighbors of a restriction R are

$$\{R \cup \{d\} | d \in \mathscr{D} \setminus R \wedge sp(R) \text{ intersects } A(d)\} \tag{7.21}$$

and

$$\{R \setminus \{d\} | d \in R\} \tag{7.22}$$

For our experiments, we have made the following implementation choices for our simulated annealing algorithm.

- **Initial solution**: As in the branch & bound approach, we start at the empty restriction $R \leftarrow \emptyset$, and compute $sp(\emptyset)$.
- **Transition probability**: Let δ be the increase in cost of the neighboring solution. We transition to this solution with probability 1 if $\delta < 0$ and with probability $e^{\frac{-\delta}{T}}$ otherwise, where T is the current temperature.

- **Temperature** T: We set the initial temperature to $\frac{-\alpha \sum_{d \in \mathscr{D}} P(d) I(d, sp(\emptyset)) M(d)^+}{\ln 0.25}$.
- **Temperature Reduction**: Every 10 "repetitions" of considering a neighboring solution we reduce the temperature by $T \leftarrow 0.9T$.
- **Freezing Point** and **Stopping Condition**: If none of the 10 repetitions resulted in a move to a different solution, we set the temperature T to 0 (thus switching to hill climbing). If none of the repetitions resulted in a move, and the temperature is already 0, we submit the restricted shortest path with the lowest cost we have encountered up till this point as the solution.

We make the following modification to the standard simulated annealing implementation: If the cable cost of a restricted shortest path $sp(R)$ exceeds $W(r^*)$, where r^* is the best route we have found up till this point, we outright reject R as a potential solution and set $\delta = \infty$. As discussed in Section 7.4.2, this prevents the algorithm from moving to restrictions that are too restrictive to improve upon r^*, while not cutting off the optimal solution from the search space. More importantly, it allows us to save computation time by setting an upper bound on the maximum cable cost and cutting of the pathfinding algorithm if this upper bound is exceeded.

7.6. EXPERIMENTS

We demonstrate our methods on the undirected Italian sub-topology of Interoute, a 25-node, 35-link network traced and made publicly available by the authors of [109]. All experiments were conducted on commodity hardware: an AMD Ryzen 7 3700X 3.6 GHz processor with 32 GB of available RAM.

We augment Interoute with respect to a publicly available earthquake dataset [109]. This dataset was purpose-built by seismologists for analyzing the resiliency of communication networks. It essentially consists of a set of 1,196,037 disk disasters, which together represent all possible earthquakes that can strike Italy.

To demonstrate the applicability of our algorithms to geographic coordinates, we do not transform coordinates to the plane. Instead, we construct a grid of cells of around 0.05 x 0.05 degree covering the longitude-latitude coordinates of all nodes padded by 0.05 degrees on all sides (resulting in a 232 by 194 grid). To compute cable costs and determine which network components are affected by a disaster, we use the great circle distance for a sphere with radius 6,371 km.

We filter out all disasters that do not damage any network components and reweight the remaining probabilities to make them sum to 1, as it is not necessary to protect the networks against these events and they would be filtered out at the penalty computation stage. This leaves us with 454,433 out of 1,196,037 disasters.

Our impact metric, M, is the number of disconnected node pairs, divided by the total amount of node pairs (i.e., 1− the ATTR). Thus, $M = 0$ if all node pairs are connected, and $M = 1$ if all node pairs are disconnected. The expected impact of earthquakes on Interoute is approximately 0.0141. This might seem small, but the disaster set is simply so extensive that it also contains many disasters that barely affect the network. In fact, the total disaster rate is 1.6006 per year. This means that on average, more than 2% of

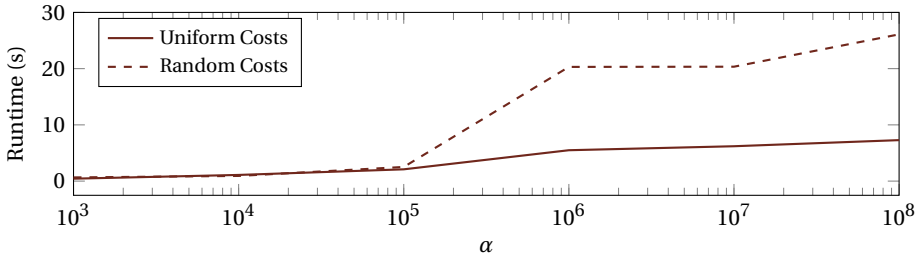
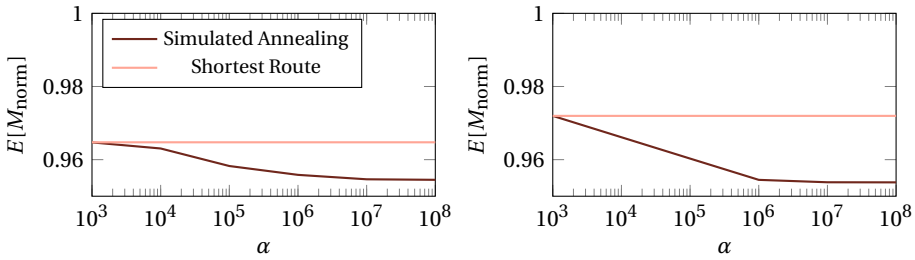


Figure 7.4: Mean computation time of a route between two random nodes (using simulated annealing).



(a) Uniform Costs

(b) Random Costs

Figure 7.5: Mean expected impact after adding a route between two random nodes divided by the expected impact on the initial network topology. The routes were computed based on a sampled set of 50,000 disasters, and evaluated on the full set of 454,433 disasters.

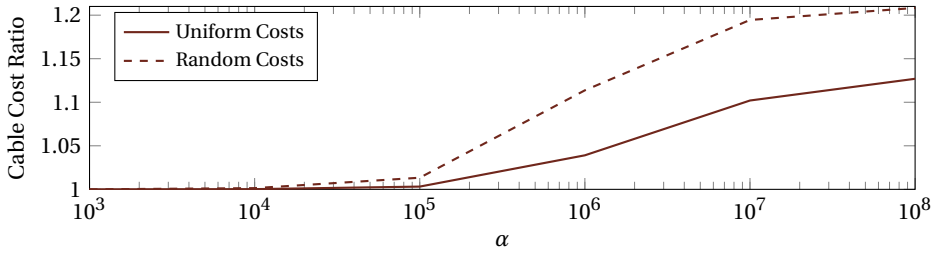


Figure 7.6: Mean cable cost of a route between two random nodes (v_1, v_2) computed by simulated annealing divided by the cable cost of the shortest route between v_1 and v_2.

connections fail due to earthquakes *per year*. In practice, there will be some events where a large part of the network is disconnected at once, and in most years none or nearly none of the connections are affected.

7.6.1. CONNECTING NODE PAIRS

We first consider sub-problem 1: finding the optimal route between two given nodes $v_1, v_2 \in \mathcal{V}$. To decrease computation times, we sample 10 sets of 50,000 disasters and use these as representative disaster sets. We randomly select 20 node pairs, and then use our

simulated annealing heuristic to compute a route between these nodes for each sampled disaster set:

- on a cost grid where cost is equal to the distance in km (**uniform cost case**);
- and on grids where the cost of traversing a cell is a random uniformly distributed value from [0,2) times the traversed distance (in km) (**random cost case**).

We repeat this experiment (for the same disaster sets and node pairs) for different α and evaluate all routes on the *full* disaster set. We do not consider the time required to compute the initial failure states and impacts[2], and compute these separately before running any experiments.

The lowest α we consider is $\alpha = 10^3$. For such an α, a network operator would only be willing to install up to around 14 km of additional cable to *completely* mitigate the impact of all potential disasters. The maximum α we consider is 10^8. Here, an operator would be willing to install around 100,000 times as much cable to achieve the same goal.

As can be seen in Fig. 7.4, computation times increase with α, but stay within 30 seconds even for an α as high as 10^8. At lower α, our approach only needs to test a few, small restrictions, but for higher α the number and size of considered restrictions, and thus the runtime, rises.

In Fig. 7.5, we compare the mean reduction in expected disaster impact due to the routes computed by our simulated annealing heuristic to that of the shortest route. As we are connecting *random* node pairs, we can not expect a major decrease in expected impact. Nevertheless, the improvement of a disaster-aware route over the shortest route is quite impressive. For random costs, the mean reduction in disaster impact due to adding a new route to the network is improved by more than 50% just by adding a small detour to the cable route. This is on top of any planned benefits of the cable in terms of, e.g., capacity. For lower α, the mean cost of connecting two random nodes is negative, and it is not worth it to deviate from the shortest path (see Fig. 7.6). However, as α increases, the simulated annealing solution starts outperforming the shortest path. Note that in practice, network operators will have enough time to compute routes based on the full representative disaster set, which would result in an even larger improvement over the shortest route.

7.6.2. GLOBAL SOLUTION

Next, we consider the Disaster-Aware Network Augmentation Problem itself and try to find optimal link triples across *all* node pairs. We run 10 experiments on the full disaster set: 5 for the uniform cost case, and 5 for the random cost case. In each experiment, we add new cable connections to the network in a greedy fashion (i.e., we iteratively compute and add the next solution for the Disaster-Aware Network Augmentation Problem to the network) until doing so would not be worth it anymore (cost $\geq \alpha E[M(\mathscr{E}^+_{\mathscr{G}}(D))]$. As before, we exclude the time required to compute initial failures.

Fig. 7.7 shows the mean computation times of the first link triple. Again, computation times increase with α. Using simulated annealing, we manage to find a solution

[2] Roughly 310 (26) seconds for the full Interoute disaster set on 1 thread (all cores), and around 35 (3) seconds on the smaller disaster sets.

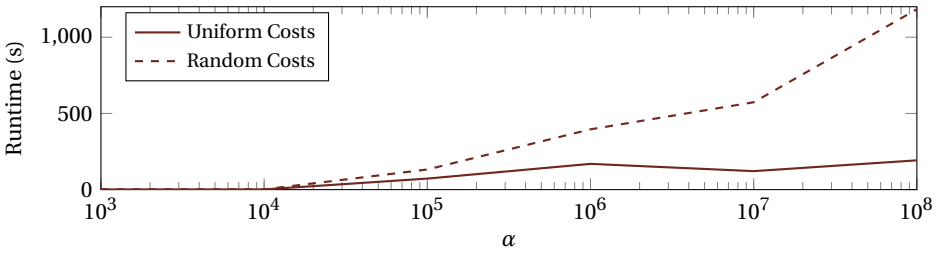


Figure 7.7: Mean simulated annealing computation time of a new link triple (Disaster-Aware Network Augmentation Problem).

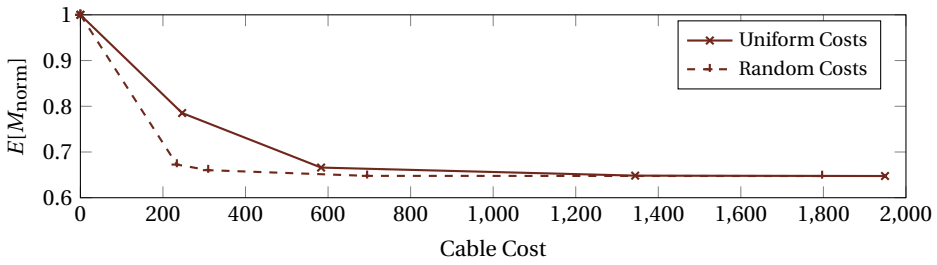


Figure 7.8: Mean expected impact after greedily augmenting the network using simulated annealing divided by the expected impact on the initial topology, against the total cable cost.

within 20 minutes, even for $\alpha = 10^8$. This is fast enough for network operators to vary α and compare different potential routes.

By repeating the same experiment for different α, we can get an idea of how much we can reduce the expected impact given a certain cable budget. Fig. 7.8 shows the mean normalized reduction in impact against the mean total cable cost of greedily augmenting Interoute. By adding new cable connections to the network, we greatly reduce the expected impact of disasters. The biggest reduction in impact comes from some cheaper, very effective cable connections. If we want to reduce the expected impact even further, we need to invest progressively more for smaller reductions in impact.

7.6.3. RESILIENCE AGAINST NEW DISASTERS

Our algorithms extend network topologies based on a set of representative disasters, $\mathscr{D}$. This raises the question of how our new routes perform on disasters that are *not* included in this input set $\mathscr{D}$. Does our approach actually increase the resilience of the network to a whole class of disasters, or does it overfit routes to the set of input disasters?

To answer this question, we take an approach that is similar to 10-fold cross-validation: (1) We randomly split our disaster set, $\mathscr{D}$, into 10 groups, or folds, of disasters; (2) For each of these groups, $\mathscr{D}^*$, we greedily augment Interoute by applying our simulated annealing algorithm to $\mathscr{D} \setminus \mathscr{D}^*$ until there is no improvement in cost anymore and (3) compute the expected impact of $\mathscr{D}^*$ on this augmented network (re-weighting probabilities where required). For the purpose of this experiment, we assume uniform cable costs and

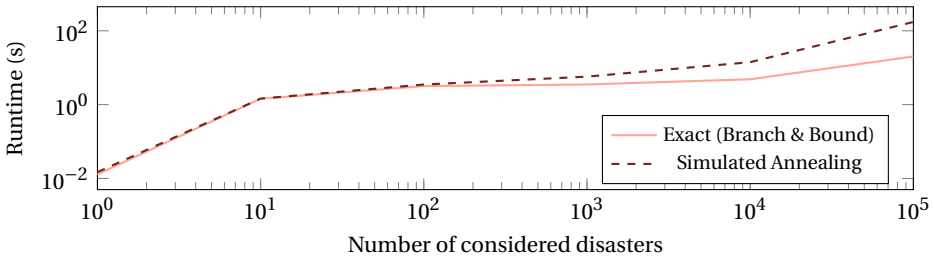


Figure 7.9: Mean computation time of the first new link triple for Interoute against the number of considered disasters. $\alpha = 5,000,000$.

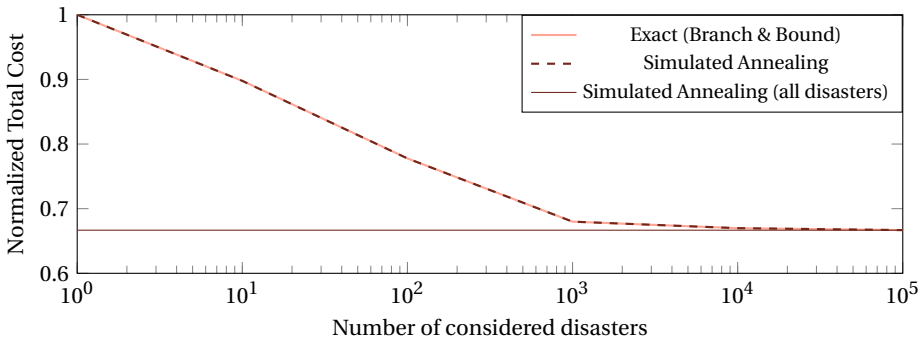


Figure 7.10: Mean total cost of a greedily computed set of new Interoute cable routes divided by the cost of not augmenting the network ($(\alpha E[M(\mathscr{E}^+_{\mathscr{G}}(D))])$) against the number of considered disasters. $\alpha = 5,000,000$.

7

set $\alpha = 5,000,000$.

If we take the full set of 454,433 disasters into account, the greedy simulated annealing approach reduces the expected impact from around 0.01411 to 0.00916 by adding 3 new cable connections with a total length of around 1235 km. In comparison, the average expected impact over all 10 folds is around 0.00917. This is achieved by adding 3 new connections with an average total length of around 1232 km. We conclude that as long as the set of input disasters is representative of the disasters we want to protect the network against, the routes we compute based on $\mathscr{D}$ also manage to effectively reduce the impact of disasters that were not included in $\mathscr{D}$.

7.6.4. NUMBER OF REPRESENTATIVE DISASTERS

We have shown that by increasing the resilience of a network to a set of representative disasters, we also increase its resilience to disasters that were not explicitly considered. But how many representative disasters should a network operator include in their input set $\mathscr{D}$? In this section, we study the effect of $|\mathscr{D}|$ on the reduction in expected impact, as well as on the computation time of our approaches.

We start with the set of 454,433 disasters, $\mathscr{D}$, and will treat this as the set of all potential disasters that can affect the network. Clearly, if we sample disasters from $\mathscr{D}$, the set of sampled disasters is representative of $\mathscr{D}$. Thus, we can create a representative disaster

set of any size simply by sampling disasters from our initial disaster set.

To be more precise, to create an input set $\mathscr{D}^*$ of N disasters, we sample N disasters from $\mathscr{D}$ with replacement (where the probability of sampling a disaster $d \in \mathscr{D}$ is $P(d)$) and assign each disaster a probability of $\frac{1}{N}$. We then apply our algorithms on $\mathscr{D}^*$ to greedily extend Interoute with new cable connections until there is no benefit in cost. Finally, we compute the expected impact of the full set $\mathscr{D}$ on this augmented network to check how well we have managed to increase the resilience of the network. We set $\alpha = 5,000,000$ and repeat this process 20 times for each N.

As can be seen in Fig. 7.9, the runtime of both algorithms increases with the number of considered disasters. Remarkably, at these representative disaster set sizes, the branch & bound algorithm is faster than the simulated annealing algorithm, and both manage to find the initial link triple in less than 3 minutes. In comparison, even if $\alpha = 100,000$, the branch & bound algorithm takes more than an hour to compute a link triple for the full 454,433 disasters input set.

Fig. 7.10 shows the relative improvement in total cost[3] over not augmenting the network against the number of considered disasters. We see that the costs of the link triples found by simulated annealing lie very close to that of the link triples found by the branch & bound algorithm. Importantly, we confirm that we do indeed need approaches that work for larger disaster set sizes (of at least 1,000 disasters), as the total cost drastically decreases as we increase the size of the representative disaster set. However, we also spot some opportunities: by reducing the disaster set size from 454,433 to 100,000 or even 10,000 we significantly reduce computation times, while not sacrificing much of the cost of the final result.

7.7. RELATED WORK

For an overview of overall strategies for increasing the survivability of communication networks to disasters, we refer the reader to a survey conducted by Gomes et al. [201].

Variations of what we call sub-problem 1, finding an optimal cable path while taking into account potential disasters, have been studied extensively [202–209]. As these works focus entirely on the path or topology of a single cable connection, they do not take into account the network-wide impact of disasters and solely consider the damage disasters can do to this single cable. Because our algorithms decouple path planning and finding the optimal restriction, approaches such as [204, 205] can be straightforwardly incorporated and used for finding restricted shortest routes. This would extend our approach by allowing network operators to include costs that depend on the length of the cable segment intersecting disaster regions, such as repair rates or even shielding [205].

Building upon the spine concept introduced in [210], Garrote et al. gave a heuristic for the obstacle-avoiding Euclidean Steiner tree problem [211]. In contrast to our approach, Garrote et al. aim to design a high-availability spine that avoids disaster-prone areas, and do not explicitly consider the impact of disasters on network connectivity metrics.

Cao et al. gave a heuristic for optimizing cable costs of a planar N-node topology under constraints on the disconnection probability of any node in the network [71]. The

[3] The expected impact on the augmented network multiplied by α plus the total cable cost of all link triples.

main difference with our work is that we consider the augmentation of an existing network, while C. Cao et al. considered the design of an entirely new submarine network topology. Furthermore, where their heuristic only considered uniformly distributed disk disasters and simplified cable costs, our disaster and cable cost models are more general.

Given a desired topology, pre-computed set of candidate cable routes, and the probabilities of failure of each of these routes, Msongaleli et al. formulated the optimization of submarine cable deployments under disasters as an integer linear optimization problem [73]. Their simulations suggest that disaster-aware submarine topologies could potentially save society billions of dollars.

Tran and Saito proposed an interesting heuristic for optimizing a weighted set of end-to-end disconnection probabilities under cable length constraints by either recomputing the routes of existing links [86] or augmenting the network by adding new links [87] based on actual seismic hazard data. To compute a set of link triples, they first compute a set of potential candidate routes for each considered link, and then feed these to a dynamic programming algorithm for the global optimization problem. Compared to our approach, their earthquake disaster model is more detailed and takes into account link failure probabilities. This does come at a cost, as their approach does not scale well to larger network sizes or disaster sets. In particular, computing the end-to-end disconnection probabilities in their evaluation metric is a well-known NP-hard problem for even a single disaster [18]. Furthermore, although the authors did limit the paths a cable was allowed to take to streets, their approach is based on a uniform cable-cost scenario.

In contrast to our approach, none of the previous work can be applied to a large set of disaster inputs *and* optimize cable routes for a network-wide impact metric, let alone incorporate detailed cable laying costs.

7.8. CONCLUSION

In this chapter, we have presented the Disaster-Aware Network Augmentation problem of finding a cost-efficient link triple (pair of nodes and a route between these nodes) to increase the resilience of networks to a large set of representative disasters. The solutions to this problem are Pareto-optimal for expected disaster impact and cable cost.

As the Disaster-Aware Network Augmentation Problem is NP-hard, we have provided both an exact algorithm, as well as a heuristic. The main idea behind our algorithms is to split the problem of finding a disaster-aware route into the problem of deciding which disasters to avoid (a *restriction*) and finding a *restricted shortest route* avoiding these disasters.

We have demonstrated the effectiveness of our algorithms by computing disaster-aware cable connections for a real topology using actual seismic hazard data. Using our approach, operators "going the extra mile" can increase the disaster-resilience of their network by adding a cost-efficient selection of new cable connections. By making disaster-aware design decisions, instead of planning based on cable costs only, network operators can simultaneously increase the capacity and disaster resilience of their network.

8

CONCLUSION

In this thesis we addressed the disaster resilience of communication networks. In particular, we considered the following problem statement:

How to create scalable, data-driven methods for assessing and improving the resilience of communication networks to natural disasters.

8.1. DATA-DRIVEN DISASTER RESILIENCE ASSESSMENT

In Chapter 2, we assessed the risk of earthquakes to Internet Exchange Points (IXPs). We found that a large number of IXP facilities are under threat of earthquakes. However, we showed that by selectively peering at multiple facilities, network operators can greatly increase the resilience of both their own networks, as well as that of the Internet as a whole. While the rest of our thesis is more general, this chapter can be seen as more of a case study in data-driven disaster resilience assessment, and showcases both the global availability of earthquake data, and how this data can be applied to the study of communication networks.

We provided both a single-disaster (Chapter 3) and successive disasters (Chapter 4) framework for assessing the resilience of communication networks to natural disasters. While these frameworks differ in both disaster model and algorithms, they both incorporate disaster data by relying on the same main principle: assessing the disaster resilience of a network based on a large set of representative disaster scenarios. The aim of this approach was to be flexible, and be applicable to a wide variety of different disasters and datasets. In this thesis, we have demonstrated the applicability of our approach to various datasets, with different properties:

- A set of 655 earthquake scenarios, obtained through J-SHIS [88] (chapters 3, 4, 5). Each scenario is provided as a map of grid cells and the intensity of shaking at each of these grid cells. We generated disaster regions by applying a threshold on the shaking intensity. The resulting disaster regions can be quite complex, and can even contain holes.

- A set of 1,649 tropical cyclone tracks from IBTrACS [168] (chapters 4, 5). Due to a lack of data, the disaster regions were estimated by mapping the strike circle to each hurricane track. The resulting disaster regions are essentially unions of hippodromes.

- A single hurricane track prediction [182] (Chapter 6). The disaster set was created by a simplified Monte Carlo process.

- A set of 1,196,037 disk-shaped regional failures, that was custom-built by seismologists for analyzing the resiliency of communication networks to earthquakes [109] (Chapter 7).

Our single-disaster framework can quickly compute the distributions of both the state of the network after and the impact of a random disaster, given an input set of disaster regions. These distributions give a wealth of information about the resilience of the network, and can be fed to visualization tools and used as input for other algorithms. In Chapter 4, we showed that the probability of successive disasters (a disaster that strikes the network while it is still being restored from a previous disaster) can be significant, and that successive disasters can be much more detrimental to the performance of a network than any single disaster. Our successive-disaster framework takes up more computation and memory resources than our single-disaster framework, but can provide essential information on the probability and impact of successive disasters on the network.

After a disaster, a network is much more vulnerable to failures. In Chapter 5, we considered the possibility of a follow-up attack after a disaster. We extended our successive disaster framework to incorporate attacks, and demonstrated that even a single targeted node failure can greatly exacerbate the impact of a natural disaster. Fortunately, our experiments indicated that network operators can protect themselves against such attacks by applying repair strategies that take potential follow-up attacks into account. This greatly reduces the impact of potential follow-up attacks, with almost no impact on network performance.

8.2. DATA-DRIVEN DISASTER RESILIENCE IMPROVEMENT

By itself, assessing the disaster resilience of a network is of limited use. To improve resilience, stakeholders also need to take action based on these assessments. Resilience assessments can help stakeholders prepare for potential disasters, and to spot any weak points in the network. Furthermore, resilience assessments can be used to evaluate and compare the effect of changes to the network, as well as compare different strategies for improving resilience.

One example of how a resilience assessment can help improve resilience is selecting a repair strategy. In chapters 4, 5, and 6, we demonstrated how one can adapt our assessment frameworks to evaluate and compare different repair strategies. Essentially, by computing the impact of a random disaster under different repair strategies or under different conditions (e.g. the number of repair crews), stakeholders can evaluate the effects of different strategies and decisions before a disaster has even struck the network.

Some decisions, however, are simply too complex to rely solely on resilience assessment. In Chapter 7, we addressed the problem of network augmentation. In particular, we studied the problem of finding cost-efficient, disaster-aware cable routes. Our aim was that, by adding these kind of connections to a network, the network would become more resilient to disasters by adding geographical redundancy. Given the enormous variety of disasters that could potentially strike a network, finding such routes manually is challenging. Thus, we proposed algorithms that, given a metric of disaster impact and cable cost, automatically find optimal disaster-aware cable routes. These algorithms take disaster data into account in the same way as our assessment frameworks do: by basing their decisions on a large set of representative disaster scenarios.

Since the final decision on adding new cable connections to a network still needs to be made by stakeholders, and not by an automated system, our algorithms can compute multiple different routes (adjusting the trade-off between cost and disaster resilience) and provide both the cable route itself, as well as an explanation of why the cable takes a specific route. In this manner, our algorithms can be seen as automated advisors, which, together with resilience assessment tools, can help stakeholders make informed, disaster-aware decisions.

8.3. SCALABILITY

Our approach for assessing the impact of a single disaster on the network (Chapter 3) and assessing repair strategies (Chapter 6) scales well with both the number of disasters, as well as the size of the network itself. In fact, it is used as an optimization step in the algorithms we proposed in Chapter 7. One of the key insights behind this approach is that the number of failure sets (i.e., the unique combinations of failed network components) tends to be much smaller than the number of considered disasters. This has proven to be true for all disaster sets we considered in this thesis.

The runtime and memory usage of the exact algorithms proposed in Chapter 4, for computing the probability and impact of successive disasters, potentially grows exponentially with the size of the network and the number of disasters. Thus, we also proposed a Monte Carlo method for estimating the probability and impact of successive disasters. The Monte Carlo method scales well with the number of disasters and the size of the network. However, both the exact algorithms and the Monte Carlo method do share one weakness: Runtime grows quickly with the probability of successive disasters. This should be of limited concern, since disasters, and by extension successive disasters, are relatively infrequent.

In Chapter 7, we proved that the Disaster-Aware Network Augmentation Problem is NP-hard. Nevertheless, our exact algorithm for this problem was successfully demonstrated on a disaster set of 10,000 disasters. For cases where the exact approach is too time-consuming, we also proposed a heuristic based on simulated annealing. Technically, the search space of this heuristic increases exponentially with the number of considered disasters. However, the heuristic has been successfully applied to the full set of 1,196,037 disasters[1].

[1] Of which 454,433 (38%) disasters were explicitly considered by the algorithm, since all other disasters did not inflict any damage to the network.

Overall, for each problem we addressed in this thesis, we either proposed an exact algorithm or a heuristic that can be applied to large disaster sets.

8.4. FUTURE WORK

Resilience Assessment A limitation of our work is that we assume that a disaster is a static, instant event. This is not an unreasonable assumption, as the timescale at which a disaster damages a network tends to be much smaller than the timescale of network repair. However, it does limit the incorporation of potential follow-up disasters that strike the network much later than the main disaster, such as some after-shocks. After the earthquake with magnitude 9.0 of the coast in Japan in 2011, the area was still struck by earthquakes with a magnitude of above 7 for up to a month later (and by smaller after-shocks for even longer).

Not all regional failures can be modeled as static. During a power outage, for example, network components may fail one by one as backup power runs out - and may recover one by one as power is gradually restored. During extreme weather, network components may be turned off as a precaution (e.g., to prevent fire during flooding), and can then be turned on again when the weather passes. Since weather gradually moves over the network region, this may be seen as a dynamic regional failure. Assessing the impact of dynamic regional failures on a network will require novel approaches that consider the hazard and the state of the network over time.

Resilience Improvement Our experiments in Chapter 5 demonstrate that the right choice of repair strategy can greatly reduce the impact of follow-up attacks. If this also applies to successive disasters, a disaster-aware repair strategy, which anticipates potential successive disasters (including after-shocks), could mitigate the impact of successive disasters on a network. Potentially, such a repair strategy would remove any need to take successive disasters into account when designing the network.

The framework we proposed in Chapter 7 only finds one cable connection at a time; collections of cable connections can only be calculated in a greedy fashion. More research is required on how to extend this approach to be able to effectively compute an optimal (or near-optimal) selection of multiple cable connections. An interesting extension to the framework would be the addition of nodes. Adding more nodes to a network does not necessarily improve its resilience, but could reduce costs when adding multiple disaster-aware cable connections (and repeaters might even be required when adding longer connections).

Scalability In this thesis, we did not address the data structure used to store disasters. Iqbal and Kuipers proposed storing disaster regions in an R-tree (a tree data structure of minimal bounding rectangles) [104]. More generally, one could store disaster regions in a Bounding Volume Hierarchy (BVH). A BVH is a tree structure, where each node of the tree is a bounding volume of all its children. The objects stored in the BHV form the leaves of the tree. Using a BHV, one can quickly retrieve all objects that intersect a given object, since, if a node does not intersect the object, neither do its children. Storing disaster regions (or network components) into a BHV could speed up the computation of

failure sets (but might make efficiently parallelizing this operation more complex). In addition, early experiments indicate that storing all disasters of interest into even a naively implemented BHV can speed up the computation of cable routes (in both algorithms proposed in Chapter 7) by around 20%.

Self-Protecting Network Lately, there is more and more interest into so-called self-driving networks [212, 213]. Analogous to self-driving cars, researchers and industry envision a network where traffic control is completely automated and tightly integrated with network measurement, relying on machine-learning and large-scale data analytics. While disaster resilience involves many actions outside the control of an automated network (such as network design itself), we envision a disaster-aware self-protecting network, which continuously anticipates potential disasters; adapts to changing circumstances; and advices network operators on current risks, network design, and disaster recovery. The methods introduced in this thesis form an early step towards this vision.

ACKNOWLEDGEMENTS

Throughout my PhD, I have had the privilege of being supported by a large number of people. Some of whom I had the pleasure of meeting as a result of my PhD itself, and some of whom I already knew before starting this journey. I would like to use this opportunity to thank all of you. Without your support, completing this thesis would have been much more difficult, if not impossible.

First, I would like to thank my promoters. When Fernando offered me a position as a PhD-student, he did so in a subtle way. Paraphrased: "Depending on your plans, I may need to make some preparations". Unfortunately for Fernando, subtlety is not my strong point, and I remember telling my friends that I thought I was just offered a position, but was not quite sure. I have not regretted the decision to accept Fernando's offer.

Fernando is an excellent supervisor, advisor, and co-author. I appreciate his honest advice. As a supervisor, it always felt like Fernando puts the interests of his PhD-students first, sometimes above those of himself. Overall, I greatly enjoy(ed) my time working under Fernando, both as a PhD-student and as a postdoc.

My first interaction with Koen was after giving a presentation about my master's thesis project. Koen gave some helpful criticism, and joked that I was lucky that he was not part of my defense committee. Of course, Fernando, not one to ignore such a great opportunity, promptly decided to add Koen to the committee. Now, approximately 5 years later, Koen is judging my work again; this time as my promoter and doctoral committee member. In a way, I have been (un)lucky twice.

Over the years, I have come to greatly appreciate Koen's feedback. Koen has a gift for quickly finding the weak points of even the most carefully drafted plan. My discussions with him have helped me improve both my research itself, as well as my ability to communicate my work to others. Koen's advice on my earliest papers still shapes my writing to this day.

I would like to thank all my colleagues at the ENS group for making my time there so enjoyable. Without you, my journey would have been much more boring. When I started my PhD, Belma, Stef, and Eric immediately made me feel welcome, advised me, and helped me integrate in the rest of the group. ENS counted 3 PhDs who tended to stay till late in the evening: Eric, Nikos, and me. I fondly remember our friendly talks and debates in Eric's office. Vito, thank you for not abandoning me on a mountain. Talia, thanks for sending me a steady supply of pictures of your pets. I strongly believe that, out of all offices in ENS, the office of Belma, Antonia, Renan, and me was the most fun and well-informed. I am looking forward to spending another year with the new members of LOIS: Gabe and Adrian. Sadly, we were not able to see each other much during the Covid lockdown. I would like to thank Jasper for joining me in forming the ultimate counterstrike team. Thanks to you, my lockdown was a lot less lonely.

I do not know if I would have been able to handle our university's bureaucratic system by myself. I would like to thank all the current and past support staff of ENS. Not

only for their ability to deal with all of my (sometimes peculiar) administrative troubles, but also for all the friendly talks we enjoyed.

Near the beginning of my PhD I got the opportunity to spend a week as a guest at the NTNU in Trondheim. I would like to thank Poul and Bjarne for hosting me. Although short, I enjoyed my time in Trondheim, and our discussions gave me new insights during a critical part of my PhD.

When I first approached Niek about designing the cover of this thesis, I confidently stated I would first try to design a cover myself. Back then, I already predicted (and told him) I would realize I need his help way too late, and would ask him to design a cover last-minute. I would like to apologize to Niek for fulfilling my own prediction and would like to thank him for doing a phenomenal job on the cover in very little time.

Throughout my PhD I have gotten the opportunity to collaborate with a number of researchers. I would like to thank each of you for making me a more complete researcher.

Originally, I never planned to specialize in networks. I would like to thank Niels, a former PhD-student under Fernando, for guiding me into this direction.

I would like to thank the committee members for their time and effort.

Last but not least, I would like to thank my family and friends. During a PhD, it helps a great deal to have people to complain to. Furthermore, all our activities, from playing games, sports, and watching television to birthday parties and dinners, have made the past years much more enjoyable. Without your support, I would not have been able to finish my PhD. In particular, I would like to thank my parents. Thank you for always supporting me in every decision I made, even when I decided to move dangerously close to Rotterdam. Without you, I would not be where I am today.

Curriculum Vitæ

Jorik Oostenbrink

27-04-1993	Born in Groningen, the Netherlands.

Education

2005–2011	VWO Dr Nassau College, Assen
2011–2015	Bachelor Industrial and Applied Mathematics Delft University of Technology
2011–2015	Bachelor Computer Science & Engineering Delft University of Technology
2015–2017	Master Computer Science Delft University of Technology

LIST OF PUBLICATIONS

10. J. Oostenbrink and F.A. Kuipers, *A Global Study of the Risk of Earthquakes to IXPs*, Proc. of IFIP Networking 2022, Catania, Italy, June 13-16, 2022

9. J. Oostenbrink and F.A. Kuipers, *Going the Extra Mile with Disaster-Aware Network Augmentation*, Proc. of IEEE INFOCOM, May 10-13, 2021

8. B. Vass, J. Tapolcai, Z. Heszberger, J. Biro, D. Hay, F.A. Kuipers, J. Oostenbrink, A. Valentini, and L. Ronyai, *Probabilistic Shared Risk Link Groups Modelling Correlated Resource Failures Caused by Disasters*, IEEE Journal on Selected Areas in Communications, vol. 39, 9, Sept. 2021

7. B. Vass, J. Tapolcai, D. Hay, J. Oostenbrink, and F.A. Kuipers, *How to Model and Enumerate Geographically Correlated Failure Events in Communication Networks*, Chapter in Guide to Disaster-Resilient Communication Networks, edited by J. Rak and D. Hutchison, Springer, July 2020, ISBN: 9783030446840

6. B. Turkovic, J. Oostenbrink, F.A. Kuipers, I. Keslassy, and A. Orda, *Sequential Zeroing: Online Heavy-Hitter Detection on Programmable Hardware*, Proc. of IFIP Networking 2020, Paris, France, June 22-25, 2020

5. J. Oostenbrink and F.A. Kuipers, *A Moment of Weakness: Protecting Against Targeted Attacks Following a Natural Disaster*, ACM SIGMETRICS Performance Evaluation Review (special issue from the 3rd ACM SIGMETRICS International Workshop on Critical Infrastructure Network Security), vol. 47, 4, pp. 12-15, April 2020

4. A. Valentini, B. Vass, J. Oostenbrink, L. Csak, F.A. Kuipers, B. Pace, D. Hay, and J. Tapolcai, *Network Resiliency Against Earthquakes*, Proc. of the 11th International Workshop on Resilient Networks Design and Modeling (RNDM 2019), October 14-16, 2019

3. J. Oostenbrink and F.A. Kuipers, *The Risk of Successive Disasters: A Blow-by-Blow Network Vulnerability Analysis*, Proc. of IFIP Networking 2019, Warsaw, Poland, May 20-22, 2019

2. J. Oostenbrink, F.A. Kuipers, P. Heegaard, and B. Helvik, *Evaluating Local Disaster Recovery Strategies*, ACM SIGMETRICS Performance Evaluation Review (special issue from the 2nd ACM SIGMETRICS International Workshop on Critical Infrastructure Network Security), vol. 46, no. 2, pp. 62-66, September 2018

1. J. Oostenbrink and F.A. Kuipers, *Computing the Impact of Disasters on Networks*, ACM SIGMETRICS Performance Evaluation Review (special issue from the 1st ACM SIGMETRICS International Workshop on Critical Infrastructure Network Security), vol. 45, no. 2, pp. 107-110, September 2017

REFERENCES

REFERENCES

[1] "Measuring digital development: Facts and figures 2020," International Telecommunication Union - Development Sector, CH-1211 Geneva Switzerland, 2020.

[2] (2013, February) Presidential policy directive – critical infrastructure security and resilience. The White House. [Online]. Available: obamawhitehouse.archives.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil

[3] "Inspectierapport onbereikbaarheid van 112 op 24 juni 2019," Inspectie Justitie en Veiligheid, Agentschap Telecom and Inspectie Gezondheidszorg en Jeugd, June 2020.

[4] (2018, October) Indonesia earthquake and tsunami: How warning system failed the victims. BBC. [Online]. Available: https://www.bbc.com/news/world-asia-45663054

[5] Sulawesi earthquake. Télécoms Sans Frontières. Accessed: 13-10-2021. [Online]. Available: https://www.tsfi.org/en/our-missions/disaster-response/sulawesi-earthquake

[6] (2017, September) De haken en ogen van de nederlandse hulp aan het getroffen sint-maarten. NOS. [Online]. Available: https://nos.nl/artikel/2191987-de-haken-en-ogen-van-de-nederlandse-hulp-aan-het-getroffen-sint-maarten

[7] (2017, September) Ongebruikelijk dat humanitaire hulp zo traag op gang komt. NOS. [Online]. Available: https://nos.nl/artikel/2192443-ongebruikelijk-dat-humanitaire-hulp-zo-traag-op-gang-komt

[8] (2021, August) How big telecom killed rules that would have prevented hurricane ida outages. Vice. [Online]. Available: https://www.vice.com/en/article/dyv9j7/how-big-telecom-killed-rules-that-would-have-prevented-hurricane-ida-outages

[9] (2021, August) Hurricane ida takes down new orleans 911. Communications Daily. [Online]. Available: https://communicationsdaily.com/news/2021/08/31/Hurricane-Ida-Takes-Down-New-Orleans-911-2108300054

[10] H. Toya and M. Skidmore, "Cellular telephones and natural disaster vulnerability," *Sustainability*, vol. 10, no. 9, 2018. [Online]. Available: https://www.mdpi.com/2071-1050/10/9/2970

[11] IPCC, *Summary for Policymakers. In: Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, V. Masson-Delmotte, P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou, Eds. Cambridge University Press, 2021.

[12] M. Timms. (2020) AT&T outage: Internet, 911 disrupted, planes grounded after nashville explosion. Nashville Tennessean. [Online]. Available: https://eu.tennessean.com/story/news/local/2020/12/25/att-outage-internet-down-hours-after-nashville-explosion/4045278001/

[13] "Stroomstoring noord-holland 27 maart 2015 - lessen uit de crisisbeheersing en telecommunicatie," Inspectie Veiligheid en Justitie and Agentschap Telecom, June 2016.

[14] M. Kazama and T. Noda, "Damage statistics (summary of the 2011 off the Pacific coast of Tohoku earthquake damage)," *Soils and Foundations*, vol. 52, no. 5, pp. 780 – 792, 2012, special Issue on Geotechnical Aspects of the 2011 off the Pacific Coast of Tohoku Earthquake. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0038080612000947

[15] Z. E. Khaled and H. Mcheick, "Case studies of communications systems during harsh environments: A review of approaches, weaknesses, and limitations to improve quality of service," *International Journal of Distributed Sensor Networks*, vol. 15, no. 2, 2019. [Online]. Available: https://doi.org/10.1177/1550147719829960

[16] A. J. Hickford, S. P. Blainey, A. O. Hortelano, and R. Pant, "Resilience engineering: theory and practice in interdependent infrastructure systems," *Environment Systems and Decisions*, vol. 38, no. 3, pp. 278–291, 2018.

[17] J. P. G. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, "Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines," *Comput. Netw.*, vol. 54, no. 8, p. 1245–1265, Jun. 2010. [Online]. Available: https://doi.org/10.1016/j.comnet.2010.03.005

[18] M. O. Ball, "Computational complexity of network reliability analysis: An overview," *IEEE Transactions on Reliability*, vol. 35, no. 3, pp. 230–239, 1986.

[19] D. Bienstock, "Some generalized max-flow min-cut problems in the plane," *Mathematics of Operations Research*, vol. 16, no. 2, pp. 310–333, 1991.

[20] S. Neumayer, A. Efrat, and E. Modiano, "Geographic max-flow and min-cut under a circular disk failure model," *Computer Networks*, vol. 77, pp. 117–127, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128614003880

[21] S. Trajanovski, F. A. Kuipers, A. Ilić, J. Crowcroft, and P. Van Mieghem, "Finding critical regions and region-disjoint paths in a network," *IEEE/ACM Transactions on Networking*, vol. 23, no. 3, pp. 908–921, 2015.

[22] P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, "Network vulnerability to single, multiple, and probabilistic physical attacks," in *2010 - MILCOM 2010 MILITARY COMMUNICATIONS CONFERENCE*, 2010, pp. 1824–1829.

[23] ——, "The resilience of wdm networks to probabilistic geographical failures," *IEEE/ACM Transactions on Networking*, vol. 21, no. 5, pp. 1525–1538, 2013.

[24] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of the fiber infrastructure to disasters," *IEEE/ACM Transactions on Networking*, vol. 19, no. 6, pp. 1610–1623, 2011.

[25] S. Banerjee, S. Shirazipourazad, and A. Sen, "Design and analysis of networks with large components in presence of region-based faults," in *2011 IEEE International Conference on Communications (ICC)*, 2011, pp. 1–6.

[26] S. Banerjee, S. Shirazipourazad, P. Ghosh, and A. Sen, "Beyond connectivity - new metrics to evaluate robustness of networks," in *2011 IEEE 12th International Conference on High Performance Switching and Routing*, 2011, pp. 171–177.

[27] X. Long, D. Tipper, and T. Gomes, "Measuring the survivability of networks to geographic correlated failures," *Optical Switching and Networking*, vol. 14, pp. 117–133, 2014, special Issue on RNDM 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1573427714000381

[28] O. Gold and R. Cohen, "Coping with physical attacks on random network structures," in *2014 IEEE International Conference on Communications (ICC)*, 2014, pp. 1166–1172.

[29] Y. Cheng and J. P. Sterbenz, "Critical region identification and geodiverse routing protocol under massive challenges," in *2015 7th International Workshop on Reliable Networks Design and Modeling (RNDM)*, 2015, pp. 14–20.

[30] J. Tapolcai, L. Rónyai, B. Vass, and L. Gyimóthi, "List of shared risk link groups representing regional failures with limited size," in *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*, 2017, pp. 1–9.

[31] ——, "Fast enumeration of regional link failures caused by disasters with limited size," *IEEE/ACM Transactions on Networking*, vol. 28, no. 6, pp. 2421–2434, 2020.

[32] B. Vass, E. Berczi-Kovacs, and J. Tapolcai, "Enumerating shared risk link groups of circular disk failures hitting k nodes," in *DRCN 2017 - Design of Reliable Communication Networks; 13th International Conference*, 2017, pp. 1–9.

[33] B. Vass, J. Tapolcai, and E. R. Bérczi-Kovács, "Enumerating maximal shared risk link groups of circular disk failures hitting k nodes," *IEEE/ACM Transactions on Networking*, vol. 29, no. 4, pp. 1648–1661, 2021.

[34] B. Vass, L. Németh, and J. Tapolcai, "The earth is nearly flat: Precise and approximate algorithms for detecting vulnerable regions of networks in the plane and on the sphere," *Networks*, vol. 75, no. 4, pp. 340–355, 2020. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/net.21936

[35] Y. Cheng, M. T. Gardner, J. Li, R. May, D. Medhi, and J. P. Sterbenz, "Analysing geopath diversity and improving routing performance in optical networks," *Computer Networks*, vol. 82, pp. 50–67, 2015, robust and Fault-Tolerant Communication Networks. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128615000699

[36] M. T. Gardner, R. May, C. Beard, and D. Medhi, "A geographic multi-topology routing approach and its benefits during large-scale geographically correlated failures," *Computer Networks*, vol. 82, pp. 34–49, 2015, robust and Fault-Tolerant Communication Networks. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128615000729

[37] J. Zhang, E. Modiano, and D. Hay, "Enhancing network robustness via shielding," *IEEE/ACM Transactions on Networking*, vol. 25, no. 4, pp. 2209–2222, 2017.

[38] Y. M. Allawi, D. Lee, and J.-K. K. Rhee, "A wireless link-up augmentation design for disaster-resilient optical networks," *J. Lightwave Technol.*, vol. 33, no. 17, pp. 3516–3524, Sep 2015. [Online]. Available: http://www.osapublishing.org/jlt/abstract.cfm?URI=jlt-33-17-3516

[39] M. F. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of disaster-resilient optical datacenter networks," *Journal of Lightwave Technology*, vol. 30, no. 16, pp. 2563–2573, 2012.

[40] M. Ju, F. Zhou, and S. Xiao, "Disaster-resilient cloud services provisioning in elastic optical inter-data center networks," in *2019 IEEE 27th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2019, pp. 116–124.

[41] Y. Liu, F. Zhou, C. Chen, Z. Zhu, T. Shang, and J.-M. Torres-Moreno, "Disaster protection in inter-datacenter networks leveraging cooperative storage," *IEEE Transactions on Network and Service Management*, vol. 18, no. 3, pp. 2598–2611, 2021.

[42] G. Grebla, A. Efrat, E. Ezra, R. Pinchasi, and S. Sankararaman, "Data recovery after geographic correlated attacks," in *2015 11th International Conference on the Design of Reliable Communication Networks (DRCN)*, 2015, pp. 65–72.

[43] S. Cai, F. Zhou, Z. Zhang, and A. Meddahi, "Disaster-resilient service function chain embedding based on multi-path routing," in *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2021, pp. 1–7.

[44] H. Yu, C. Qiao, V. Anand, X. Liu, H. Di, and G. Sun, "Survivable virtual infrastructure mapping in a federated computing and networking system under single regional failures," in *2010 IEEE Global Telecommunications Conference GLOBECOM 2010*, 2010, pp. 1–6.

[45] G. Sun, H. Yu, L. Li, V. Anand, H. Di, and X. Gao, "Efficient algorithms for survivable virtual network embedding," in *Network Architectures, Management, and Applications VIII*. Optical Society of America, 2010, p. 79890K. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=ACP-2010-79890K

[46] F. Gu, K. Shaban, N. Ghani, S. Khan, M. R. Naeini, M. M. Hayat, and C. Assi, "Survivable cloud network mapping for disaster recovery support," *IEEE Transactions on Computers*, vol. 64, no. 8, pp. 2353–2366, 2015.

[47] J. Tapolcai, Z. L. Hajdú, A. Pašić, P.-H. Ho, and L. Rónyai, "On network topology augmentation for global connectivity under regional failures," in *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*, 2021, pp. 1–10.

[48] F. Iqbal, S. Trajanovski, and F. Kuipers, "Detection of spatially-close fiber segments in optical networks," in *2016 12th International Conference on the Design of Reliable Communication Networks (DRCN)*, 2016, pp. 95–102.

[49] A. d. Sousa, D. Santos, and P. Monteiro, "Determination of the minimum cost pair of d-geodiverse paths," in *DRCN 2017 - Design of Reliable Communication Networks; 13th International Conference*, 2017, pp. 1–8.

[50] P. D. Joshi, A. Sen, D. F. Hsu, S. Hamdioui, and K. Bertels, "Region based containers — a new paradigm for the analysis of fault tolerant networks," in *2017 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, 2017, pp. 1–4.

[51] A. de Sousa, T. Gomes, R. Girão-Silva, and L. Martins, "Minimization of the network availability upgrade cost with geodiverse routing for disaster resilience," *Optical Switching and Networking*, vol. 31, pp. 127 – 143, 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1573427718300687

[52] A. de Sousa and D. Santos, "The minimum cost d-geodiverse anycast routing with optimal selection of anycast nodes," in *2019 15th International Conference on the Design of Reliable Communication Networks (DRCN)*, 2019, pp. 21–28.

[53] T. Feyessa and M. Bikdash, "Geographically-sensitive network centrality and survivability assessment," in *2011 IEEE 43rd Southeastern Symposium on System Theory*, 2011, pp. 18–23.

[54] Y. Cheng, D. Medhi, and J. P. G. Sterbenz, "Geodiverse routing with path delay and skew requirement under area-based challenges," *Networks*, vol. 66, no. 4, pp. 335–346, 2015. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/net.21660

[55] M. T. Gardner, Y. Cheng, C. Beard, J. P. Sterbenz, and D. Medhi, "Provisioning dynamic and critical demand structures for geographically correlated failures," *Annals of Telecommunications*, vol. 73, pp. 111–125, 2018. [Online]. Available: https://doi.org/10.1007/s12243-017-0618-z

[56] M. W. Ashraf, S. M. Idrus, F. Iqbal, and R. A. Butt, "On spatially disjoint lightpaths in optical networks," *Photonic Network Communications*, vol. 36, pp. 11–25, 2018. [Online]. Available: https://doi.org/10.1007/s11107-018-0764-x

[57] D. Santos, T. Gomes, and D. Tipper, "Sdn controller placement with availability upgrade under delay and geodiversity constraints," *IEEE Transactions on Network and Service Management*, vol. 18, no. 1, pp. 301–314, 2021.

[58] J. Wang, J. Bigham, and C. Phillips, "A geographical proximity aware multi-path routing mechanism for resilient networking," *IEEE Communications Letters*, vol. 21, no. 7, pp. 1533–1536, 2017.

[59] R. Girão-Silva, B. Nedic, M. Gunkel, and T. Gomes, "Shared risk link group disjointness and geodiverse routing: A trade-off between benefit and practical effort," *Networks*, vol. 75, no. 4, pp. 374–391, 2020. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/net.21931

[60] S. Neumayer and E. Modiano, "Network reliability under geographically correlated line and disk failure models," *Computer Networks*, vol. 94, pp. 14–28, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128615004740

[61] ——, "Assessing the effect of geographically correlated failures on interconnected power-communication networks," in *2013 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2013, pp. 366–371.

[62] X. Wang, M. Chen, and S. Lu, "Modeling geographically correlated failures to assess network vulnerability," *IEEE Transactions on Communications*, vol. 66, no. 12, pp. 6317–6328, 2018.

[63] J. Liu, X. Jiang, H. Nishiyama, and N. Kato, "Reliability assessment for wireless mesh networks under probabilistic region failure model," *IEEE Transactions on Vehicular Technology*, vol. 60, no. 5, pp. 2253–2264, 2011.

[64] J. Rak, "Measures of region failure survivability for wireless mesh networks," *Wireless Networks*, vol. 21, no. 2, pp. 673–684, 2015. [Online]. Available: https://doi.org/10.1007/s11276-014-0806-y

[65] M. T. Gardner and C. Beard, "Evaluating geographic vulnerabilities in networks," in *2011 IEEE International Workshop Technical Committee on Communications Quality and Reliability (CQR)*, 2011, pp. 1–6.

[66] M. T. Gardner, R. May, C. Beard, and D. Medhi, "Determining geographic vulnerabilities using a novel impact based resilience metric," *Journal of Network*

and Systems Management, vol. 24, pp. 711–745, 2016. [Online]. Available: https://doi.org/10.1007/s10922-016-9383-y

[67] H. Saito, "Geometric evaluation of survivability of disaster-affected network with probabilistic failure," in *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*, 2014, pp. 1608–1616.

[68] ——, "Spatial design of physical network robust against earthquakes," *Journal of Lightwave Technology*, vol. 33, no. 2, pp. 443–458, 2015.

[69] M. Rahnamay-Naeini, J. E. Pezoa, G. Azar, N. Ghani, and M. M. Hayat, "Modeling stochastic correlated failures and their effects on network reliability," in *2011 Proceedings of 20th International Conference on Computer Communications and Networks (ICCCN)*, 2011, pp. 1–6.

[70] P. Das, M. Rahnamay-Naeini, N. Ghani, and M. M. Hayat, "On the vulnerability of multi-level communication network under catastrophic events," in *2017 International Conference on Computing, Networking and Communications (ICNC)*, 2017, pp. 912–916.

[71] C. Cao, M. Zukerman, W. Wu, J. Manton, and B. Moran, "Survivable topology design of submarine networks," *J. Lightwave Technol.*, vol. 31, no. 5, pp. 715–730, Mar 2013. [Online]. Available: http://jlt.osa.org/abstract.cfm?URI=jlt-31-5-715

[72] C. Colman-Meixner, F. Dikbiyik, M. F. Habib, M. Tornatore, C.-N. Chuah, and B. Mukherjee, "Disaster-survivable cloud-network mapping," *Photonic network communications*, vol. 27, no. 3, pp. 141–153, 2014. [Online]. Available: https://doi.org/10.1007/s11107-014-0434-6

[73] D. L. Msongaleli, F. Dikbiyik, M. Zukerman, and B. Mukherjee, "Disaster-aware submarine fiber-optic cable deployment for mesh networks," *J. Lightwave Technol.*, vol. 34, no. 18, pp. 4293–4303, Sep 2016. [Online]. Available: http://jlt.osa.org/abstract.cfm?URI=jlt-34-18-4293

[74] F. Dikbiyik, M. Tornatore, and B. Mukherjee, "Minimizing the risk from disaster failures in optical backbone networks," *J. Lightwave Technol.*, vol. 32, no. 18, pp. 3175–3183, Sep 2014. [Online]. Available: http://www.osapublishing.org/jlt/abstract.cfm?URI=jlt-32-18-3175

[75] S. Yang, S. Trajanovski, and F. A. Kuipers, "Availability-based path selection and network vulnerability assessment," *Networks*, vol. 66, no. 4, pp. 306–319, 2015. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/net.21652

[76] H.-W. Lee, E. Modiano, and K. Lee, "Diverse routing in networks with probabilistic failures," *IEEE/ACM Transactions on Networking*, vol. 18, no. 6, pp. 1895–1907, 2010.

[77] O. Diaz, F. Xu, N. Min-Allah, M. Khodeir, M. Peng, S. Khan, and N. Ghani, "Network survivability for multiple probabilistic failures," *IEEE Communications Letters*, vol. 16, no. 8, pp. 1320–1323, 2012.

[78] S. S. Savas, F. Dikbiyik, M. F. Habib, M. Tornatore, and B. Mukherjee, "Disaster-aware service provisioning with manycasting in cloud networks," *Photonic Network Communications*, vol. 28, pp. 123–134, 2014. [Online]. Available: https://doi.org/10.1007/s11107-014-0457-z

[79] S. Ferdousi, F. Dikbiyik, M. F. Habib, M. Tornatore, and B. Mukherjee, "Disaster-aware datacenter placement and dynamic content management in cloud networks," *J. Opt. Commun. Netw.*, vol. 7, no. 7, pp. 681–694, Jul 2015. [Online]. Available: http://www.osapublishing.org/jocn/abstract.cfm?URI=jocn-7-7-681

[80] S. S. Savas, M. Tornatore, M. F. Habib, P. Chowdhury, and B. Mukherjee, "Disaster-resilient control plane design and mapping in software-defined networks," in *2015 IEEE 16th International Conference on High Performance Switching and Routing (HPSR)*, 2015, pp. 1–6.

[81] E. K. Cetinkaya, D. Broyles, A. Dandekar, S. Srinivasan, and J. P. G. Sterbenz, "Modelling communication network challenges for future internet resilience, survivability, and disruption tolerance: A simulation-based approach," *Telecommunication Systems*, vol. 52, no. 2, pp. 751–766, 2013. [Online]. Available: https://doi.org/10.1007/s11235-011-9575-4

[82] P. E. Heegaard and K. S. Trivedi, "Network survivability modeling," *Computer Networks*, vol. 53, no. 8, pp. 1215–1234, 2009, performance Modeling of Computer Networks: Special Issue in Memory of Dr. Gunter Bolch. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128609000425

[83] M. Gajić, M. Furdek, and P. Heegaard, "A framework for spatial and temporal evaluation of network disaster recovery," in *2020 32nd International Teletraffic Congress (ITC 32)*, 2020, pp. 37–45.

[84] L. Zhong, K. Takano, F. Jiang, X. Wang, Y. Ji, and S. Yamada, "Spatio-temporal data-driven analysis of mobile network availability during natural disasters," in *2016 3rd International Conference on Information and Communication Technologies for Disaster Management (ICT-DM)*, 2016, pp. 1–7.

[85] L. Ma, X. Jiang, B. Wu, A. Pattavina, and N. Shiratori, "Probabilistic region failure-aware data center network and content placement," *Computer Networks*, vol. 103, pp. 56–66, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128616300755

[86] P. N. Tran and H. Saito, "Geographical route design of physical networks using earthquake risk information," *IEEE Communications Magazine*, vol. 54, no. 7, pp. 131–137, 2016.

[87] P. N. Tran and H. Saito, "Enhancing physical network robustness against earthquake disasters with additional links," *J. Lightwave Technol.*, vol. 34, no. 22, pp. 5226–5238, Nov 2016. [Online]. Available: http://www.osapublishing.org/jlt/abstract.cfm?URI=jlt-34-22-5226

[88] Japan seismic hazard information station. [Online]. Available: http://www.j-shis.bosai.go.jp/en/

[89] Shakemap. USGS. [Online]. Available: https://earthquake.usgs.gov/data/shakemap/

[90] B. Eriksson, R. Durairajan, and P. Barford, "Riskroute: A framework for mitigating network outage threats," in *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '13. New York, NY, USA: Association for Computing Machinery, 2013, p. 405–416. [Online]. Available: https://doi.org/10.1145/2535372.2535385

[91] Q. Zheng, G. Cao, T. La Porta, and A. Swami, "Optimal recovery from large-scale failures in ip networks," in *2012 IEEE 32nd International Conference on Distributed Computing Systems*, 2012, pp. 295–304.

[92] S. S. Savas, M. F. Habib, M. Tornatore, F. Dikbiyik, and B. Mukherjee, "Network adaptability to disaster disruptions by exploiting degraded-service tolerance," *IEEE Communications Magazine*, vol. 52, no. 12, pp. 58–65, 2014.

[93] A. Izaddoost and S. S. Heydari, "Risk-adaptive strategic network protection in disaster scenarios," *Journal of Communications and Networks*, vol. 19, no. 5, pp. 509–520, 2017.

[94] B. J. Liu, P. Yu, Q. Xue-song, and L. Shi, "Survivability-aware routing restoration mechanism for smart grid communication network in large-scale failures," *EURASIP Journal on Wireless Communications and Networking*, 2020. [Online]. Available: https://doi.org/10.1186/s13638-020-1653-4

[95] M. Oguz, F. Dikbiyik, and H. S. Kuyuk, "Earthquake preparedness strategies for telecom backbone with integration of early warning systems and optical wdm networks," in *2016 8th International Workshop on Resilient Networks Design and Modeling (RNDM)*, 2016, pp. 181–188.

[96] N.-H. Bao, M. F. Habib, M. Tornatore, C. U. Martel, and B. Mukherjee, "Global versus essential post-disaster re-provisioning in telecom mesh networks," *J. Opt. Commun. Netw.*, vol. 7, no. 5, pp. 392–400, May 2015. [Online]. Available: http://www.osapublishing.org/jocn/abstract.cfm?URI=jocn-7-5-392

[97] N.-H. Bao, M. Tornatore, C. U. Martel, and B. Mukherjee, "Fairness-aware degradation based multipath re-provisioning strategy for post-disaster telecom mesh networks," *J. Opt. Commun. Netw.*, vol. 8, no. 6, pp. 441–450, Jun 2016. [Online]. Available: http://www.osapublishing.org/jocn/abstract.cfm?URI=jocn-8-6-441

[98] C. Colman-Meixner, M. Tornatore, and B. Mukherjee, "Cloud-network disaster recovery against cascading failures," in *2015 IEEE Global Communications Conference (GLOBECOM)*, 2015, pp. 1–5.

[99] R. Zou, H. Hasegawa, and S. Subramaniam, "Drama: Disaster management algorithm with mitigation awareness for elastic optical networks," in *2021 17th International Conference on the Design of Reliable Communication Networks (DRCN)*, 2021, pp. 1–7.

[100] L. Guillen, S. Izumi, T. Abe, and T. Suganuma, "A resilient mechanism for multi-controller failure in hybrid sdn-based networks," in *2021 22nd Asia-Pacific Network Operations and Management Symposium (APNOMS)*, 2021, pp. 285–290.

[101] F. Iqbal and F. Kuipers, "Spatiotemporal risk-averse routing," in *2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2016, pp. 395–400.

[102] L. Ma, W. Su, B. Wu, B. Yang, and X. Jiang, "Early warning disaster-aware service protection in geo-distributed data centers," *Computer Networks*, vol. 180, p. 107419, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128620300074

[103] H. Honda and H. Saito, "Nation-wide disaster avoidance control against heavy rain," *IEEE/ACM Transactions on Networking*, vol. 27, no. 3, pp. 1084–1097, 2019.

[104] F. Iqbal and F. Kuipers, "On centrality-related disaster vulnerability of network regions," in *2017 9th International Workshop on Resilient Networks Design and Modeling (RNDM)*, 2017, pp. 1–6.

[105] J. Tapolcai, B. Vass, Z. Heszberger, J. Bíró, D. Hay, F. A. Kuipers, and L. Rónyai, "A tractable stochastic model of correlated link failures caused by disasters," in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, 2018, pp. 2105–2113.

[106] G. A. Beletsioti, G. I. Papadimitriou, P. Nicopolitidis, and A. N. Miliou, "Earthquake tolerant energy aware algorithms: A new approach to the design of wdm backbone networks," *IEEE Transactions on Green Communications and Networking*, vol. 2, no. 4, pp. 1164–1173, 2018.

[107] A. Agrawal, V. Bhatia, and S. Prakash, "Network and risk modeling for disaster survivability analysis of backbone optical communication networks," *Journal of Lightwave Technology*, vol. 37, no. 10, pp. 2352–2362, 2019.

[108] S. Esposito, A. Botta, M. De Falco, I. Iervolino, A. PESCAPÈ, and A. Santo, "Seismic risk analysis of data communication networks: a feasibility study," in *16th European Conference on Earthquake Engineering*, 2018.

[109] A. Valentini, B. Vass, J. Oostenbrink, L. Csák, F. Kuipers, B. Pace, D. Hay, and J. Tapolcai, "Network resiliency against earthquakes," in *2019 11th International Workshop on Resilient Networks Design and Modeling (RNDM)*, 2019, pp. 1–7.

[110] H. Talebiyan, K. Leelardcharoen, L. Dueñas-Osorio, B. J. Goodno, and J. I. Craig, "Congestion and observability across interdependent power and telecommunication networks under seismic hazard," *Earthquake Spectra*, vol. 37, no. 4, pp. 2892–2919, 2021. [Online]. Available: https://doi.org/10.1177/87552930211026690

[111] K. Mitchell-Wallace, M. Jones, J. Hillier, and M. Foote, *Natural catastrophe risk management and modelling: A practitioner's guide*. John Wiley & Sons, 2017, ch. 4, pp. 297–388.

[112] J. Oostenbrink and F. Kuipers, "A Global Study of the Risk of Earthquakes to IXPs," in *2022 IFIP Networking Conference (IFIP Networking)*, 2022, pp. 1–9.

[113] W. Qiu. (2011) Submarine cables cut after taiwan earthquake in dec 2006. Submarine Cable Networks. [Online]. Available: www.submarinenetworks.com/en/news/cables-cut-after-taiwan-earthquake-2006

[114] "Special report - the NTT groups response to the Great East Japan Earthquake," June 2016.

[115] S. Giovinazzi, A. Austin, R. Ruiter, C. Foster, M. Nayyerloo, N.-K. Nair, and L. Wotherspoon, "Resilience and fragility of the telecommunication network to seismic events," *Bulletin of the New Zealand Society for Earthquake Engineering*, vol. 50, no. 2, pp. 318–328, Jun. 2017. [Online]. Available: https://www.bulletin.nzsee.org.nz/index.php/bnzsee/article/view/84

[116] M. Di Bartolomeo, G. Di Battista, R. di Lallo, and C. Squarcella, "Is it really worth to peer at IXPs? A comparative study," in *2015 IEEE Symposium on Computers and Communication (ISCC)*, 2015, pp. 421–426.

[117] J. Oostenbrink and F. Kuipers, "Computing the impact of disasters on networks," *SIGMETRICS Perform. Eval. Rev.*, vol. 45, no. 2, p. 107–110, Oct. 2017. [Online]. Available: https://doi.org/10.1145/3152042.3152075

[118] S. A. Jyothi, "Solar superstorms: Planning for an internet apocalypse," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, ser. SIGCOMM '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 692–704. [Online]. Available: https://doi.org/10.1145/3452296.3472916

[119] S. Anderson, C. Barford, and P. Barford, "Five alarms: Assessing the vulnerability of us cellular communication infrastructure to wildfires," in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 162–175. [Online]. Available: https://doi.org/10.1145/3419394.3423663

[120] R. Durairajan, S. Ghosh, X. Tang, P. Barford, and B. Eriksson, "Internet atlas: A geographic database of the internet," in *Proceedings of the 5th ACM Workshop on HotPlanet*, ser. HotPlanet '13. New York, NY, USA: Association for Computing Machinery, 2013, p. 15–20. [Online]. Available: https://doi.org/10.1145/2491159.2491170

[121] R. Durairajan, C. Barford, and P. Barford, "Lights out: Climate change risk to internet infrastructure," in *Proceedings of the Applied Networking Research Workshop*, ser. ANRW '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 9–15. [Online]. Available: https://doi.org/10.1145/3232755.3232775

[122] J. Mayer, V. Sahakian, E. Hooft, D. Toomey, and R. Durairajan, "On the resilience of internet infrastructures in pacific northwest to earthquakes." in *Passive and Active Measurement: 22nd International Conference*, 2021, pp. 247–265.

[123] The CAIDA UCSD IXPs Dataset, 2021-07. CAIDA. [Online]. Available: https://www.caida.org/catalog/datasets/ixps

[124] OpenQuake Engine. Global Earthquake Model Foundation. [Online]. Available: https://github.com/gem/oq-engine

[125] J. Woessner, D. Laurentiu, D. Giardini, H. Crowley, F. Cotton, G. Grünthal, G. Valensise, R. Arvidsson, R. Basili, M. B. Demircioglu, S. Hiemer, C. Meletti, R. Musson, A. Rovida, K. Sesetyan, and M. Stucchi, "The 2013 european seismic hazard model: key components and results," *Bulletin of Earthquake Engineering*, vol. 13, pp. 3553–3596, 2015. [Online]. Available: https://doi.org/10.1007/s10518-015-9795-1

[126] Hazard Model for South America. Global Earthquake Model Foundation. [Online]. Available: https://www.globalquakemodel.org/product/south-america-model-2016

[127] T. Allen, J. Griffin, and D. Clark. (2019) The 2018 National Seismic Hazard Assessment for Australia: Model input files. Record 2018/032. Geoscience Australia. [Online]. Available: http://dx.doi.org/10.11636/Record.2018.032

[128] N. Ackerley. Indian Subcontinent PSHA. [Online]. Available: https://github.com/nackerley/indian-subcontinent-psha

[129] Hazard Model for Southeast Asia (2018). Global Earthquake Model Foundation. Maintained by Chan/Ornammatah. [Online]. Available: https://www.globalquakemodel.org/product/southeast-asia-model

[130] Hazard Model for Canada (2015). Global Earthquake Model Foundation. Maintained by NRCan. [Online]. Available: https://www.globalquakemodel.org/product/canada2015-model

[131] Global Earthquake Model Foundation. Maintained by Public Works Department. [Online]. Available: https://www.globalquakemodel.org/product/indonesia-model

[132] Hazard Model for Western Africa. Global Earthquake Model Foundation. [Online]. Available: https://www.globalquakemodel.org/product/western-africa-model

[133] L. Danciu, K. Sesetyan, M. Demircioglu, M. Erdik, and D. Giardini. (2016) Openquake input files of the seismogenic source model of the 2014 earthquake model of the middle east (emme-project).

[134] Hazard Model for Eastern Sub-Saharan Africa (2018). Global Earthquake Model Foundation. [Online]. Available: https://www.globalquakemodel.org/product/sub-saharan-africa-model

[135] Hazard Model for the Caribbean and Central America. Global Earthquake Model Foundation. [Online]. Available: https://www.globalquakemodel.org/product/ccara2018-model

[136] Hazard Model for South Africa. Global Earthquake Model Foundation. Maintained by GeoScience Council. [Online]. Available: https://www.globalquakemodel.org/product/south-africa-model

[137] S. Ullah, K. Abdrakhmatov, A. Sadykova, R. Ibragimov, A. Ishuk, D. Laurentiu, S. Parolai, D. Bindi, M. Wieland, and M. Pittore. (2015) Emca central asia seismic source model. v. 1.1. [Online]. Available: https://doi.org/10.5880/GFZ.EWS.2015.002

[138] Hazard Model for the Philippines (2018). Global Earthquake Model Foundation. Maintained by PHIVOLCS/GEM. [Online]. Available: https://www.globalquakemodel.org/product/philippines-models

[139] Hazard Model for the Arabian Peninsula. Global Earthquake Model Foundation. Maintained by Saudi Geological Survey. [Online]. Available: https://www.globalquakemodel.org/product/arabia-model

[140] Hazard Model for Taiwan. Global Earthquake Model Foundation. Maintained by TEM. [Online]. Available: https://www.globalquakemodel.org/product/taiwan-model

[141] Hazard Model for Northern Africa (2018). Global Earthquake Model Foundation. [Online]. Available: https://www.globalquakemodel.org/product/northern-africa-model

[142] Papua New Guinea Seismic Hazard Assessment. Geoscience Australia. [Online]. Available: https://github.com/GeoscienceAustralia/PNGSHA

[143] M. Stirling, G. McVerry, M. Gerstenberger, N. Litchfield, R. Van Dissen, K. Berryman, P. Barnes, L. Wallace, P. Villamor, R. Langridge, G. Lamarche, S. Nodder, M. E. Reyners, B. Bradley, D. A. Rhoades, W. D. Smith, A. Nicol, J. Pettinga, K. J. Clark, and K. Jacobs, "National seismic hazard model for new zealand: 2010 update," *Bulletin of the Seismological Society of America*, vol. 102, no. 4, pp. 1514–1542, 2012.

[144] Earthquake Hazards 201 - Technical Q&A. USGS. Accessed: 2021-07-20. [Online]. Available: https://www.usgs.gov/natural-hazards/earthquake-hazards/science/earthquake-hazards-201-technical-qa

[145] E. H. Field, T. H. Jordan, and C. A. Cornell, "Opensha: A developing community-modeling environment for seismic hazard analysis," *Seismological Research Letters*, vol. 74, no. 4, pp. 406–419, 2003.

[146] *The OpenQuake-engine User Manual. Global Earthquake Model (GEM) OpenQuake Manual for Engine version 3.11.2*, GEM, 2021.

[147] K. S. Rukstales and M. D. Petersen. Data Release for 2018 Update of the U.S. National Seismic Hazard Model: U.S. Geological Survey data release. USGS. [Online]. Available: https://doi.org/10.5066/P9WT5OVB

[148] M. Caprio, B. Tarigan, C. B. Worden, S. Wiemer, and D. J. Wald, "Ground motion to intensity conversion equations (gmices): A global relationship and evaluation of regional dependency," *Bulletin of the Seismological Society of America*, vol. 105, no. 3, pp. 1476–1490, 2015.

[149] C. Worden, M. Gerstenberger, D. Rhoades, and D. Wald, "Probabilistic relationships between ground-motion parameters and modified mercalli intensity in california," *Bulletin of the Seismological Society of America*, vol. 102, no. 1, pp. 204–221, 2012.

[150] M. Pagani, J. Garcia-Pelaez, R. Gee, K. Johnson, V. Poggi, V. Silva, M. Simionato, R. Styron, D. Viganò, L. Danciu *et al.*, "The 2018 version of the global earthquake model: hazard component," *Earthquake Spectra*, vol. 36, no. 1_suppl, pp. 226–251, 2020.

[151] T. I. Allen, S. Halchuk, J. Adams, and G. A. Weatherill, "Forensic psha: Benchmarking canada's fifth generation seismic hazard model using the openquake-engine," *Earthquake Spectra*, vol. 36, no. 1_suppl, pp. 91–111, 2020. [Online]. Available: https://doi.org/10.1177/8755293019900779

[152] *ISO 3166*, International Organization for Standardization Std. [Online]. Available: https://www.iso.org/iso-3166-country-codes.html

[153] H. Fujiwara, S. Kawai, N. M. Shin Aoi, S. Senna, N. Kudo, M. Ooi, K. X. Hao, K. Wakamatsu, Y. Ishikawa, T. Okumura, T. Ishii, S. Matsushima, Y. Hayakawa, N. Toyama, and A. Narita, "Technical reports on national seismic hazard maps for japan," NIED, Tech. Rep. 336, November 2009.

[154] "Standard grid square and grid square code used for the statistics (announcement no. 143 by the administrative management agency on july, 12, 1973)," the Administrative Management Agency, July 1973. [Online]. Available: www.stat.go.jp/english/data/mesh/02.html

[155] S. Liew and K. Lu, "A framework for characterizing disaster-based network survivability," *IEEE Journal on Selected Areas in Communications*, vol. 12, no. 1, pp. 52–58, 1994.

[156] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The internet topology zoo," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 9, pp. 1765–1775, 2011.

[157] J. Oostenbrink and F. Kuipers, "The risk of successive disasters: A blow-by-blow network vulnerability analysis," in *2019 IFIP Networking Conference (IFIP Networking)*, 2019, pp. 1–9.

[158] P. J. Klotzbach and M. M. Bell. (2017, November) Summary of 2017 Atlantic tropical cyclone activity and verification of authors' seasonal and two-week forecasts. Department of Atmospheric Science - Colorado State University. [Online]. Available: http://tropical.colostate.edu/media/sites/111/2017/11/2017-11.pdf

[159] E. S. Blake and D. A. Zelinsky. (2018, May) Tropical cyclone report Hurricane Harvey. National Hurricane Center. [Online]. Available: https://www.nhc.noaa.gov/data/tcr/AL092017_Harvey.pdf

[160] J. P. Cangialosi, A. S. Latto, and R. Berg. (2018, June) Tropical cyclone report Hurricane Irma. National Hurricane Center. [Online]. Available: https://www.nhc.noaa.gov/data/tcr/AL112017_Irma.pdf

[161] (2017) Costliest U.S. tropical cyclones tables updated. NHC. [Online]. Available: https://www.nhc.noaa.gov/news/UpdatedCostliest.pdf

[162] (2018) U.S. billion-dollar weather and climate disasters. NOAA National Centers for Environmental Information (NCEI). [Online]. Available: https://www.ncdc.noaa.gov/billions/

[163] A. Witze. (2017, September) Pair of deadly Mexico quakes puzzles scientists. Nature News. [Online]. Available: https://www.nature.com/news/pair-of-deadly-mexico-quakes-puzzles-scientists-1.22650

[164] D. Agren, N. Lakhani, R. Carroll, and S. Jones. (2017, September) At least 225 dead after powerful earthquake hits central Mexico. The Guardian. [Online]. Available: https://www.theguardian.com/world/2017/sep/19/mexico-city-earthquake-anniversary-1985

[165] "Weather, climate & catastrophe insight," Aon Benfield, 2017. [Online]. Available: http://thoughtleadership.aonbenfield.com/Documents/20180124-ab-if-annual-report-weather-climate-2017.pdf

[166] J. W. Baker, "An introduction to probabilistic seismic hazard analysis," *Report for the US Nuclear Regulatory Commission*, 2008.

[167] R. W. Katz, "Stochastic modeling of hurricane damage," *Journal of Applied Meteorology*, vol. 41, no. 7, pp. 754 – 762, 2002. [Online]. Available: https://journals.ametsoc.org/view/journals/apme/41/7/1520-0450_2002_041_0754_smohd_2.0.co_2.xml

[168] K. R. Knapp, M. C. Kruk, D. H. Levinson, H. J. Diamond, and C. J. Neumann, "The international best track archive for climate stewardship (IBTrACS)," *Bulletin of the American Meteorological Society*, vol. 91, no. 3, pp. 363–376, 2010. [Online]. Available: https://doi.org/10.1175/2009BAMS2755.1

[169] (2017, June) Glossary of NHC terms. [Online]. Available: www.nhc.noaa.gov/aboutgloss.shtml

[170] K. M. Sullivan and J. Cole Smith, "Exact algorithms for solving a euclidean maximum flow network interdiction problem," *Networks*, vol. 64, no. 2, pp. 109–124, 2014. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/net.21561

[171] A. Sen, S. Murthy, and S. Banerjee, "Region-based connectivity - a new paradigm for design of fault-tolerant networks," in *2009 International Conference on High Performance Switching and Routing*, 2009, pp. 1–7.

[172] J. Oostenbrink and F. A. Kuipers, "A moment of weakness: Protecting against targeted attacks following a natural disaster," *SIGMETRICS Perform. Eval. Rev.*, vol. 47, no. 4, p. 12–15, Apr. 2020. [Online]. Available: https://doi.org/10.1145/3397776.3397780

[173] (2016, March) NSA chief worries about cyber attack on US infrastructure. [Online]. Available: https://www.securityweek.com/nsa-chief-worries-about-cyber-attack-us-infrastructure

[174] R. Cohen, K. Erez, D. ben Avraham, and S. Havlin, "Breakdown of the internet under intentional attack," *Phys. Rev. Lett.*, vol. 86, pp. 3682–3685, Apr 2001. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.86.3682

[175] C. M. Schneider, A. A. Moreira, J. S. Andrade, S. Havlin, and H. J. Herrmann, "Mitigation of malicious attacks on networks," *Proceedings of the National Academy of Sciences*, vol. 108, no. 10, pp. 3838–3841, 2011. [Online]. Available: https://www.pnas.org/content/108/10/3838

[176] S. Iyer, T. Killingback, B. Sundaram, and Z. Wang, "Attack robustness and centrality of complex networks," *PLOS ONE*, vol. 8, no. 4, pp. 1–17, 04 2013. [Online]. Available: https://doi.org/10.1371/journal.pone.0059613

[177] L. C. Freeman, "A set of measures of centrality based on betweenness," *Sociometry*, vol. 40, no. 1, pp. 35–41, 1977. [Online]. Available: http://www.jstor.org/stable/3033543

[178] J. Oostenbrink, F. A. Kuipers, P. E. Heegaard, and B. E. Helvik, "Evaluating local disaster recovery strategies," *SIGMETRICS Perform. Eval. Rev.*, vol. 46, no. 2, p. 62–66, Jan. 2019. [Online]. Available: https://doi.org/10.1145/3305218.3305241

[179] T. Sakano, Z. M. Fadlullah, T. Ngo, H. Nishiyama, M. Nakazawa, F. Adachi, N. Kato, A. Takahara, T. Kumagai, H. Kasahara, and S. Kurihara, "Disaster-resilient networking: a new vision based on movable and deployable resource units," *IEEE Network*, vol. 27, no. 4, pp. 40–46, 2013.

[180] J. Wang, C. Qiao, and H. Yu, "On progressive network recovery after a major disruption," in *2011 Proceedings IEEE INFOCOM*, 2011, pp. 1925–1933.

[181] "TIGER/Line with selected demographic and economic data," https://www.census.gov/geographies/mapping-files/time-series/geo/tiger-data.2010.html.

[182] (2005) Tropical storm Katrina discussion number 7. NHC. [Online]. Available: https://www.nhc.noaa.gov/archive/2005/dis/al122005.discus.007.shtml?

[183] K. Miranda, A. Molinaro, and T. Razafindralambo, "A survey on rapidly deployable solutions for post-disaster networks," *IEEE Communications Magazine*, vol. 54, no. 4, pp. 117–123, 2016.

[184] K. Genda and S. Kamamura, "Multi-stage network recovery considering traffic demand after a large-scale failure," in *2016 IEEE International Conference on Communications (ICC)*, 2016, pp. 1–6.

[185] K. A. Sabeh, M. Tornatore, and F. Dikbiyik, "Progressive network recovery in optical core networks," in *2015 7th International Workshop on Reliable Networks Design and Modeling (RNDM)*, 2015, pp. 106–111.

[186] M. Pourvali, C. Cavdar, K. Shaban, J. Crichigno, and N. Ghani, "Post-failure repair for cloud-based infrastructure services after disasters," *Computer Communications*, vol. 111, pp. 29–40, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0140366417307831

[187] A. Mazumder, C. Zhou, A. Das, and A. Sen, "Progressive recovery from failure in multi-layered interdependent network using a new model of interdependency," in *International Conference on Critical Information Infrastructures Security*, 2014, pp. 368–380. [Online]. Available: https://doi.org/10.1007/978-3-319-31664-2_38

[188] S. Ferdousi, M. Tornatore, F. Dikbiyik, C. U. Martel, S. Xu, Y. Hirota, Y. Awaji, and B. Mukherjee, "Joint progressive network and datacenter recovery after large-scale disasters," *IEEE Transactions on Network and Service Management*, vol. 17, no. 3, pp. 1501–1514, 2020.

[189] S. Ciavarella, N. Bartolini, H. Khamfroush, and T. La Porta, "Progressive damage assessment and network recovery after massive failures," in *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*, 2017, pp. 1–9.

[190] G. Ishigaki, S. Devic, R. Gour, and J. P. Jue, "Deeppr: Progressive recovery for interdependent vnfs with deep reinforcement learning," *IEEE Journal on Selected Areas in Communications*, vol. 38, no. 10, pp. 2386–2399, 2020.

[191] C. Ma, C. Colman-Meixner, M. Tornatore, Y. Zhao, J. Zhang, and B. Mukherjee, "Multiple traveling repairmen problem with virtual networks for post-disaster resilience," in *2016 IEEE International Conference on Communications (ICC)*, 2016, pp. 1–6.

[192] N. Bartolini, S. Ciavarella, T. F. La Porta, and S. Silvestri, "On critical service recovery after massive network failures," *IEEE/ACM Transactions on Networking*, vol. 25, no. 4, pp. 2235–2249, 2017.

[193] D. Zad Tootaghaj, N. Bartolini, H. Khamfroush, and T. La Porta, "On progressive network recovery from massive failures under uncertainty," *IEEE Transactions on Network and Service Management*, vol. 16, no. 1, pp. 113–126, 2019.

[194] S. Xu, N. Yoshikane, M. Shiraiwa, T. Tsuritani, Y. Awaji, and N. Wada, "Multicarrier-collaboration-based emergency packet transport network construction in disaster recovery," in *2018 10th International Workshop on Resilient Networks Design and Modeling (RNDM)*, 2018, pp. 1–7.

[195] S. Xu, N. Yoshikane, M. Shiraiwa, T. Tsuritani, H. Harai, Y. Awaji, and N. Wada, "Multi-carrier interconnection-based emergency packet transport network planning in disaster recovery," in *DRCN 2017 - Design of Reliable Communication Networks; 13th International Conference*, 2017, pp. 1–8.

[196] S. Xu, N. Yoshikane, M. Shiraiwa, T. Tsuritani, X. Zhang, Y. Awaji, and N. Wada, "A novel carrier-cooperation scheme with an incentive to offer emergency lightpath support during disaster recovery," *Photonic Network Communications*, vol. 40, pp. 175–193, 2020. [Online]. Available: https://doi.org/10.1007/s11107-020-00898-5

[197] J. Oostenbrink and F. Kuipers, "Going the extra mile with disaster-aware network augmentation," in *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*, 2021, pp. 1–10.

[198] W. Qiu. (2011) Submarine cables cut by taiwan earthquake and typhoon morakot. Submarine Cable Networks. [Online]. Available: www.submarinenetworks.com/en/news/cables-cut-by-taiwan-earthquake-and-typhoon-morakot

[199] R. M. Karp, "Reducibility among combinatorial problems," in *Complexity of computer computations*. Springer, 1972, pp. 85–103.

[200] K. A. Dowsland and J. M. Thompson, *Simulated Annealing*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 1623–1655. [Online]. Available: https://doi.org/10.1007/978-3-540-92910-9_49

[201] T. Gomes, J. Tapolcai, C. Esposito, D. Hutchison, F. Kuipers, J. Rak, A. de Sousa, A. Iossifides, R. Travanca, J. André, L. Jorge, L. Martins, P. O. Ugalde, A. Pašić, D. Pezaros, S. Jouet, S. Secci, and M. Tornatore, "A survey of strategies for communication networks to protect against large-scale natural disasters," in *2016 8th International Workshop on Resilient Networks Design and Modeling (RNDM)*, 2016, pp. 11–22.

[202] C. Cao, Z. Wang, M. Zukerman, J. H. Manton, A. Bensoussan, and Y. Wang, "Optimal cable laying across an earthquake fault line considering elliptical failures," *IEEE Transactions on Reliability*, vol. 65, no. 3, pp. 1536–1550, 2016.

[203] M. Zhao, T. W. S. Chow, P. Tang, Z. Wang, J. Guo, and M. Zukerman, "Route selection for cabling considering cost minimization and earthquake survivability via a semi-supervised probabilistic model," *IEEE Transactions on Industrial Informatics*, vol. 13, no. 2, pp. 502–511, 2017.

[204] Z. Wang, Q. Wang, M. Zukerman, J. Guo, Y. Wang, G. Wang, J. Yang, and B. Moran, "Multiobjective path optimization for critical infrastructure links with consideration to seismic resilience," *Computer-Aided Civil and*

Infrastructure Engineering, vol. 32, no. 10, pp. 836–855, 2017. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/mice.12287

[205] Z. Wang, Q. Wang, B. Moran, and M. Zukerman, "Application of the fast marching method for path planning of long-haul optical fiber cables with shielding," *IEEE Access*, vol. 6, pp. 41 367–41 378, 2018.

[206] Q. Wang, J. Guo, Z. Wang, E. Tahchi, X. Wang, B. Moran, and M. Zukerman, "Cost-effective path planning for submarine cable network extension," *IEEE Access*, vol. 7, pp. 61 883–61 895, 2019.

[207] Z. Wang, Q. Wang, B. Moran, and M. Zukerman, "Terrain constrained path planning for long-haul cables," *Opt. Express*, vol. 27, no. 6, pp. 8221–8235, Mar 2019. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-27-6-8221

[208] Z. Wang, Q. Wang, B. Moran, and M. Zukerman, "Optimal submarine cable path planning and trunk-and-branch tree network topology design," *IEEE/ACM Transactions on Networking*, pp. 1–11, 2020.

[209] T. Tsubaki, M. Ishizuka, and S. Yasukawa, "A new algorithm of route design against large-scale disasters," in *NOMS 2018 - 2018 IEEE/IFIP Network Operations and Management Symposium*, 2018, pp. 1–5.

[210] A. Alashaikh, T. Gomes, and D. Tipper, "The spine concept for improving network availability," *Computer Networks*, vol. 82, pp. 4 – 19, 2015, robust and Fault-Tolerant Communication Networks. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1389128615000687

[211] L. Garrote, L. Martins, U. J. Nunes, and M. Zachariasen, "Weighted euclidean steiner trees for disaster-aware network design," in *2019 15th International Conference on the Design of Reliable Communication Networks (DRCN)*, 2019, pp. 138–145.

[212] W. Kellerer, P. Kalmbach, A. Blenk, A. Basta, M. Reisslein, and S. Schmid, "Adaptable and data-driven softwarized networks: Review, opportunities, and challenges," *Proceedings of the IEEE*, vol. 107, no. 4, pp. 711–731, 2019.

[213] N. Feamster and J. Rexford, "Why (and how) networks should run themselves," *arXiv preprint arXiv:1710.11583*, 2017.