On the Robustness of Network Infrastructures to Disasters and Physical Attacks

by

Sebastian James Neumayer

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2013

Author ....
Department of Electrical Engineering and Computer Science
October 9th, 2012

Certified by ....
Eytan Modiano
Professor
Thesis Supervisor

Accepted by ....
Leslie A. Kolodziejski
Chair, Department Committee on Graduate Studies

On the Robustness of Network Infrastructures to Disasters and Physical Attacks

by

Sebastian James Neumayer

Submitted to the Department of Electrical Engineering and Computer Science
on October 9th, 2012, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Networks are vulnerable to natural disasters, such as earthquakes or floods, as well as to physical attacks, such as an Electromagnetic Pulse (EMP) attack. Such real-world events happen in specific geographical locations and disrupt specific parts of the network. Therefore, the geographical layout of the network determines the impact of such events on the network's connectivity. We focus on network analysis and design under a geographic failure model of (geographical) networks to such disasters.

Initially, we aim to identify the most vulnerable parts of data networks to attack. That is, the locations of a disaster that would have the maximum disruptive effect on a network in terms of capacity and connectivity. We consider graph models in which nodes and links are geographically located on a plane, and model the disaster event as a line segment or circular disk. We develop polynomial time algorithms for finding the worst possible cut in this setting. Then, we obtain numerical results for a specific backbone network, thereby demonstrating the applicability of our algorithms to real-world networks.

We also develop tools to calculate network metrics after a 'random' geographic disaster. The random location of the disaster allows us to model situations where the physical failures are not targeted attacks. In particular, we consider disasters that take the form of a 'random' circular disk or line in a plane. Using results from geometric probability, we are able to calculate some network performance metrics to such a disaster in polynomial time. In particular, we can evaluate average two-terminal reliability in polynomial time under these 'random' cuts. This is in contrast to the case of independent link failures for which there exists no known polynomial time algorithm to calculate this reliability metric. We present some numerical results to show the significance of geometry on the survivability of the network. This motivates the formulation of several network design problems in the context of randomly located disasters.

We also study some min-cut and max-flow problems in a geographical setting. Specifically, we consider the problem of finding the minimum number of failures, modeled as circular disks, to disconnect a pair of nodes and the maximum number

3

of failure disjoint paths between a pair of nodes. This model applies to the scenario where an adversary is attacking the network *multiple* times with intention to reduce its connectivity. We present a polynomial time algorithm to solve the geographic min-cut problem and develop an ILP formulation, an exact algorithm, and a heuristic algorithm for the geographic max-flow problem.

Finally, we study the reliability of power transmission networks under regional disasters. Initially, we quantify the effect of large-scale non-targeted disasters and their resulting cascade effects on power networks. We then model the dependence of data networks on the power systems and consider network reliability in this dependent network setting. Our novel approach provides a promising new direction for modeling and designing networks to lessen the effects of geographical disasters or attacks.

Thesis Supervisor: Eytan Modiano
Title: Professor

Acknowledgments

Almost all of my acknowledgement goes to my advisor, Professor Eytan Modiano. His support and encouragement throughout my years at MIT have been essential to my research as well as to my personal development.

The remaining measure of my acknowledgement goes to the following awesome people:

Swimming: John '1.O' Kelleher, Kelly Benedict, **Santiago Lima**, Jen-Jen Lao, Todd Herman, Josh Gonzalez, Mike Baier, Zack Cordero, John '3.O' Wang, and **Coach Bill Paine**

Board Game Crew: Chris Evans, Andrew Rader, Joe Sikora, Prashant Luitel, **Pouya Kheradpour**, Matt Rasmussen, and Jessica

Childhood Friends: Bjorn Johnson, Mark Kittisopikul, Justin Daun, and Steve Shrygler

UIUC: Maneesh Shanbhag, Vinh Lam, Karl Schmidt, Mike Molinaro, Dave Reyna, Mike Leavitt, and Jon Michaels

Random Hall: Nina and all the great people who live there

Disney World: Lucky and Coach Cathy

Family: Mom, Pop, Rachel, Rebecca, and Skippy Neumayer (and Grandma too)

Other Notable: Srujan Linga, Sungwon Chung, Greg Kuperman, Greg Monigold, John Nikolai, Lauren Hendrix, CNRG labmates, Coaches Bonebrake and Keller, Profs. Zussman, Efrat, Meyn, Wyatt, Staelin, Jaillet, and Tsitsiklis

Contents

Chapter 1

Introduction

The global communications infrastructure is primarily based on fiber-optic networks, and as such has physical vulnerabilities. Similarly, power transmission networks, critical to the operation of data networks, use high-voltage power lines and are vulnerable to physical failures. Fiber links and power lines can be destroyed by anything from Electromagnetic Pulse (EMP) attacks [40, 60] to natural disasters such as hurricanes or earthquakes [38, 42]. Such real-world disasters happen in specific geographic locations, and therefore, the geographical layout of the network affects their impact. For example, an Electromagnetic Pulse (EMP) is a large burst of electromagnetic energy that can disable electronics over a large geographic region [70]. Hence, such an attack over a city which is a telecommunications hub would have a disastrous impact on the telecommunications infrastructure. In this thesis we develop the necessary theory to evaluate network performance metrics under a geographic failure model. This allows us to begin developing some network design tools that can mitigate the effects of regional disasters.

There are several works on the topology of the Internet as a random graph [12] and on the effect of link failures in these graphs [27, 48] (for more details see Section 1.2). However, most of these works are motivated by failures of routers due to logical attacks (e.g., viruses and worms), and thereby, focus on the logical Internet topology. There have also been some attempts to model the Internet using geographical notions [41, 72] (see Fig. 1-1 for a fiber map). Yet, these works do not consider the effect of failures

Figure 1-1: The fiber backbone operated by a major U.S. network provider [46].

that are geographically correlated.

Since disasters affect a specific geographical area, they will result in failures of neighboring network components. Therefore, one has to consider the effect of disasters on the physical layer rather than on the network layer (e.g., the effect on the fibers rather than on the logical links). Again, these fibers are subject to regional failures resulting from events such as earthquakes, floods, and even an EMP attack; as these may lead to failure of the electrical circuits (e.g., amplifiers) that are needed to operate the fiber plant [70]. Similarly, a geomagnetic storm may damage lines in the power grid and cause power loss over a large geographic region [3].

Our goal is to understand the effect of a regional failure on the bandwidth and connectivity of the Internet as well as the reliability of the power transmission network, and to expose the design tradeoffs related to network survivability under a disaster with regional implications. Such tradeoffs may imply that in certain cases there may be a need to redesign parts of the network while in other cases there is a need to protect electronic components in critical areas (e.g., protecting against EMP attacks by shielding).

In the remainder of this chapter we review the problems considered in this thesis, our contributions, and related work.

1.1 Problem Descriptions and Contributions

We now give an overview of the problems considered and contributions of the thesis. Initially motivated by targeted attacks on data networks, in chapter 2, we consider the location of geographical disasters that have the maximum effect on a network, in terms of capacity and connectivity. That is, we want to identify the worst-case location for a disaster or attack with respect to certain connectivity metrics. Then, we turn our attention to the effects of non-targeted attacks such as natural disasters; in chapter 3 we analyze the effect of 'randomly' located regional failures on a network. Specifically, we introduce methods to calculate relevant network connectivity metrics after such an event. Motivated by the effects of *multiple* disasters, in chapter 4 we consider the geographic min-cut and max-flow problem. In chapter 5 we analyze the effects of large scale disasters on power grids and consider some power-data network dependency problems in the context of network survivability. Finally, in chapter 6 we conclude and discuss possible extensions to this research. In the following subsections we give a more detailed description of each problem and the contributions made.

1.1.1 Targeted Attacks

In the context of disasters that cause the failure of multiple links in a geographic region, chapter 2 focuses on the worst-case location for a disaster to occur. This can model a scenario where an adversary that knows the network topology and geography is attacking the network.

We consider two graph models which serve as an abstraction of the continental/undersea fiber plant. In our graph models, nodes are represented by points on the plane and links are represented by line segments connecting these points. Let a cut denote a geometric shape located on the plane, such as a line segment or disk. We assume that a cut (which represents a regional disaster) affects the electronic components of the network within its particular region. Hence, the fibers that pass through the cut are assumed to be effectively destroyed and removed from the network. See Fig. 1-2 for an example of the effects of a particular disk cut.
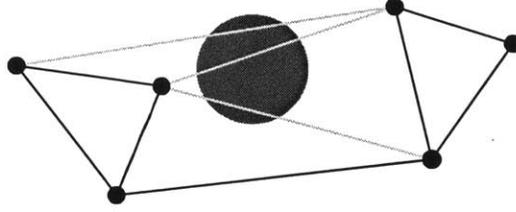
Figure 1-2: The black dots represent node locations in the network and the line segments between these points represent links. The red disk represents a circular cut (which may model the effect of some large disaster). The three grey links represent link failures that occur due to this cut. In chapter 2 we find the worst case location for these types of cuts with respect to certain connectivity metrics.

We first study a bipartite graph model (in the topological and geographical sense). This model is analogous to the east and west coasts of the U.S., where nodes on the left and right sides of the graph represent west and east coast cities (respectively) and the cities within the continent are ignored. Similarly, it can represent transatlantic or transpacific cables. Since vertical line segment cuts are simpler to analyze, we focus first on such cuts.

However, the bipartite model does not consider the impact on nodes located within the continent; nor does it consider the impact of a disaster that is not simply a vertical cut. We relax the *bipartite* graph and *vertical* cut assumptions by considering a general model where nodes can be (almost) arbitrarily located on the plane. Under this model, we consider two problems. In the first one, disasters are modeled as line segment cuts (not necessarily vertical) in the network graph. In the second one, disasters are modeled as circular disks in which the links and nodes are affected. These general problems can be used to study the impact of disasters such as EMP attacks (disks) and natural disasters (line segments) more realistically.

We consider various performance metrics for the effect of a cut. We consider the following: (i) the capacity of the removed links, (ii) the fraction of pairs of nodes that remain connected (termed average two-terminal reliability or $ATTR$), (iii) the maximum possible flow between a given source-destination pair, and (iv) the average maximum flow between pairs of nodes. We show that although there are infinite number of possible cut locations, only a polynomial number of candidate locations have to be considered in order to identify a worst-case cut for the aforementioned performance metrics. Thus, we are able to show that the location of a worst-case cut

Figure 1-3: Line segments cuts minimizing the fraction of pairs of nodes that remain connected. The red cut minimizes this metric and the black segments nearly minimize it.

can be found by polynomial time algorithms.

We then present numerical results which demonstrate the use of these algorithms. We identify the locations of the worst-case line segment and circular cuts in the network presented in Fig.1-1.[1] In particular, we illustrate the locations of cuts that optimize the different performance metrics described above. See Fig. 1-3 for an example.

The main contributions of chapter 2 are the formulation of a new problem (termed as the *geographical network inhibition* problem), the design of algorithms for its solution, and the demonstration of the obtained numerical solutions on an example infrastructure. To the best of our knowledge, our work is among the first to study this problem.

1.1.2 Non-targeted Attacks

Motivated by the effects of non-targeted failures such as natural disasters (e.g. hurricanes and floods) or collateral damage from an attack, in chapter 3 we consider a failure model where a disk or line is 'randomly' placed in the plane. Our goal is to calculate the expected value of relevant network connectivity metrics after such an event. As in the geographical network inhibition problem above, we assume our network nodes and links are represented by points and line segments in the plane.

---

[1]We present results only for one major operator. The same methodologies can be used in order to obtain results for all other major operators.
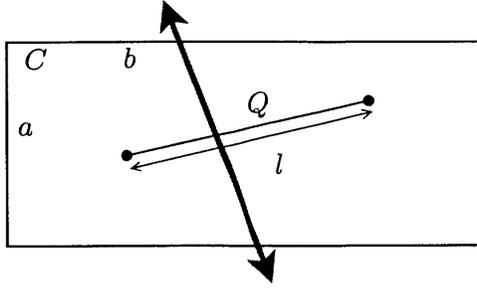
Figure 1-4: In the example above we consider the probability that a 'random' line that intersects the rectangle C will also intersect the line segment link Q. In chapter 3 using tools and measures from geometric probability we will show this probability is given by the ratio of the perimeters of the line segment and rectactangle, that is $\frac{l}{a+b}$.

We assume that any links which are intersected by the randomly placed line or disk are removed from the network.

In order to obtain probabilities of relevant failure events we use geometric probability to assign a measure to sets of lines and disks that intersect some set of line segments (e.g. a set of segments that disconnects the network). See Fig. 1-4 for an example. Using these tools, we are able to calculate certain network performance metrics to a randomly located geographic disaster in polynomial time. To the best of our knowledge this is the first attempt to apply geometric probability techniques to network survivability ([47] applies similar techniques to detection in sensor networks). In particular, we can calculate average two-terminal reliability of a network in polynomial time with respect to a randomly located line or disk. This is a significant contribution because calculating this metric assuming independent link failures in known to be NP-hard [10]. We then present some numerical results to show the significance of geometry on the survivability of the network (see Fig. 1-5) and consider a few network design examples.

1.1.3 Geographic Min-Cut and Max-Flow

In chapter 4 we consider the problem of finding the minimum number of failures, modeled as circular disks, to disconnect two nodes and the maximum number of failure disjoint paths between two nodes. This models the scenario where an adversary is attacking the network *multiple* times (with geographic scale attacks) with intention to reduce its capacity or connectivity. These problems may also be useful in the context
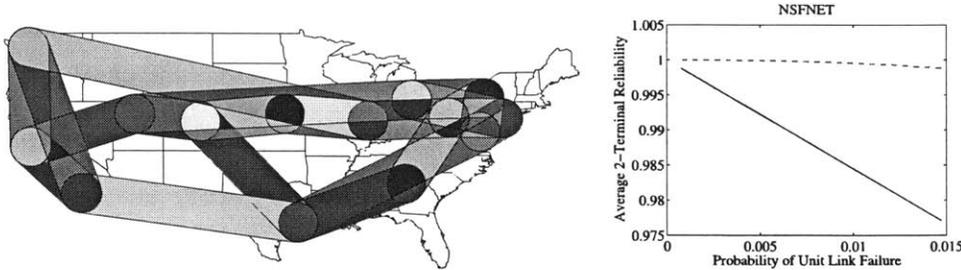
Figure 1-5: The colored areas in the figure on the left represent relevant regions with respect to NSFNET [51] and a 'randomly' located circular cut of a particular radius. In the right figure, the solid line shows the fraction of pairs of nodes that remain connected ($ATTR$) versus the probability a unit (latitude/longitude) of fiber is cut by a random disk. The dashed line shows $ATTR$ assuming links fail *independently* such that links fail with the same probability as in the random disk-cut case. The difference in these curves shows the significance of geometry on the survivability of the network.
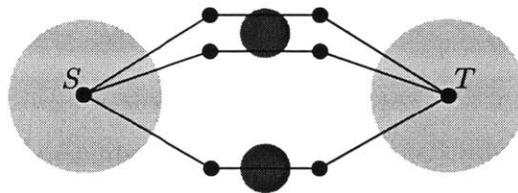
Figure 1-6: The light gray area (yellow area in online color version) above represents the protected zone that no circular failure may be centered. The gray disks (red disks in online color version) represent disasters that remove links (of unit capacity) they intersect. Two disasters are required to disconnect the two nodes S and T (shown above), so the geographic min-cut is two. Also, since the top pair of paths can be intersected by the same failure, the geographic max-flow is two; two failure disjoint paths are given by the topmost and bottommost path. In contrast, the standard min-cut and max-flow is three.

of path protection algorithms to ensure at least some of the primary and backup paths survive a large scale failure.

We first consider a geographical variant of the min-cut problem. Given a set of points on the plane, each of which represents a node, and (non-overlapping) line segments between these points representing links, what is the minimum number of circular failures such that two nodes, S and T, are disconnected from each other. If applied to the national fiber plant, the solution to this problem is the number of failures required to disconnect two cities. If we do not restrict the locations of potential failure sites, the geographic min-cut will be at most one because nodes S or T can trivially be eliminated with a single failure. In order to make the problem more interesting and realistic we can restrict potential failure locations (see Fig. 1-6). This can represent fiber that has been hardened against EMP attacks or a well defended city. We extend our arguments from chapter 2 to show we only need to consider a polynomial number of possible failure sites, thus reducing the geographic

17
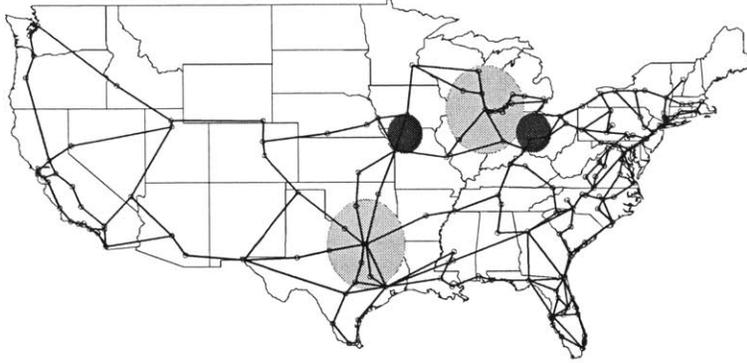
Figure 1-7: A solution to the geographic min-cut problem on a particular infrastructure. The disaster radius is about 78 miles and the protection radius around Chicago and Dallas is about 180 miles. The gray disks (red disks in the online color version) represent the hole locations and the light gray disks (yellow disks in the online color version) represent the protected zones. Only two disasters, located at 'choke' points to the east and west of Chicago, are required to disconnect these cities. We note that the standard (non-geographical) min-cut solution is 4, but because the disks remove multiple links at a time there are only 2 disasters in a geographical min-cut solution.

min-cut to a discrete problem. Then applying the methods in [18], we show how to find a solution in polynomial time. We obtain numerical results for a specific backbone network, thereby demonstrating the applicability of our min-cut algorithm to a real-world network. See Fig. 1-7 for an example.

Next, in the context of geographic failures and path-protection algorithms we study the geographic max-flow problem; what is the maximum number of paths between nodes S and T such that no two paths can be intersected by the same failure. The solution to this problem gives us the maximum number of paths that are geographically disjoint with respect to disasters of a particular radius. In other words, we are interested in finding the maximum number of backup paths between a pair of nodes such that a disaster intersecting one of the paths does not affect the connection of the other paths (see Fig. 1-6). The solution to this problem gives us failure disjoint backup paths. Again, to avoid triviality we restrict the locations of potential failure sites so that nodes S or T cannot simply be eliminated with a single failure. We then develop an integer linear program (ILP) formulation, an exact algorithm, and a heuristic algorithm for this geographic max-flow problem. See Fig. 1-8 for an illustration of the heuristic on a real-world network.

In the final part of chapter 4, we explore the analogue to the min-cut max-flow theorem in the geographic setting. In particular, we show that the cardinality of

18
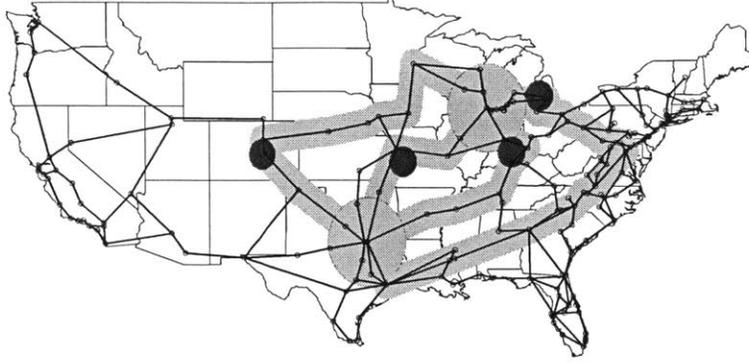
Figure 1-8: The four gray disks (red disks in the online color version) represent the hole locations in a geographic min-cut and the light gray disks (yellow disks in the online color version) represent the protected zones. The disaster radius is about 60 miles and the protection radius around Chicago and Dallas is about 180 miles. The four light gray 'worms' (teal 'worms' in the online color version) correspond to hole disjoint paths found using the heuristic algorithm developed for the geographic max-flow problem. We note that since the cardinality of a max-flow solution must be less than a min-cut solution, our developed heuristic has found an optimal solution to the geographic max-flow problem for this particular instance. Also, comparing with Fig. 1-7 we observe that the min-cut increases from 2 to 4 as the radius of disaster decreases from 78 to 60 miles.

the solutions to these geographic min-cut and max-flow problems are not the same. Supported by simulation results, we conjecture this difference is no greater than one, i.e. max-flow $\leq$ min-cut $\leq$ max-flow $+1$.

1.1.4 Applications to Power Networks

Similar to fiber infrastructures, power transmission networks are vulnerable to large-scale natural disasters or attacks, such as hurricanes or geomagnetic storms [3, 22]. Beyond the effects of physical disasters, power networks are also vulnerable to *cascading failures*. Cascading failures occur when an initial failure in the network changes power flows, which must obey physical law constraints, such that additional lines overload and fail. This in turn causes the power flows to change again; this process will continue until some stability is reached. A well known example of a cascading failure is the 2003 blackout [6]. In the following we describe the two failure models presented in chapter 5. The first model considers power networks with with respect to geographic disasters and cascading failures. The second model builds on the first; we describe a dependency between power and data networks and consider the connectivity of data networks in this context.

Motivated by the effects of natural disasters and cascading failures, in chapter 5

CDF of Yield on HVIET network with radius ≈ 50km, Average Yield = 0.78336
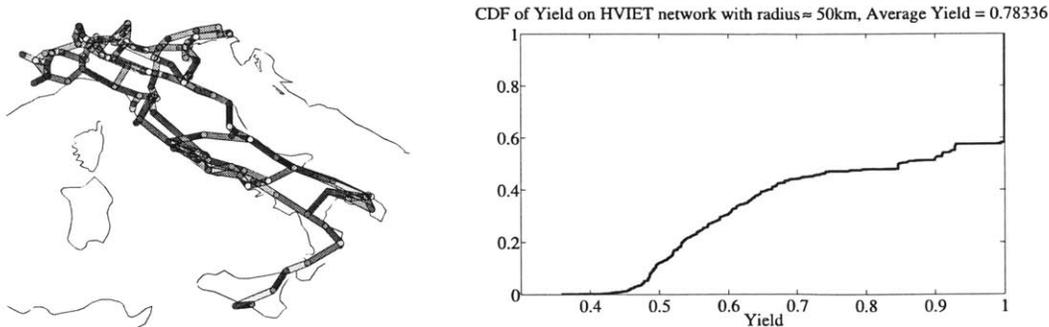
Figure 1-9: In the left figure every shaded region represents a set of disk centers whose radius is about 8 kilometers and only intersects a particular set of links in the Italian high-voltage electrical transmission network (HVIET) [63, 64]. In our model, the area of each of these regions is proportional to the probability a randomly located disk will remove a particular set of power lines. The right figure shows the CDF of the yield (total fraction of demand satisfied) on the HVIET network [63, 64] under our two-stage failure model (a randomly located disk followed by cascading failures). Note that there is a significant probability the yield is 1; this is mainly caused by disks centered within the region of interest but do not intersect the network.

we consider a two-stage failure model for power networks. The first stage removes power lines that intersect a randomly located disk. The second stage then calculates the cascading failure that occurs due to the removal of the initial links. By using the tools developed in chapter 3 and using the cascading failure model developed in [16], we are able to calculate the effect of this type of failure in power networks. To the best of our knowledge, [14] is the only other work to look at the effect of geographically correlated failures on power networks. See Fig. 1-9 for an example of a numerical result.

Then motivated by the effects of power loss on data networks [30], in the final part of chapter 5 we consider the survivability of data networks with respect to power networks. We assume data nodes rely on the operation of the closest power demand nodes to function. See Fig. 1-10 for an example. We note that network connectivity in this particular example is significantly lower when power network dependency is considered; this implies power network effects have a significant impact on the survivability of real-world data networks.

1.2 Related Work

The issue of network survivability and resilience has been extensively studied in the past (e.g., [15, 35, 49, 74] and references therein). However, most of the previous work
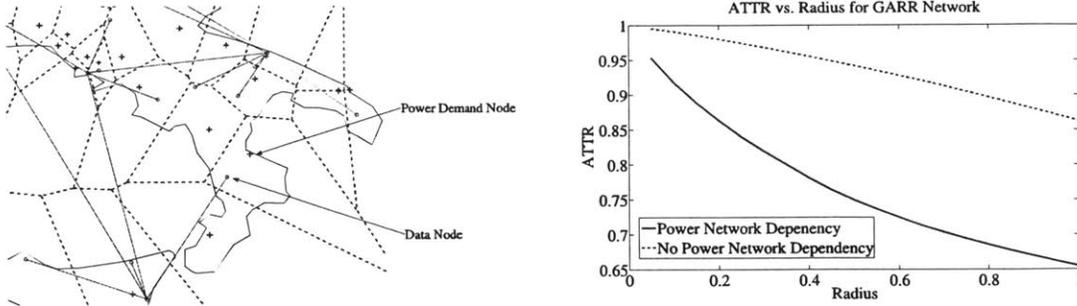
Figure 1-10: In the left figure part of the backbone of the Italian research network (GARR) [63, 64] is shown by solid line segments representing links and circles representing nodes. The dashed segments represent the Voronoi cells based on the locations of power supply nodes, shown by crosses above, in the Italian high-voltage electrical transmission network (HVIET) [63, 64]. Our model assumes that data nodes extract power from the closest power demand node; when a power demand node no longer receives power, data nodes located within its Voronoi cell are assumed to fail as well. The right figure shows $ATTR$ in the GARR network versus the effect of a randomly located disk. The solid line assumes power network dependency effects and the dashed line assumes only the data network itself is affected by the disaster (the same model considered in chapter 3). We note that $ATTR$ is significantly lower when power network dependency is considered; this implies power network effects have a significant impact on the survivability of real-world data networks.

in this area and in particular in the area of physical topology and fiber networks (e.g., [53, 54]) focused on a *small number of fiber failures*. On the contrary, in this thesis we focus on events that cause a large number of failures in a specific geographical region (e.g., [7, 20, 40, 60]). To the best of our knowledge, [36] is among the first papers that considered geographically correlated failures. Yet, it focused on a specific routing solution.

A theoretical problem closely related to some of the problems considered in this thesis is known as the *network inhibition problem* [57]. Under that problem, each edge in the network has a destruction cost, and a fixed budget is given to attack the network. A feasible attack removes a subset of the edges, whose total destruction cost is no greater than the budget. The objective is to find an attack that minimizes the value of a maximum flow in the graph after the attack. Several variants of this problems have been studied in the past (see for example [58] and the review in [26]). However, as mentioned above, the removal of (geographically) neighboring links has not been considered. Perhaps the closest to this concept is the problem formulated in [67] which considers geographical failures in a wireless network setting.

When the logical (i.e., IP) topology is considered, wide-spread failures have been extensively studied [27, 28, 34, 48]. Most of these works consider the topology of the Internet as a random graph [12] and use percolation theory to study the effects of

21

random link and node failures on these graphs and whether the resulting network has a large connected component.

The focus on the logical topology rather than on the physical topology is motivated by failures of routers due to attacks by software viruses and worms. Based on various measurements (e.g., [33]), it has been recently shown that the topology of the Internet is influenced by geographical factors such as population density [11, 41, 72]. These observations motivated the modeling of the Internet as a scale free geographical graph [69, 73].

Geographic min-cut and max-flow problems similar to the ones described in this chapter have received some attention. Recently Sen [67] [66] has proposed the idea of a geographic max-flow and min-cut in a wireless network setting. In [52] the problem of finding the maximum number of geographically disjoint paths with total minimum cost is discussed in the continuous setting where paths may be placed anywhere within a polygonal domain. Additionally, Bienstock has analyzed similar problems to the ones described above [18]. The key difference is that we assume disasters are circular and may be placed almost anywhere on the plane; in [18] disasters may be of various shapes but can only be placed in a finite number of locations.

Power network resilience has been considered in the past [9, 17], however so far only [14] has considered the effects of a targeted geographic failure model. In this work we consider the effect of non-targeted geographic attacks on the power network. Some recent work has modeled the interdependence between data and power networks and demonstrated asymptotic percolation results [23]; however they did not consider power flows or geography in their models. Additionally, [63] considered a geographic dependence model but did not consider failures which were geographically correlated.

The rest of the thesis is organized as follows. In chapter 2 we consider the problem of finding worst case locations for disaster or attack on a network. In chapter 3 we consider randomly located attacks on the network. In chapter 4 we discuss the geographic min-cut and max-flow problem. Then, in chapter 5 we apply our developed tools to assess the survivability of power networks. Finally, in chapter 6 we conclude and propose future research directions.

Chapter 2

Targeted Attacks

Motivated by the effects of large-scale failures on the fiber infrastructure, in this chapter we are interested in the location of geographical disasters that have the maximum effect on a data network, in terms of capacity and connectivity. That is, we want to identify the worst-case location for a disaster or an attack as well as its effect on the network.

2.1 Introduction

Fiber links in backbone data networks have geographic location and as such have physical vulnerabilities. For example, natural disasters (e.g. earthquakes) as well as EMP attacks may cause the failure of several links over a large geographic area [38, 40, 42, 60]. Modeling networks and attacks as geometric objects on the plane, we consider the problem of finding the worst-case location for a disaster to occur. We will now introduce our models in more detail.

The global fiber plant has a complicated structure. For example, Fig. 2-1 presents the fiber backbone operated by a major network provider in the U.S. (node locations are approximate and point-to-point fibers are represented by straight lines). We consider two graph models which serve as an abstraction of the continental/undersea fiber plant. In these models, nodes, links, and cuts are geographically located on a plane. Nodes are represented as points and links are represented as line segments

Figure 2-1: The fiber backbone operated by a major U.S. network provider [46]. Node locations are approximate and point-to-point fibers are represented by straight lines.

between these points. We first study a bipartite graph model (in the topological and geographical sense). That model is analogous to the east and west coasts of the U.S., where nodes on the left and right sides of the graph represent west and east coast cities (respectively) and the cities within the continent are ignored. Similarly, it can represent transatlantic or transpacific cables. Since vertical line segment cuts are simpler to analyze, we focus initially on such cuts and provide some motivating examples.

However, the bipartite model does not consider the impact on nodes located within the continent; nor does it consider the impact of a disaster that is not simply a vertical cut. Therefore, we later relax the *bipartite* graph and *vertical* cut assumptions by considering a general model where nodes can be arbitrarily located on the plane. Under this model, we consider two problems. In the first one, disasters are modeled as line segment cuts (not necessarily vertical) in the network graph. In the second one, disasters are modeled as circular areas in which the links and nodes are affected. These general problems can be used to study the impact of disasters such as EMP attacks (circular disks) and tornadoes (line segments) more realistically.

We assume that a regional disaster affects the electronic components of the network within a certain region. Hence, the fibers that pass through that region are effectively removed due to such a disaster. There are various performance measures for the effect of a cut. We consider the following: (i) the expected capacity of the removed links, (ii) the fraction of pairs of nodes that remain connected, (iii) the

24

maximum possible flow between a given source-destination pair, and (iv) the average maximum flow between pairs of nodes. We show that although there are infinite number of cut locations, only a polynomial number of candidate cuts have to be considered in order to identify a worst-case cut for these performance measures in any of the problems above. Thus, we are able to show that the location of a worst-case cut can be found by polynomial time algorithms. It should be noted that any other quantity that can be calculated in polynomial time may be used as a performance measure. Hence, measures such as concurrent maximum flow and other measures that are derived from multicommodity flow problems may also be used.

Finally, we present numerical results and demonstrate the use of these algorithms. We identify the locations of the worst-case line segment and circular cuts in the network presented in Fig. 2-1.[1] In particular, we illustrate the locations of cuts that optimize the different performance measures described above.

The main contributions of this chapter are the formulation of a new problem (termed as the *geographical network inhibition* problem), the design of algorithms for its solution, and the demonstration of the obtained numerical solutions on a U.S. infrastructure. To the best of our knowledge, we are the first to attempt to study this problem.

This chapter is organized as follows. We briefly discuss related work in Section 1.2. In Section 2.3, we introduce the network models and formulate the geographical network inhibition problems. In Section 2.4, we consider a simple case of the bipartite model and provide numerical examples that provide insight into the location of a worst-case cut. In Section 2.5, we develop a polynomial-time algorithm for finding the worst-case cuts in the bipartite model. In Sections 2.6 and 2.7 we study the general model with line segment and circular cuts. In Section 2.8 we present numerical results. We conclude and discuss future research directions in Section 2.9.

---

[1]We present results only for one major operator. The same methodologies can be used in order to obtain results for all other major operators.

25

2.2 Related Work

The issue of network survivability and resilience has been extensively studied in the past (e.g., [15,35,49,74] and references therein). However, most of the previous work in this area and in particular in the area of physical topology and fiber networks (e.g., [53,54]) focused on a *small number of fiber failures* or on the concept of *Shared Risk Link Group* (SRLG) [39]. On the contrary, in this chapter we focus on events that cause a large number of failures in a specific geographical region (e.g., [7,20,40,60]). To the best of our knowledge, before our work, *geographically correlated failures* have been considered only in a few papers and under very specific assumptions [8,36,71].

The theoretical problem most closely related to the problem we consider is known as the *network inhibition problem* [57]. Under that problem, each edge in the network has a destruction cost, and a fixed budget is given to attack the network. A feasible attack removes a subset of the edges, whose total destruction cost is no greater than the budget. The objective is to find an attack that minimizes the value of a maximum flow in the graph after the attack. A few variants of this problems were studied in the past (e.g., [24,26,58]). However, as mentioned above, the removal of (geographically) neighboring links has been rarely considered [18,66]. One of the first and perhaps the closest to this concept is the problem studied in [67].

When the logical (i.e., IP) topology is considered, wide-spread failures have been extensively studied [27,28,34,48]. Most of these works consider the topology of the Internet as a random graph [12] and use percolation theory to study the effects of random link and node failures on these graphs. The focus on the logical topology rather than on the physical topology is motivated by failures of routers due to attacks by viruses and worms. Based on various measurements (e.g., [33]), it has been recently shown that the topology of the Internet is influenced by geographical concepts [11, 41,72]. These observations motivated the modeling of the Internet as a scale free geographical graph [69,73]. Although these models may prove useful in generating *logical* network topologies, we decided to present numerical results based on *real physical* topologies (i.e., the topology presented in Fig. 2-1).

2.3 Model and Problem Formulation

In this section we present three geographical network inhibition problems. The first problem assumes that the network is bipartite in the topological and geographic sense and that the cuts are vertical line segments. We then present two problems where network links can be in almost arbitrary locations on the plane. In one of the problems, the disasters correspond to line segment cuts in any direction. In the other, the cuts are modeled by arbitrarily placed circular disks on the plane.

2.3.1 Bipartite Model with Vertical Line Segment Cuts

We now define the *geometric* bipartite graph. It has a width of 1 and height (south-to-north) of h_G. The height of a left (west) node i is denoted by l_i. Similarly, the height of a right (east) node j is denoted by r_j. Nodes cannot overlap and must have non-negative height; that is $r_i \neq r_j \geq 0 \ \forall \ i, j$ and $l_i \neq l_j \geq 0 \ \forall \ i, j$. Denote the total number of nodes on the left and right side by N. We denote a link from node i to node j as (i, j) and let (i, j) be represented by a line segment from $[0, l_i]$ to $[1, r_j]$. We define p_{ij} as the probability that link (i, j) exists, and c_{ij} as the capacity of link (i, j) where $c_{ij} \in [0, \infty)$. To avoid considering the trivial case in which there are no links with positive capacity, we assume that there exist some i and j for which $c_{ij} p_{ij} > 0$. We assume that the disaster results in a vertical line segment cut of height h whose lowest point is at point $[x, y]$. We denote this cut by $\text{cut}_h(x, y)$. Such a cut removes all links that intersect it. For clarity, in this chapter we refer to the start and the end of a link as nodes and the start and the end of a cut as endpoints. Fig. 2-2 demonstrates a specific construction of the model and an example of a cut.

There are many ways to define the effect of a cut on the loss of communication capability in a network. We define the performance measures and the worst-case cut as follows.

Definition 1 (Performance Measures). *The performance measures of a cut are (the last 3 are defined as the values after the removal of the intersected links):*

- *TEC* - The total expected capacity of the intersected links.

27
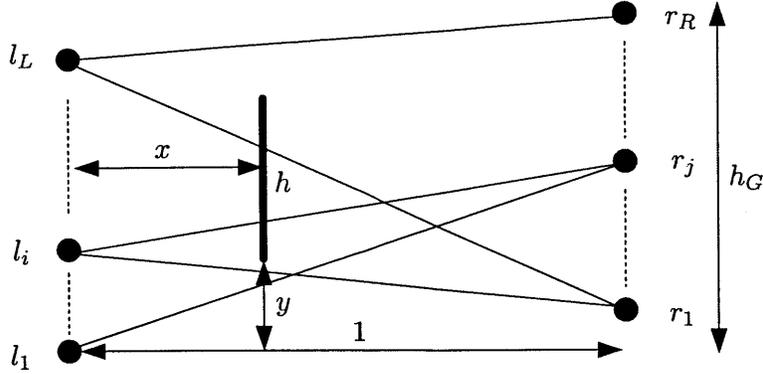
Figure 2-2: A bipartite network and an example of a cut.

- *ATTR* - The fraction of pairs of nodes that remain connected (this is similar to the average two-terminal reliability of the network[2]).

- *MFST* - The maximum flow between a given pair of nodes s and t.

- *AMF* - The average value of maximum flow between all pairs of nodes.

Definition 2 (Worst-Case Cut). *Under a specific performance measure, a worst-case cut, denoted by* $\mathrm{cut}_h(x^*, y^*)$, *is a cut which maximizes/minimizes the value of the performance measure.*[3]

We now demonstrate the formulation of the following optimization problem using the *TEC* performance measure.

Bipartite Geographical Network Inhibition (BGNI) Problem: *Given a bipartite graph, cut height, link probabilities, and capacities, find a worst-case vertical line segment cut under performance measure* TEC.

We define the following $(0, 1)$ variables:

$$z_{ij}(x, y) \quad = \quad \begin{cases} 1 & \text{if } (i, j) \text{ is removed by } \mathrm{cut}_h(x, y) \\ 0 & \text{otherwise} \end{cases}$$

---

[2]The two-terminal reliability between two nodes is the probability they remain connected after random independent link failures [61].

[3]For performance measure *TEC*, the worst-case cut obtains a maximum value, while for the rest, it obtains a minimum value.

A solution to the BGNI optimization problem below is an endpoint of a worst-case cut.

$$\max \sum_{(i,j)} p_{ij} c_{ij} z_{ij}(x, y)$$

$$\text{such that}$$

$$0 \leq x \leq 1$$

$$-h \leq y \leq h_G \qquad (2.1)$$

The above optimization problem can be formulated as a Mixed Integer Linear Program (MILP) as follows. Define the following (0,1) variables:

$$u_{ij} = \begin{cases} 1 & \text{if } (i,j) \text{ crosses the cut location } (x) \text{ above } y \\ 0 & \text{otherwise} \end{cases}$$

$$d_{ij} = \begin{cases} 1 & \text{if } (i,j) \text{ crosses the cut location } (x) \text{ below } y + h \\ 0 & \text{otherwise} \end{cases}$$

For $h_G \leq 1$, the solution to the MILP below is a worst-case cut.

$$\max \sum_{(i,j)} p_{ij} c_{ij} z_{ij}$$

$$\text{such that}$$

$$(r_j - l_i)x - (y - l_i) \geq u_{ij} - 1 \quad \forall i, j$$

$$(y + h - l_i) - (r_j - l_i)x \geq d_{ij} - 1 \quad \forall i, j$$

$$u_{ij} + d_{ij} \geq 2z_{ij} \quad \forall i, j$$

$$0 \leq x \leq 1$$

$$-h \leq y \leq h_G$$

$$u_{ij}, d_{ij}, z_{ij} \in \{0, 1\}$$

29

Solving integer programs can be computationally intensive. Yet, the geographical (geometric) nature of the BGNI Problem lends itself to relatively low complexity algorithms (see Section 2.5). Although we initially focus only on the *TEC* measure, variants of the BGNI Problem can be formulated for performance measures *ATTR*, *MFST*, and *AMF* (by definition, when computing these measures we assume that $p_{ij} \in \{0, 1\} \, \forall i, j$). In the bipartite model, the worst-case cut under some of these measures is trivial. However, in the general model, a worst-case cut is non-trivial.

2.3.2 General Model

The general geometric graph model contains N non-overlapping nodes on a plane. Let the location of node i be given by the cartesian pair $[x_i, y_i]$. Assume the points representing the nodes are in general form, that is no three points are collinear. Denote a link from node i to node j as (i, j) and let (i, j) be represented by a line segment from $[x_i, y_i]$ to $[x_j, y_j]^4$. We define p_{ij} as the probability of (i, j) existing and c_{ij} as the capacity of (i, j) where $c_{ij} \in [0, \infty)$. We again assume that $c_{ij} p_{ij} > 0$ for some i and j. We now define two types of cuts and the corresponding problems.

When dealing with *Arbitrary Line Segment Cuts* we assume that a disaster results in a line segment cut of length h which starts at $[x, y]$ and contains the point $[v, w]$ (with $[x, y] \neq [v, w]$). We define this cut as $\text{cut}_h([x, y], [v, w])$ (note there can be infinitely many ways to express a single cut). A cut removes all links which intersect it. For brevity, we sometimes denote the worst-case cut $\text{cut}_h([x^*, y^*], [v^*, w^*])$ as cut_h^*. We now define the following problem and demonstrate its formulation.

Geographical Network Inhibition by Line Segments (GNIL) Problem: *Given a graph, cut length, link probabilities, and capacities, find a worst-case cut under per-*

---

[4]Notice that the assumption that links are represented by line segments is an approximation of the real deployments (e.g., [46]) in which links may not be linear.

formance measure TEC.

We define the following (0,1) variable:

$$z_{ij}([x,y],[v,w]) = \begin{cases} 1 & \text{if } (i,j) \text{ is removed} \\ & \text{by } \text{cut}_h([x,y],[v,w]) \\ 0 & \text{otherwise} \end{cases}$$

A solution to the GNIL optimization problem below is a worst-case cut.

$$\max \sum_{(i,j)} p_{ij} c_{ij} z_{ij}([x,y],[v,w])$$

such that

$$[x,y] \neq [v,w]$$

$$\sqrt{(x-v)^2 + (y-w)^2} \leq h$$

$$x_i \leq x \leq x_j \text{ for some } i \text{ and } j$$

$$y_i \leq y \leq y_j \text{ for some } i \text{ and } j \tag{2.2}$$

When dealing with *Circular Cuts* we assume that a disaster results in a cut of radius r which is centered at $[x,y]$. We define this cut as $\text{cut}_r(x,y)$. Such a cut removes all links which intersect it (including the interior of the disk). We call the set of points for which the Euclidean distance is r away from $[x,y]$ the boundary of $\text{cut}_r(x,y)$. For brevity, we sometimes denote the worst-case cut $\text{cut}_r(x^*,y^*)$ as cut_r^*. We now define the following problem and demonstrate its formulation.

Geographical Network Inhibition by Circular Cuts (GNIC) Problem: *Given a graph, cut radius, link probabilities, and capacities, find a worst-case circular cut under performance measure TEC.*

We define the following (0,1) variable:

$$z_{ij}(x,y) = \begin{cases} 1 & \text{if } (i,j) \text{ is removed by } \text{cut}_r(x,y) \\ 0 & \text{otherwise} \end{cases}$$
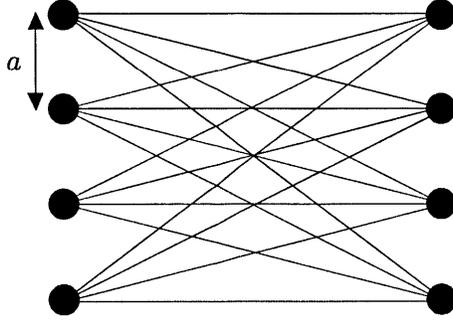
Figure 2-3: An example of a complete bipartite graph with $N = 8$.

A solution to the GNIC optimization problem below is the center of a worst-case cut.

$$\max \sum_{(i,j)} p_{ij} c_{ij} z_{ij}([x, y])$$

such that

$$x_i \leq x \leq x_j \text{ for some } i \text{ and } j$$

$$y_i \leq y \leq y_j \text{ for some } i \text{ and } j \tag{2.3}$$

Similar GNIL and GNIC problems can be formulated for performance measures *ATTR*, *MFST*, and *AMF* (for these measures we assume that $p_{ij} \in \{0, 1\} \, \forall i, j$). For example, under *MFST*, flow conversation constraints should be added to the set of constraints, the flow through links for which $z_{ij}([x, y], [v, w]) = 1$ is 0, and the flow between s and t has to be maximized. In sections 2.6 and 2.7 we use the geometric nature of the GNIL and GNIC problems to show that under all these measures, we only need to check a polynomial number of locations in order to find a worst-case cut.

2.4 A Motivating Example

In this section, we consider a simple case of the bipartite model in which the network is represented as a complete bipartite graph, each side has $N/2$ nodes, $p_{ij} = 1$, and $c_{ij} = 1$. We also place nodes evenly on each side such that they are separated by distance a. An example is shown in Fig. 2-3. We first obtain a lower bound for the BGNI problem by considering cuts down the center. Then, we provide numerical

32
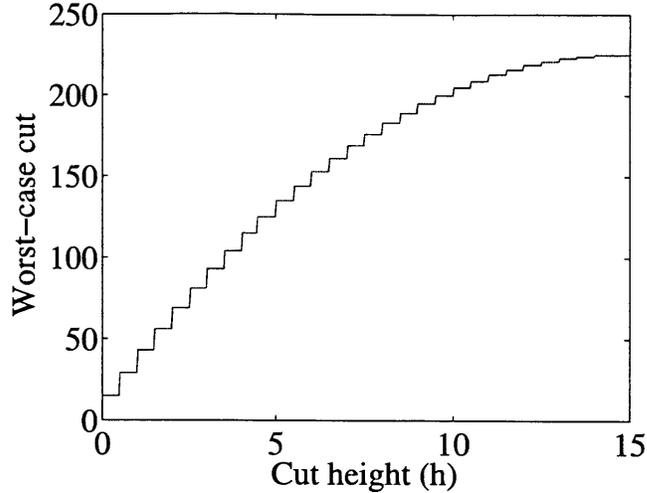
Figure 2-4: Number of links intersected (*TEC*) by a worst-case cut ($cut_h(x^*, y^*)$) as a function of the cut height (h) in a bipartite graph with 15 nodes on each side ($N = 30$).

results for the BGNI problem.

2.4.1 A Lower Bound

In this simple model, we can bound the value of *TEC* for the worst-case cut by considering cuts with endpoints at $x = 0.5$. In the very center of the graph there is an intersection of $N/2$ links. $a/2$ units vertically up and down from this point, an additional $(N/2) - 1$ links intersect. Another $a/2$ units up and down from these points, another $(N/2) - 2$ links intersect. This pattern continues until all of the links are included. Therefore, the capacity removed by a worst-case cut of height h for $h \leq h_G$ is lower bounded by:

$$\frac{N}{2} + \sum_{i=1}^{\lfloor \frac{2h}{a} \rfloor} (\frac{N}{2} - 1 - \lfloor \frac{i-1}{2} \rfloor). \tag{2.4}$$

2.4.2 Intuition from Numerical Results

We now describe numerical solutions obtained for the BGNI problem (2.1).[5] We obtained solutions for a network with 15 nodes on each side ($N = 30$) and with $a = 1$

---

[5]These solutions were initially obtained using MATLAB's genetic algorithms and later on verified using the algorithm described in Section 2.5.
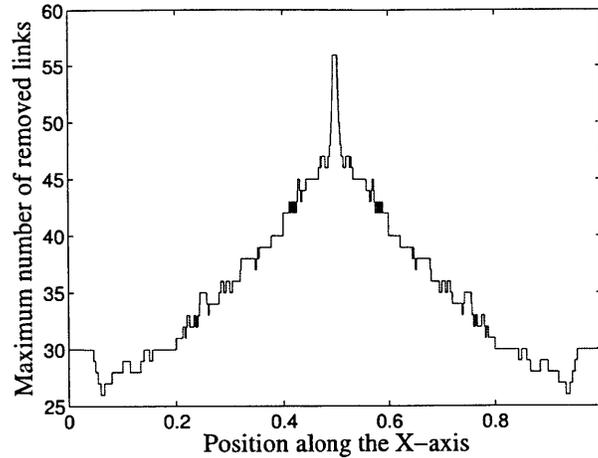
Figure 2-5: The maximum number of removed links (TEC) as a function of the x-location of the cut for $h = 1.6$. Note that the results were relatively monotonic, with the worst-case cut appearing at the center.
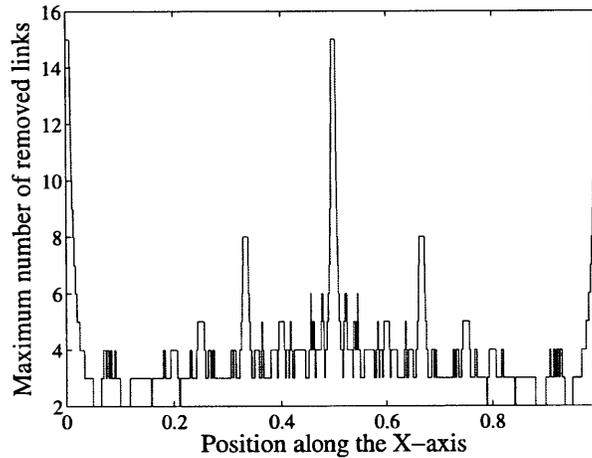
Figure 2-6: The maximum number of removed links (TEC) as a function of the x-location of the cut for $h = 0.1$. Note the two 'spikes' in the function at $x \approx .3$ and $x \approx .7$.

($h_G = 14$). Fig. 2-4 describes the values of TEC under the worst-case cut for different cut heights, h (notice that for $p_{ij} = 1$ and $c_{ij} = 1$, TEC is equivalent to the number of removed links). The result is identical to the lower bound for the center cuts in (2.4). This implies that a worst-case cut is located at the center of the graph.

Next, we study the effect of the horizontal cut location on TEC (the number of removed links) on the same network. Figures 2-5 and 2-6 illustrate the maximum number of removed links versus the horizontal (x) position of the cut on the network.

For a given cut height (h), the maximum number of removed links at each horizontal position (x) is not decreasing monotonically as we move away from the center. With

34

$h = 1.6$ the results were relatively monotonic, with the worst-case cut appearing at the center while the number of removed links more or less descends from there (Fig. 2-5). When the cut height is reduced to 0.1, significant local maxima begin to appear (Fig. 2-6). It seems the smaller the cut height, the more pronounced these local maxima are. This possibly results from large intersections of links crossing at different horizontal locations in the graph. Small cuts can cut these off-center intersections and remove a large number of links but these small cuts are not as effective elsewhere in the graph (where links do not intersect).

The results above motivate us to analytically study the effect of the cut location on the removed capacity. In the following sections, we focus on developing polynomial-time algorithms for identifying a worst-case cut.

2.5   Worst-Case Cuts - Bipartite Model

In this section we present an $O(N^6)$ algorithm for solving the BGNI problem. The main underlying idea is that the algorithm only needs to consider cuts which have an endpoint on a link intersection or a node. Before proceeding, we note that the objective function takes on a finite number of bounded values. This leads to the following observation.

Observation 1. *There always exists an optimal solution to (2.1) (i.e., a worst-case cut).*

Below, we present the algorithm which finds a worst-case cut. It can be seen that the complexity of Algorithm WCBG is $O(N^6)$. This results from the following facts: (i) links are line segments and a pair of line segments can have at most one intersection point (no three nodes are collinear), resulting in at most $O(N^4)$ link intersections; (ii) there are two candidate cuts per link intersection or node (cuts have two endpoints), and therefore, the total number of candidate cuts is at most $O(N^4)$; (iii) since evaluating $\mathbf{1}_{y_k \leq (r_j - l_i)x_k + l_i}\mathbf{1}_{y_k + h \geq (r_j - l_i)x_k + l_i}$ (Line 8) takes $O(1)$ time and it has to be evaluated for all (i, j), finding the capacity of a candidate cut takes

35

$O(N^2).$[6]

Algorithm 1 Worst-Case Cut in a Bipartite Graph (WCBG)

1: input: h, height of cut
2: worstCaseCapacityCut $\leftarrow 0$
3: **for** every node location and link intersection $[x_k, y_k]$ **do**
4:     **call evaluateCapacityofCut**(x_k, y_k)
5:     **call evaluateCapacityofCut**$(x_k, y_k - h)$
Procedure evaluateCapacityofCut(x_k, y_k)
6: capacityCut $\leftarrow 0$
7: **for** every (i, j) **do**
8:     **if** $1_{y_k \leq (r_j - l_i)x_k + l_i} 1_{y_k + h \geq (r_j - l_i)x_k + l_i} = 1$ **then**
9:         capacityCut $\leftarrow$ capacityCut $+ c_{ij}p_{ij}$
10: **if** capacityCut $\geq$ worstCaseCapacityCut **then**
11:     $x^* \leftarrow x_k$
12:     $y^* \leftarrow y_k$
13:     worstCaseCapacityCut $\leftarrow$ capacityCut

We now use a number of steps to prove the theorem below.

Theorem 1. *Algorithm WCBG finds a worst-case cut which is a solution to the optimization problem in (2.1).*

Before proving the theorem, we introduce some useful terminology and prove two supporting lemmas. If $\text{cut}_h(x, y)$ intersects any links, the links which are intersected closest to the endpoint $[x, y]$ are denoted by (i_α, j_α) and the point where they intersect the cut is denoted by $[x_\alpha, y_\alpha]$ (see Fig. 2-7 for an example). Let those links which intersect $\text{cut}_h(x, y)$ furthest from the endpoint $[x, y]$ be given by (i_ω, j_ω) and let the point where they intersect the cut be given by $[x_\omega, y_\omega]$. Note that (i_ω, j_ω) or (i_α, j_α) need not be unique. This is because $[x_\omega, y_\omega]$ or $[x_\alpha, y_\alpha]$ can be a link intersection. It should be noted that since the model assumes that there exists a link with $p_{ij}c_{ij} > 0$ for some i and j, all worst-case cuts must intersect at least one link. This implies (i_ω, j_ω) and (i_α, j_α) exist for all worst-case cuts.

---

[6]Computational geometry results can probably be used to reduce the complexity of Algorithm WCBG. Particularly, [25] (based on [13]), enables counting and locating all the intersections of N^2 line segments in $O(N^2 \log N + I)$ time, where I is the number of line segment intersections. A modified version of the algorithm of [25] can be used within Algorithm WCBG.
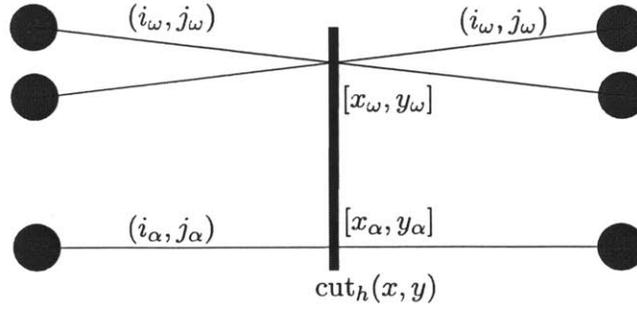
Figure 2-7: Example showing (i_ω, j_ω) and (i_α, j_α). (i_α, j_α) is the lowest link intersected by the cut and this intersection is at $[x_\alpha, y_\alpha]$. (i_ω, j_ω) are the highest links intersected by the cut and this intersection is at $[x_\omega, y_\omega]$. Note (i_ω, j_ω) is not unique.
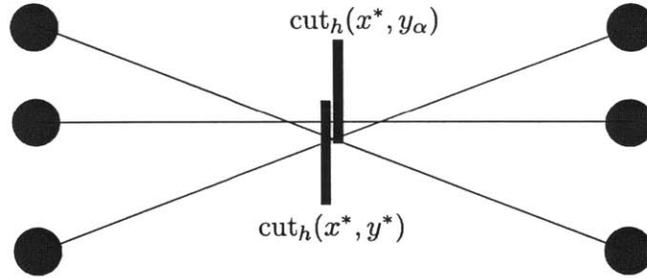
Figure 2-8: Example showing how $\text{cut}_h(x^*, y_\alpha)$ is a 'slid up' version of $\text{cut}_h(x^*, y^*)$. $\text{cut}_h(x^*, y_\alpha)$, which has an endpoint on a link intersection, is guaranteed to intersect every link $\text{cut}_h(x^*, y^*)$ does because there exist no links at x^* from y^* to y_α.

Lemma 1. *If there exists a worst-case cut, $\text{cut}_h(x^*, y^*)$, such that either (i_ω, j_ω) is not unique, (i_α, j_α) is not unique, or $x^* \in \{0, 1\}$, then there exists a worst-case cut that has an endpoint on a node or a link intersection.*

Proof. Assume (i_α, j_α) is not unique or $x^* \in \{0, 1\}$ ($[x^*, y_\alpha]$ is a node or link intersection). Consider $\text{cut}_h(x^*, y_\alpha)$ which is a 'slid up' version of the worst-case cut $\text{cut}_h(x^*, y^*)$. $\text{cut}_h(x^*, y_\alpha)$ intersects at least the same links as $\text{cut}_h(x^*, y^*)$ since, by definition of $[x_\alpha, y_\alpha]$, there exist no links at x^* from y^* to y_α. Thus, $\text{cut}_h(x^*, y_\alpha)$ is also a worst-case cut and has an endpoint on a node or link intersection. For an example, see Fig. 2-8. The case where (i_ω, j_ω) is not unique is analogous except that $\text{cut}_h(x^*, y_\omega - h)$, which is a 'slid down' version of $\text{cut}_h(x^*, y^*)$, is considered. □

Lemma 2. *If there exists a worst-case cut, $\text{cut}_h(x^*, y^*)$, such that both (i_ω, j_ω) and (i_α, j_α) are unique, then there exists a worst-case cut that has an endpoint on a link intersection or node.*

37

Proof. Let $y_\omega(x) = (r_\omega - l_\omega)x + l_\omega$ be the equation of (i_ω, j_ω) on $x \in [0, 1]$. Let $y_\alpha(x) = (r_\alpha - l_\alpha)x + l_\alpha$ be the equation of (i_α, j_α) on $x \in [0, 1]$. Let $y_{ij}(x) = (r_j - l_i)x + l_i$ be the equation of (i, j) on $x \in [0, 1]$.

Consider the slopes of $y_\omega(x)$ and $y_\alpha(x)$. There are two cases:

1. The slope of $y_\omega(x)$ is smaller or equal to the slope of $y_\alpha(x)$: $r_\omega - l_\omega \leq r_\alpha - l_\alpha$.

2. The slope of $y_\omega(x)$ is greater or equal to the slope of $y_\alpha(x)$: $r_\omega - l_\omega \geq r_\alpha - l_\alpha$.

We consider now the first case. Let:

$$
x' = \begin{cases}
\min x & \text{such that } x^* \leq x \leq 1 \text{ and} \\
& y_{ij}(x) = y_\alpha(x) \text{ for any } y_{ij} \text{ not } y_\alpha \text{ or} \\
& y_{ij}(x) = y_\omega(x) \text{ for any } y_{ij} \text{ not } y_\omega \\
1 & \text{if the } x \text{ above does not exist}
\end{cases}
$$

Essentially, x' is the first x-location after x^* where $y_\omega(x)$ or $y_\alpha(x)$ intersect another link. If $y_\omega(x)$ or $y_\alpha(x)$ do not intersect another link after x^*, then $x' = 1$.

We now show that x' is an x-location where it is possible to cut all the links which intersect $\text{cut}_h(x^*, y^*)$. Since links are line segments, we know $y_{ij}(x') = y_{ij}(x^*) + (x' - x^*)(r_j - l_i)$ $\forall i, j$. Since we know $y_\omega(x^*) \leq y_\alpha(x^*) + h$ $\big(\text{cut}_h(x^*, y^*)$ intersects both $y_\omega(x)$ and $y_\alpha(x)\big)$ and $(r_\omega - l_\omega)(x' - x^*) \leq (r_\alpha - l_\alpha)(x' - x^*)$ $\big(\text{case 1 above and } x' - x^* \geq 0\big)$, we have $y_\omega(x^*) + (r_\omega - l_\omega)(x' - x^*) \leq y_\alpha(x^*) + (r_\alpha - l_\alpha)(x' - x^*) + h$. Thus $y_\omega(x') \leq y_\alpha(x') + h$. See Fig. 2-9.

This means $\text{cut}_h(x', y_\alpha(x'))$ will intersect both (i_ω, j_ω) and (i_α, j_α). Since both these links do not intersect another link on $x^* \leq x < x'$, links which are intersected by $\text{cut}_h(x^*, y^*)$ are also intersected by $\text{cut}_h(x', y_\alpha(x'))$ (they are 'trapped' between (i_ω, j_ω) and (i_α, j_α) on $x^* \leq x < x'$).

Now we know $\text{cut}_h(x', y_\alpha(x'))$ is a worst-case cut and $x' = 1$, $[x', y_\alpha(x')]$ is a link intersection, or $[x', y_\omega(x')]$ is a link intersection. Therefore, by Lemma 1, we know there exists a worst-case cut which has an endpoint on a link intersection or node. The second case follows in an analogous fashion. □
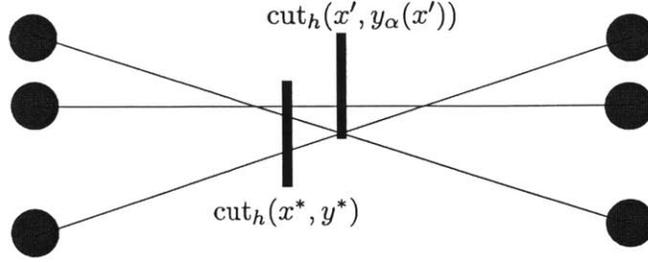
Figure 2-9: $\mathrm{cut}_h(x^*, y^*)$ is a worst-case cut and has a unique (i_ω, j_ω) and (i_α, j_α). From this we are able to find $\mathrm{cut}_h(x', y_\alpha(x'))$, a worst-case cut that has an endpoint on a link intersection.

Basically, according to Lemma 2, if (i_ω, j_ω) and (i_α, j_α) are both unique for a worst-case cut, we can find another worst-case cut such that it has at least one endpoint on a link intersection or node (see Fig. 2-9).

Using the above lemmas, we now prove Theorem 1.

Proof of Theorem 1: Since (i_ω, j_ω) and (i_α, j_α) exist for all worst-case cuts, Lemmas 1 and 2 imply that we need only check cuts which have endpoints at nodes or link intersections to find a worst-case cut. Algorithm 1 checks all possible nodes and intersections as endpoints, and therefore will necessarily find also a worst-case cut.

We note that although algorithm WCBG finds a worst-case cut, there may be other worst-case cuts with the same value. The endpoints of these cuts do not necessarily have to be on a link intersection or a node. However, there cannot be a cut with a higher value than the one obtained by the algorithm.

2.6 Worst-Case Line Segment Cut – General Model

In this section, we present a polynomial time algorithm for finding the solution of the GNIL Problem; i.e., for finding a worst-case line segment cut in the general model. We show that we only need to consider a polynomial-sized subset of all possible cuts. We first focus on the *TEC* performance measure and then discuss how to obtain a worst-case cut for other measures. Our methods are similar to the approach for solving the BGNI Problem, described in Section 2.5. In this section, a worst-case cut

39

refers to a worst-case *line segment* cut.

2.6.1 *TEC* **Performance Measure**

Before proceeding, note that the objective function in (2.2) takes on a finite number of bounded values. This leads to the following observation.

Observation 2. *There always exists an optimal solution to (2.2) (i.e., a worst-case cut).*

Below we present an algorithm that finds a worst-case line segment cut under the *TEC* measure in the general model. This algorithm considers all cuts that (i) have an endpoint on a link intersection and contain a node not at the intersection, (ii) have an endpoint on a link intersection and another endpoint on a link, (iii) contain two distinct nodes and have an endpoint on a link, and (iv) contain a node and have both endpoints on links.

We now use a number of steps to prove the theorem below.

Theorem 2. *Algorithm WLGM has a running time of $O(N^8)$ and finds a worst-case line segment cut that is a solution to the GNIL Problem.*

Before proving the theorem we present some lemmas to reduce the set of candidate worst-case cuts.

Lemma 3. *There exists a worst-case cut that contains a node or has an endpoint at a link intersection.*

Proof. Let cut_h^* be a worst-case cut with endpoints given by $[x^*, y^*]$ and $[v^*, w^*]$. We now define some useful terminology. Let the links that intersect cut_h^* closest to the endpoint $[x^*, y^*]$ be given by (i_α, j_α) and let the closest point to $[x^*, y^*]$ where (i_α, j_α) intersects cut_h^* be given by $[x_\alpha, y_\alpha]$. Let those links which intersect cut_h^* furthest from the endpoint $[x^*, y^*]$ be given by (i_ω, j_ω) and let the closest point to $[v^*, w^*]$ where (i_ω, j_ω) intersects cut_h^* be given by $[x_\omega, y_\omega]$. We consider two cases, one where either (i_α, j_α) or (i_ω, j_ω) are not unique and the other where (i_α, j_α) and (i_ω, j_ω) are unique.

40

Algorithm 2 Worst-Case Line Segment Cut in the General Model (WLGM)

1: input: h, length of cut
2: worstCaseCapacityCut $\leftarrow 0$
3: $L \leftarrow \{\}$
4: **for** every link intersection $[x_k, y_k]$ **do**
5:     **for** every node i such that $[x_i, y_i] \neq [x_k, y_k]$ **do**
6:         $L = L \cup \{$cut that has an endpoint at $[x_k, y_k]$ and contains $[x_i, y_i]\}$
7:     **for** every (i, j) **do**
8:         $L = L \cup \{$cuts that have an endpoint at $[x_k, y_k]$ and another endpoint on $(i, j)\}$
9: **for** every (i, j) and node k **do**
10:     **for** every node l such that $k \neq l$ **do**
11:         $L = L \cup \{$cuts that have an endpoint on (i, j) and contain $[x_k, y_k]$ and $[x_l, y_l]\}$
12:     **for** every (m, n) **do**
13:         $L = L \cup \{$cuts that have an endpoint on (i, j), another endpoint on (m, n), and contain $[x_k, y_k]\}$
14: **for** every $\text{cut}_h([x_k, y_k], [v_k, w_k]) \in L$ **do**
15:     **call evaluateCapacityofCut**(x_k, y_k, v_k, w_k)
16: **return** cut_h^*
Procedure evaluateCapacityofCut(x_k, y_k, v_k, w_k)
17: capacityCut $\leftarrow 0$
18: **for** every (i, j) **do**
19:     **if** $z_{ij}([x_k, y_k], [v_k, w_k]) = 1$ **then**
20:         capacityCut $\leftarrow$ capacityCut $+ c_{ij}p_{ij}$
21: **if** capacityCut $\geq$ worstCaseCapacityCut **then**
22:     $\text{cut}_h^* \leftarrow \text{cut}_h([x_k, y_k], [v_k, w_k])$
23:   worstCaseCapacityCut $\leftarrow$ capacityCut
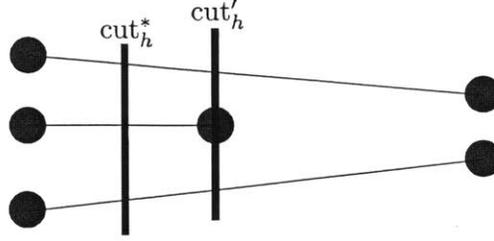
Figure 2-10: cut'_h contains a node as well as intersects all links which cut^*_h does.

In the first case, either (i_α, j_α) or (i_ω, j_ω) are not unique for cut^*_h. Without loss of generality, we assume (i_α, j_α) is not unique. We consider cut'_h which is a translated version of cut^*_h such that it has an endpoints on $[x_\alpha, y_\alpha]$ and on $[v^* + x_\alpha - x^*, w^* + y_\alpha - y^*]$. Since there exist no links between $[x^*, y^*]$ and $[x_\alpha, y_\alpha]$, we know cut'_h intersects at least as many links as cut^*_h and thus is a worst-case cut that satisfies the lemma. Fig. 2-8 shows the analogous case for the bipartite model.

In the second case, (i_α, j_α) and (i_ω, j_ω) are both unique for cut^*_h. If cut^*_h contains a node, the lemma is satisfied. In the following, assume cut^*_h does not contain a node. Now we consider $cut'_h([x^* + a, y^* + b], [v^* + a, w^* + b])$ and $cut''_h([x^* - c, y^* - d], [v^* - c, w^* - d])$ to be translated versions of cut^*_h such that (i) $sign(a) = sign(c)$ and $sign(b) = sign(d)$, (ii) there does not exist any nodes in the parallelogram defined by cut^*_h and cut'_h (which we denote "parallelogram B") except those contained in cut'_h and in the parallelogram defined by cut^*_h and cut''_h (which we denote "parallelogram C") except those contained in cut''_h, and (iii) no link intersects (i_α, j_α) or (i_ω, j_ω) in either parallelogram except on cut'_h or cut''_h. Since a node does not exist within the interior of either parallelogram all links intersected by cut^*_h must also cut one of the other three edges of each parallelogram.

Now choose the maximum a and c such that the edge $([x^*, y^*], [x^* + a, y^* + b])$ of parallelogram B and the edge $([x^*, y^*], [x^* - c, y^* - d])$ of parallelogram C are both parallel to the link (i_α, j_α) and the parallelograms satisfy the constraints in the paragraph above. This implies both cut'_h and cut''_h contain a node or contain a point where (i_α, j_α) or (i_ω, j_ω) intersects a link. Since (i_α, j_α) is parallel to both edges $([x^*, y^*], [x^* + a, y^* + b])$ and $([x^*, y^*], [x^* - c, y^* - d])$ and since (i_ω, j_ω) can cut at

42

most one of the edges $([v^*, w^*], [v^* + a, w^* + b])$ and $([v^*, w^*], [v^* - c, w^* - d])$ or be parallel to them (as they both lay on the same straight line), we know at least one of cut'_h or cut''_h intersects the same links that are intersected by cut^*_h. Therefore, we can choose a, b, c, and d such that either cut'_h or cut''_h is a worst-case cut and (i) contains a node (Fig. 2-10) or (ii) contains a point where (i_α, j_α) or (i_ω, j_ω) intersects a link. In the latter case, we can translate this worst-case cut in a similar fashion to the first case to construct a worst-case cut which satisfies the lemma. $\qquad\square$

We now consider two cases of worst-case cuts. The first case is a worst-case cut that has an *endpoint* at a link intersection. The second case is a worst-case cut that *contains* a node. In both cases, let the node or link intersection that is in the cut be denoted by A. Lemma 4 handles the first case where A can be considered as a link intersection.

Lemma 4. *If there exists a worst-case cut that has an endpoint on point A, then (i) there exists a worst-case cut that has an endpoint on A and has its other endpoint on a link or (ii) there exists a worst-case cut that has an endpoint on A and contains a node that is not A.*

Proof. Assume there exists a worst-case cut with endpoint A, denoted by cut^*_h. Therefore, the other endpoint of cut^*_h must be on a circle of radius h. Denote by θ the angle of cut^*_h in some coordinate system. Denote by θ_i the angles from A to all nodes inside the circle and all intersections of links with the circle (including links tangent to the circle). Choose $\theta' = \theta_j$ such that $j = \arg\min_i |\theta - \theta_i|$. Choose cut'_h to be the cut with endpoint at A and having length h and angle θ'. By definition of θ' and the θ_i's, all links intersecting cut^*_h must also intersect cut'_h (because between θ and θ' no link intersects with the circle and there exists no node within the interior of that sector). Thus, cut'_h is a worst-case cut (see Fig. 2-11). $\qquad\square$

The following two lemmas handle the second case where A can be considered as a node.

Lemma 5. *If there exists a worst-case cut that contains point A, then there exists a worst-case cut that contains A and has an endpoint on some link.*
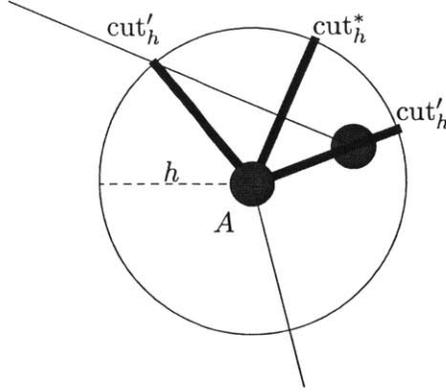
Figure 2-11: Translate an endpoint of cut_h^* along the circumference of the circle until the cut intersects a node or the translated endpoint intersects a link; call this new cut cut_h'. Since every link which intersects cut_h^* intersects cut_h', cut_h' is a worst-case cut.
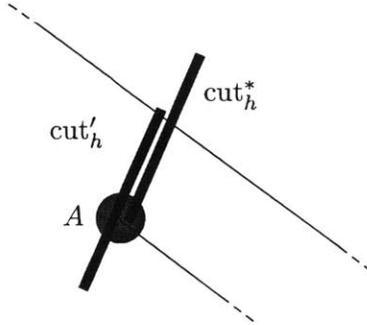
Figure 2-12: Translate cut_h^* along the line which contains it until one of its endpoints intersects a link; we call this new cut cut_h'. cut_h' intersects all links cut_h^* intersects.

Proof. Let cut_h^* be a worst-case cut that intersects A with endpoints given by $[x^*, y^*]$ and $[v^*, w^*]$. Let the links that intersect cut_h^* closest to the endpoint $[x^*, y^*]$ be given by (i_α, j_α) and let the closest point to $[x^*, y^*]$ where (i_α, j_α) intersects cut_h^* be given by $[x_\alpha, y_\alpha]$. We consider cut_h' which is a translated version of cut_h^* such that it has endpoints at $[x_\alpha, y_\alpha]$ and at $[v^* + x_\alpha - x^*, w^* + y_\alpha - y^*]$. Since there exist no links between $[x^*, y^*]$ and $[x_\alpha, y_\alpha]$, and because the same line contains both cut_h^* and cut_h', we know that every link which intersects cut_h^* also intersects cut_h' in the same location (see Fig. 2-12). Thus, cut_h' is a worst-case cut that contains A and has an endpoint on a link (this endpoint is $[x_\alpha, y_\alpha]$). $\qquad\square$

Lemma 6. *If there exists a worst-case cut that contains A and has an endpoint on a link, then there exists a worst-case cut that contains A, has an endpoint on a link,*
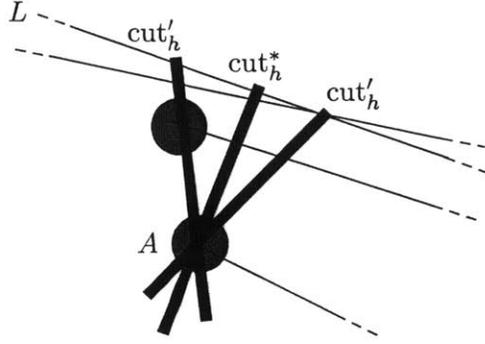
Figure 2-13: Translate an endpoint of cut$_h^*$ right along L until it intersects a link intersection. This new cut is the cut$_h'$ on the right. We can also translate an endpoint of cut$_h^*$ left along L until it intersects a node. This new cut is the cut$_h'$ on the left.

and at least one of the following holds: (i) the cut contains a node that is not A, (ii) one of the cut endpoints is also a link intersection that is not A , or (iii) the cut has both endpoints on links.

Proof. Let cut$_h^*$ be a worst-case cut such that it contains A and has an endpoint on a link. If cut$_h^*$ has an endpoint on A, then Lemma 4 implies Lemma 6. Assume cut$_h^*$ contains A and has an endpoint on a link and does *not* have an endpoint on A. Denote the link which contains this endpoint by L, and one of its endpoints by $[x_1, y_1]$. Denote the point at which cut$_h^*$ intersects L by $[x_0, y_0]$. Now translate the endpoint of cut$_h^*$ along L so that this new cut still contains A. That is, consider the cut, of length h, with endpoint at $[ax_1 + (1-a)x_0, ay_1 + (1-a)y_0]$ and passing through A, for $0 \leq a \leq 1$. For $a = 0$ this is just cut$_h^*$. We increase a until a new cut, called cut$_h'$, either has an endpoint that is h away from A (we cannot translate further) or cut$_h'$ can no longer satisfy $\sum_{(i,j)} p_{i,j} c_{i,j} \text{cut}_h' = \sum_{(i,j)} p_{i,j} c_{i,j} \text{cut}_h^*$. In the first case, the cut has both endpoints on links. In the second case cut$_h'$ satisfies at least one of the following: cut$_h'$ has an endpoint on L that is a link intersection (considered in Lemma 4), cut$_h'$ intersects a node which is not A, or cut$_h'$ has an endpoint on L and the other endpoint on a link. The first two possibilities are demonstrated in Fig. 2-13. Fig. 2-14, which demonstrates the third possibility, shows cut$_h'$ that contains A and has both endpoints on links. $\square$

Using the lemmas above we now prove Theorem 2.
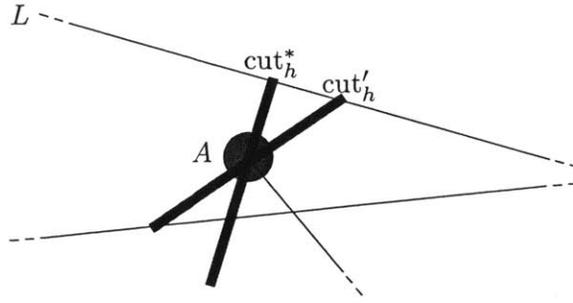
Figure 2-14: Translate an endpoint of cut_h^* along L until it can no longer intersect the bottom link. This new cut is cut_h'.

Proof of Theorem 2: The lemmas presented in this section imply we only need to consider a polynomially sized set of cuts. By Lemma 3 there are two possible cases of worst-case cuts. The first case is a worst-case cut which has a endpoint at a link intersection. The second case is a worst-case cut which contains a node. In the first case, Lemma 4 implies that for every link intersection, $O(N^4)$, there exists a possible worst-case cut for every link and node, $O(N^2)$. In the second case, Lemmas 5 and 6 imply that for every node-link pair (A and some link L), $O(N^3)$, there exist several possible worst-case cuts for every node and link, $O(N^2)$. Since naively checking each cut for the total cut capacity takes $O(N^2)$, the algorithm has a total running time of $O(N^8)$ (the first case provides the greatest running time).

It should be noted that similarly to the bipartite case, although the algorithm finds a worst-case cut, there may be other worst-case cuts with the same value. However, there cannot be a cut with a better value than the one obtained by the algorithm.

2.6.2 *ATTR*, *MFST*, and *AMF* Performance Measures

As mentioned in Section 2.3.2, the formulation of the GNIL Problem, presented in (2.2) should be slightly modified in order to accommodate the *ATTR*, *MFST*, and *AMF* performance measures. We now briefly discuss how the algorithm has to be modified in order to obtain results for these problems. In Section 2.8, we present numerical results obtained using these modified algorithms. Using the lemmas and theorem above, it is easy to show that only a polynomial number of candidate cuts

need to be checked in order to find the worst-case cut under any of the performance measures. This is due to the fact that the performance measures are monotonic. Therefore, any additional link removed/added only increases/decreases the measure and all the arguments supporting our lemmas still hold.

For each potential cut some links and/or nodes are removed. Hence, one has to update the network adjacency matrix. Then, different operations have to be performed for each measure:

- *ATTR* - If the network is fully connected, the value of *ATTR* is 1. Otherwise, one has to sum over all components the value of $k(k - 1)$, where k is the number of nodes in each of the components. Then the sum has to be divided by $N(N - 1)$. In order to verify connectivity or to count the number of nodes in each component, Breadth First Search (BFS) algorithm or the adjacency matrix eigenvalues and eigenvectors can be used.

- *MFST* - Run a max-flow algorithm (e.g., $O(N^3)$ [2]).

- *AMF* - Run a max-flow algorithm for every node pair.

2.7   Worst-Case Circular Cut − General Model

In this section we present a polynomial time algorithm for finding a solution of the GNIC Problem; i.e., for finding a worst-case circular cut in the general model. We show that we only need to consider a polynomial-sized subset of all possible cuts. We focus on the *TEC* performance measure and then briefly discuss how to obtain a worst-case cut for the other performance measures. In this section, a cut refers to a *circular* cut of a particular radius.

Before proceeding, note that the objective function in (2.3) takes on a finite number of bounded values. This leads to the following observation.

Observation 3. *There always exists an optimal solution to (2.3) (i.e., a worst-case cut).*

Above, we present an algorithm which finds a worst-case circular cut under the *TEC* measure in the general model.

Algorithm 3 Worst-Case Circular Cut in the General Model (WCGM)

1: input: r, radius of cut
2: worstCaseCapacityCut $\leftarrow 0$
3: $L \leftarrow \{\}$
4: **for** every (i,j) **do**
5:     $L = L \cup \{$cuts that intersect (i,j) at exactly one point and are centered on the line which contains $(i,j)\}$
6:     **for** (k,l) such that $(i,j) \neq (k,l)$ **do**
7:         **if** (i,j) is parallel to (k,l) **then**
8:             $L = L \cup \{$cuts that contain node i or j on its boundary and intersect (k,l) at exactly one point$\}$
9:         **else**
10:             $L = L \cup \{$cuts that intersect (i,j) and (l,k) at exactly one point each such that these points are distinct$\}$
11: **for** every cut$_r(x_k, y_k) \in$ L **do**
12:     **call evaluateCapacityofCut**(x_k, y_k)
13: **return** cut$_r^*$
Procedure evaluateCapacityofCut(x_k, y_k)
14: capacityCut $\leftarrow 0$
15: **for** every (i,j) **do**
16:     **if** minimum distance from (i,j) to $[x_k, y_k]$ is $\leq r$ **then**
17:         capacityCut $\leftarrow$ capacityCut $+ c_{ij}p_{ij}$
18: **if** capacityCut $\geq$ worstCaseCapacityCut **then**
19:     cut$_r^* \leftarrow$ cut$_r(x_k, y_k)$
20:     worstCaseCapacityCut $\leftarrow$ capacityCut

Theorem 3. *Algorithm WCGM has a running time of $O(N^6)$ and finds a worst-case circular cut which is a solution to the GNIC Problem.*

Before proving the theorem, we present a useful lemma about cuts and line segments and then present some lemmas to reduce the set of candidate cuts.

Lemma 7. *If a line segment intersects only the boundary of a cut, then the line segment and cut intersect at exactly one point.*

Proof. Proof by contradiction. Assume a line segment intersects only the boundary of a cut and this intersection contains more than one point. Since a line segment and a cut region are both convex, their intersection must be convex as well. However,

48
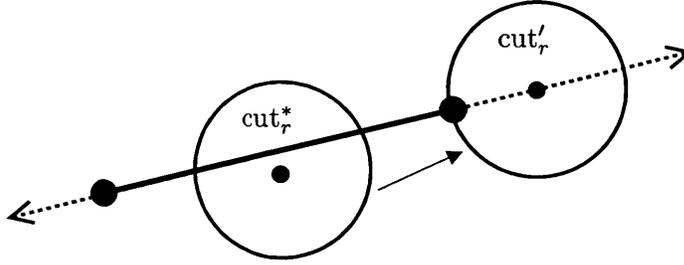
Figure 2-15: An example illustrating Lemma 8. cut'_r is a translated version of cut^*_r such that $[x', y']$ lies on the line which contains the intersected link and cut'_r intersects the link at exactly one point (recall $[x', y']$ is the center of cut'_r).

we assumed at least two points on the boundary of the cut are in the intersection. The fact that the intersection must be convex implies the chord connecting these two points must be in the intersection as well. Since part of the chord is in the interior of the cut, this leads to a contradiction. $\square$

Lemma 8. *If there exists a worst-case cut, denoted by* cut^*_r, *which intersects exactly one link, then there exists a worst-case cut, denoted by* cut'_r, *that intersects this link at exactly one point such that* $[x', y']$ *lies on the line which contains the link (recall* $[x', y']$ *is the center of* cut'_r).

Proof. Since cut^*_r is a worst-case cut and only intersects a single link, any cut that intersects the same link is also a worst-case cut. See Fig. 2-15. $\square$

Lemma 9. *If there exists a worst-case cut, denoted by* cut^*_r, *that intersects at least two links, then there exists a worst-case cut, denoted by* cut'_r, *that intersects at least two links at exactly one point each and at least one of the following holds: (i) at least two of the points are distinct and are not diametrically opposite, (ii) at least two of the points are distinct and one of them is a node, or (iii)* $[x', y']$ *lies on a line containing one of the two links.*

The proof of the lemma above is similar to the proofs of the lemmas in Section 2.6. Essentially, it is shown that we can translate a worst-case cut such that it remains a worst-case cut and satisfies the properties in the lemma.

Proof. Assume a link that intersects cut^*_r has node locations given by $[x_i, y_i]$ and $[x_j, y_j]$. Consider $\text{cut}_r[x^* + h(x_j - x_i), y^* + h(y_j - y_i)]$ where h is the minimum non-

49

negative value such that only the boundaries of this cut and some link intersect. Denote this translation of cut_r^* by cut_r'' and note by Lemma 7 this cut must intersect at least one link at exactly one point. Every link which is intersected by cut_r^* must intersect cut_r'' because as a line segment and a cut are continuously translated away from each other, the last non-empty intersection is an intersection of their boundaries. Thus, cut_r'' is also a worst-case cut. In the proceeding we consider two cases. In the first case we assume cut_r'' intersects at least two links at exactly one point each and in the second case we assume cut_r'' intersects exactly one link at exactly one point.

We first consider the case where cut_r'' intersects at least two links at exactly one point each (in addition to possibly other links that intersect the interior of cut_r''). Denote one of the points by A and another by B. If A and B are distinct and not diametrically opposite, the conditions in the lemma are satisfied. Now we will consider two sub-cases. In the first sub-case, we assume A and B reside in two diametrically opposing points on cut_r'' and in the second sub-case we assume A and B are not distinct. In the first sub-case, if either A or B is a node, the lemma holds true. If neither A or B are nodes, then A and B are diametrically opposing points where parallel links are tangent to cut_r''. Denote one of these parallel links by (i, j). Now consider $\text{cut}_r[x'' + h(x_j - x_i), y'' + h(y_j - y_i)]$ where h is the minimum nonnegative value such that two links intersect only the boundary of this cut at distinct and non-diametrically opposing points or two links intersect only the boundary of this cut and one of these intersection points is a node. Denote this translated cut by cut_r'. Now, by Lemma 7 one of the following must hold: either cut_r' intersects the parallel links at exactly one point each where one of these points is a node, or a link which intersected the interior of cut_r'' now intersects cut_r' at exactly one point such that cut_r' intersects two links at exactly one point each such that they are not diametrically opposite and distinct.

In second sub-case, two links intersect cut_r'' at a single point, C. This implies C is a node of at least one of these links. Now choose a link with a node given by C and denote the link by (k, l). Let $\mathbf{p}(t)$ be a continuous parameterized closed curve which is always a distance r from (k, l) such that $\mathbf{p}(0) = [x'', y'']$ and $\mathbf{p}(t_C)$ where $t_C > 0$ is

50
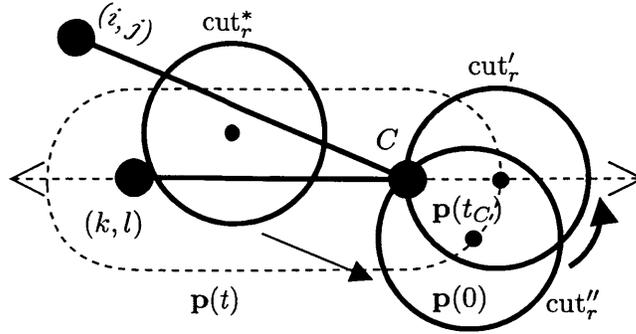
Figure 2-16: A case in the proof of Lemma 9. cut_r^* is first translated in the direction of (i,j) to become cut_r'' which intersects (k,l) at exactly one point and intersects another link (in this case (i,j)) at exactly the same point. Then cut_r'' is translated along $\mathbf{p}(t)$ towards $\mathbf{p}(t_C)$ to cut_r' such that $[x',y']$ lies on the line which contains (k,l).

the point on $\mathbf{p}(t)$ closest to C that intersects the line containing (k,l) (see Fig. 2-16). Additionally, we require that $\mathbf{p}(t)$ is exactly r units away from C for $0 \le t \le t_C$. Let $p_x(t)$ and $p_y(t)$ denote the x and y components of $\mathbf{p}(t)$ respectively. Since cut_r'' intersects C, we know $[x'',y'']$ is on a semi-circular shaped part of $\mathbf{p}(t)$ (these are the only parts of $\mathbf{p}(t)$ that are r units away from an endpoint of (k,l)). Now consider $\text{cut}_r[p_x(t),p_y(t)]$ where t is the minimum value such that two links intersect only the boundary of this cut and these intersection points are distinct or $t = t_C$. Denote this translated cut by cut_r'. If $t = t_C$ we know cut_r' is centered on the line which contains (k,l). As before, we know every link which is intersected by cut_r'' must intersect cut_r'. This is because as a line segment and a cut are continuously translated away from each other, the last non-empty intersection is an intersection of their boundaries. Also, the links that intersect cut_r'' at C remain intersected throughout the translation because $\text{cut}_r[p_x(t),p_y(t)]$ intersects C on $0 \le t \le t_C$. Thus, cut_r' is a worst-case cut and by Lemma 7 we know two links intersect this cut at exactly one point each and one of the following: i) these points are distinct and one of them is a node given by C or ii) $[x',y']$ lies on a line that contains (k,l) ($[x',y'] = \mathbf{p}(t_C)$).

Now we consider the case where cut_r'' intersects exactly one link at exactly one point (in addition to other links that intersect the interior of cut_r''). Similarly as above, denote this link by (k,l). Let $\mathbf{p}(t)$ be a continuous parameterized closed curve which is always a distance r from (k,l) such that $\mathbf{p}(0) = [x'',y'']$ (see Fig. 2-17). Consider $\text{cut}_r[p_x(t),p_y(t)]$ where t is the minimum nonnegative value such that two

51
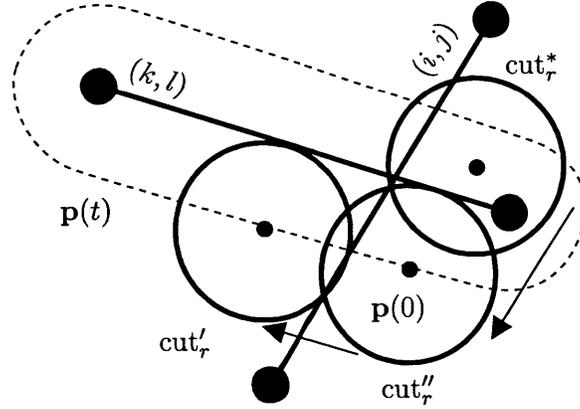
Figure 2-17: A case in the proof of Lemma 9. cut_r^* is first translated in the direction of (i,j) to become cut_r'' which intersects (k,l) at exactly one point. Then cut_r'' is translated along $p(t)$ to cut_r' where (i,j) and (k,l) each intersect cut_r' at exactly one point.

links intersect only the boundary of this cut (we assume cut_r^* intersects at least two links). By Lemma 7 we know these two links intersect this cut at exactly one point each. So this case reduces to the first case for which we know the lemma holds. $\square$

Lemma 10. *There are at most* 20 *cuts of radius* r *that intersect two non-parallel line segment links at exactly one point each such that these points are distinct.*

Proof. If a link intersects a cut at exactly one point, then either a node of the link intersects the boundary of the cut or the link is tangent to the cut (we call a link tangent to a cut if the line containing the link is tangent to the boundary of the cut). For a particular pair of links, this implies a cut that satisfies the lemma falls into at least one of three cases: i) the boundary of the cut intersects two distinct nodes (one from each link), ii) the boundary of the cut intersects a node of one link and the cut is tangent to the other link, or iii) both links are tangent to the cut.

In case one, by geometry we know there are at most two cuts of radius r whose boundary contains two distinct nodes. In case two, given a node and a link, we know by geometry there are at most two cuts of radius r that the link is tangent to and whose boundary contains the node. In case three, given two non-parallel links, the lines containing these segments divide the plane into four pieces. There exist at most one cut tangent to both lines in each of these pieces. Thus, there are at most four cuts tangent to both links. Since for a pair of non-parallel links there are four pairs

52

of nodes to consider (with at most two cuts per pair that satisfy the lemma), four endpoint-link pairs (with at most two cuts per pair that satisfy the lemma), and one link-link pair (with at most four cuts per pair that satisfy the lemma), we know there exists at most 20 cuts that satisfy the lemma. □

Note that the bound above is a simple upper bound on the number of possible cuts and can possibly be further reduced.

Using the above lemmas, we now prove Theorem 3.

Proof of Theorem 3: The lemmas presented in this section imply there exists a worst-case cut that intersects a link at exactly one point such that the center of this cut lies on the line containing this link or there exists a worst-case cut that intersects two links at exactly one point each and at least one of the following: (i) at least two of the points are distinct and are not diametrically opposite or (ii) at least two of the points are distinct and one of them is a node. Algorithm WCGM enumerates all these possible cuts. It considers each link, $O(N^2)$, and finds both cuts that intersect the link at exactly one point and whose center lies on the line which contains this link. Then, it considers every combination of two links, $O(N^4)$, and if the links are not parallel it finds every cut (if any exist) which intersect each of the two links at exactly one point such that these points are distinct. By Lemma 10 we know there are at most 20 of these cuts for every pair of links. If the links are parallel, we need only consider cuts that intersect one of the links at exactly one point and whose boundary intersects the other links endpoint. In total, Algorithm WCGM considers $O(N^4)$ cuts and since naively checking each cut for the total expected capacity removed takes $O(N^2)$, the algorithm has a total running time of $O(N^6)$.

As mentioned in Section 2.3.2, the formulation of the GNIC Problem, presented in (2.3), can be slightly modified in order to accommodate the *ATTR*, *MFST*, and *AMF* performance measures. This modification is done in exactly the same way as it was done for the GNIL Problem (see Section 2.6.2).

It should be noted that we can also consider the case of an *elliptic cut* with fixed axis (that is, no rotation of the ellipse is considered). This disaster model more closely resembles the effect of an EMP. This case can be solved by applying an affine

Figure 2-18: Line segments cuts optimizing TEC for $h = 2$ - the red cuts maximize TEC and the black segments are nearly worst-case cuts.

transformation to the network node locations and then running WCGM.

2.8 Numerical Results

In this section we present numerical results that demonstrate the use of the algorithms presented in sections 2.6 and 2.7. These results shed light on the vulnerabilities of a specific fiber network. Clearly, the algorithms can be used in order to obtain results for additional networks or for a combined fiber plant of several operators. The results were obtained using MATLAB.

We used Algorithm WLGM, presented in Section 2.6, to compute worst-case cuts under the TEC, $ATTR$, $MFST$, and AMF performance measures for a fiber plant of a major network provider [46]. In all cases, we found that the results are intuitive. We also used Algorithm WCGM, presented in Section 2.7, to compute worst-case circular cuts under the $MFST$ performance measure for the same fiber plant. We found these circular cuts are in similar locations to their line segment counterparts. All distance units mentioned in this section are in longitude and latitude coordinates (one unit is approximately 60 miles) and for simplicity we assume latitude and longitude coordinates are projected directly to $[x, y]$ pairs on the plane. We also assume that all the link capacities are equal to 1.

Fig. 2-18 presents line segment cuts of $h = 2$ which maximize the TEC performance measure. As expected, we find that TEC is large in areas of high link density,

54

Figure 2-19: Line segments cuts optimizing the $ATTR$ for $h = 2$ - the red cuts minimize $ATTR$ and the black segments are nearly worst-case cuts.

such as areas in Florida, New York, and around Dallas. Fig. 2-19 presents line segment cuts of $h = 2$ which minimize the $ATTR$ performance measure. $ATTR$ is smallest where parts of the network are disconnected, such as at the southern tip of Texas, Florida and most of New England. This is intuitive since in order to decrease the $ATTR$, the graph must be split and under a small cut, only small parts of the graph can be removed.

Fig. 2-20 illustrates line segment cuts of $h = 4$ which minimize the $MFST$ performance measure between Los Angeles (s) and New York City (NYC) (t). Removal of the s and t nodes themselves is not considered as this is a trivial worst-case cut. We found that $MFST$ is smallest directly around Los Angeles and NYC as well as in Colorado, Utah, Arizona, New Mexico, and Texas. There are also cuts in the East Coast which completely disconnect NYC from Los Angeles without actually going through NYC. The cuts in the southwest are intuitive since the network in that area is very sparse. In some sense, the fact that in this case we obtain expected results validates the assumptions and approximations.

We note that different networks (e.g., networks in Europe or Asia) have a different structure than the sparse structure of the southwest U.S. network. In such cases, the solution will not be straightforward. In order to demonstrate it, we will discuss below the $MSFT$ measure between NYC and Forth-Worth. Before that, we present in Fig. 2-21 line segment cuts of $h = 2$ which minimize the AMF performance measure. The AMF values are minimized by cuts in the southwest as well as in Florida and
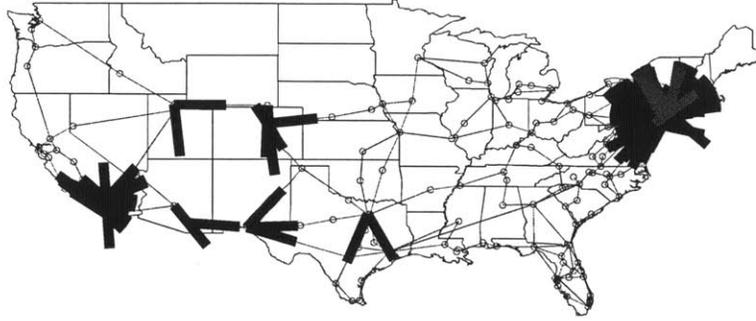
Figure 2-20: Line segments cuts optimizing $MFST$ between Los Angeles and NYC for $h = 4$ - the red cuts minimize $MFST$ and the black segments are nearly worst-case cuts. Cuts which intersect the nodes representing Los Angeles or NYC are not shown.
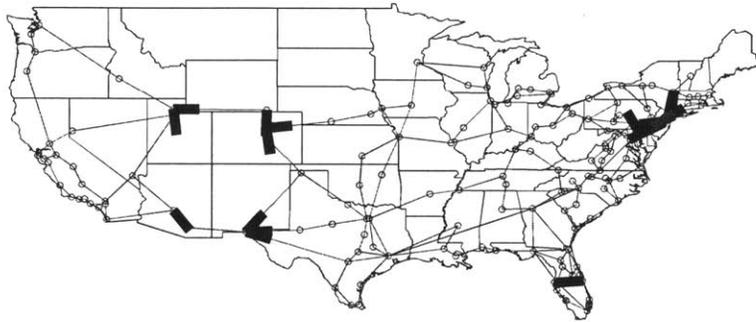
Figure 2-21: Line segments cuts optimizing the AMF for $h = 2$ - the red cuts minimize AMF and the black segments are nearly worst-case cuts.

New York.

Finally, we tested how line segment cuts compare to circular cuts. Using Algorithm WCGM we found circular cuts of $r = 2$ which minimize the $MFST$ performance measure between Los Angeles and NYC (see Fig. 2-22). Our results were similar to the line segment case; worst-case circular cuts were found close to both to Los Angeles and NYC. The southwest area also appeared to be vulnerable, just as in the line segment case.

As mentioned above, we tested the $MFST$ measure for circular cuts between Fort Worth and NYC (see Fig. 2-23). Due to the complexity of the network along the east coast, the results were less straightforward than in the Los Angles-NYC case.

Finally, for a circular cut in the fiber plant illustrated in Fig. 2-1, we computed the maximum value of TEC (removed capacity) as a function of the cut radius. The results are illustrated in Fig. 2-24. As expected, the maximum value of TEC

56
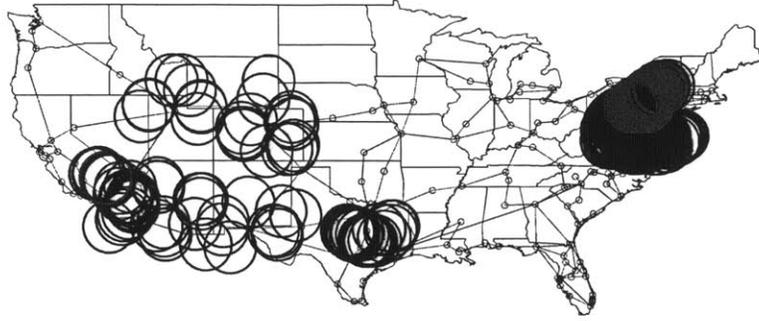
Figure 2-22: The impact of circular cuts of radius 2 on the $MFST$ between Los Angeles and NYC. Red circles represent cuts that result in $MFST = 0$ and black circles result in $MFST = 1$. Cuts which intersect the nodes representing Los Angeles or NYC are not shown.
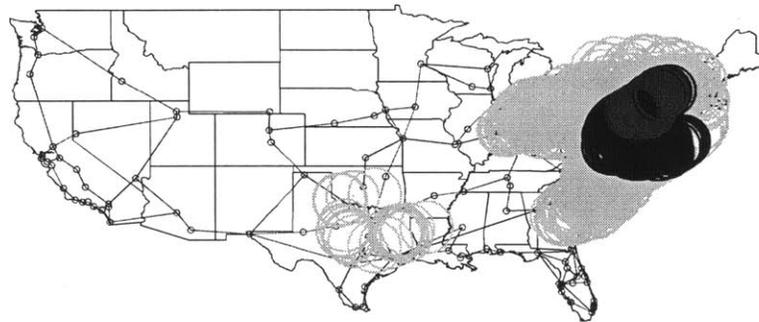
Figure 2-23: The impact of circular cuts of radius 2 on the $MFST$ between Fort Worth and NYC. Red circles represent cuts that result in $MFST = 0$, black circles result in $MFST = 1$, and yellow circles result in $MFST = 2$. Cuts which intersect the nodes representing Fort Worth or NYC are not shown.

monotonically increases with the cut radius. This implies that the minimum radius that guarantees a certain level of a specific performance measure (e.g., finding the radius of a circular cut that ensures that $AMF \leq 3$) can be found by using binary search along with the methods described in Section 2.7.

2.9 Conclusions

Motivated by applications in the area of network robustness and survivability, in this chapter, we focused on the problem of geographical network inhibition. Namely, we studied the properties and impact of geographical disasters that can be represented by either a line segment cut or a circular cut in the physical network graph. We considered a simple bipartite graph that abstracts the fiber links between the east
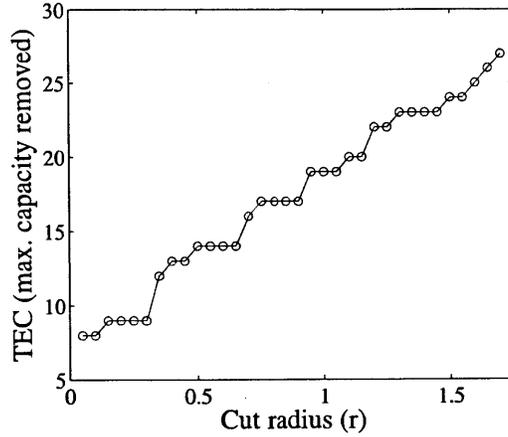
Figure 2-24: The maximum value of *TEC* as a function of the cut radius for a circular cut in the fiber plant illustrated in Fig. 2-1.

and west coasts in the U.S. or transatlantic/pacific links. Then, we considered a general graph model in which nodes are located on the Euclidian plane and studied two related problems in which cuts are modeled as line segments or as circular disks. For all cases, we developed polynomial-time algorithms for finding worst-case cuts. We then used the algorithms to obtain numerical results for various performance measures.

Our approach provides a fundamentally new way to look at network survivability under disasters or attacks that takes into account the geographical correlation between links. Some future research directions include the analytical consideration of arbitrarily shaped cuts and the use of computational geometric tools for the design of efficient algorithms.

Chapter 3

Non-Targeted Attacks

In chapter 2 we considered the problem of finding the worst-case location for a geographic failure (represented as a disk or line segment) in a geographic network with respect to certain network connectivity measures. On the other hand, in this chapter we consider the impact of a randomly located disaster (represented as a disk or line) on network connectivity.

3.1 Introduction

In this chapter we consider non-targeted geographic attacks; this can model failures resulting from natural disasters such as hurricanes or collateral (non-targeted) damage in an EMP attack. It can also model manufacturing errors in a VLSI ship or damage to a printed circuit board. This chapter first considers the effect of a random line cut, then it considers the effect of a random disk cut. Any links (modeled as line segments) that are intersected by a cut are removed from the network.

Our method is to use geometric probability to assign a measure to sets of lines and disks in the plane that intersect some set of line segments. Using these basic tools which are introduced and explained in sections 3.2 and 3.7, we are able to calculate network performance metrics to these random cuts in polynomial time. We also present numerical results that demonstrate the significance of geometry on network survivability.

To the best of our knowledge we are among the first to apply geometric probability techniques to network survivability. In [71] the survivability of undersea cables with respect to a randomly located disk is studied, however only a two node topology was considered. Also, [47] applied geometric probability techniques to detection in sensor networks.

A notable contribution of this chapter is the development of an algorithm to calculate the average two-terminal reliability of a network in polynomial time with respect to non-targeted line or disk failure model. This result is significant because calculating this metric assuming independent link failures in known to be NP-hard [10]. We also present some numerical results that demonstrate the significance of geometry on the survivability of the network and discuss network design in the context of random line and disk cuts.

This chapter is organized as follows: In sections 3.2 and 3.3 we introduce geometric probability with respect to lines and present an algorithm that allows us to evaluate joint link failure probabilities after a random line cut. In section 3.4 we use these results to demonstrate how to evaluate average two-terminal reliability to a random line cut (among other metrics). In section 3.5 we present some numerical results to show the significance of geometry on the survivability of the network and then in section 3.6 we present some network design problems in the context of random-line cuts. We then shift focus to random disk cuts. In sections 3.7 and 3.8 we present a method that allows us to approximate joint link failure probabilities under a random disk-cut. Then in section 3.9 we use these tools to show how to approximate certain metrics under this geographic failure model. In section 3.10 we again present some numerical results that make clear that geographically correlated failures are fundamentally different from independent failures. Then in section 3.11 we present some network design problems in the context of random-disk cuts. We conclude and discuss future research directions in section 3.12.

60

3.2 Modeling Random Line Cuts in Geographic Networks

In this section we describe how to model random disasters using geometric probability. For simplicity, we focus only on disasters which remove links along a random line. For example, damage to communication infrastructure can be a result of a natural disaster such as a tornado or collateral damage in an EMP attack. As these disasters are not targeted, these events can be modeled as breaking fiber along random line. After introducing some basic definitions from geometric probability, we review classic results which allow us to find single and pairwise link failure probabilities. These results are requisite for Section 3.3 where we show how to find joint link failure probabilities to random line-cuts.

3.2.1 Geometric Probability

Geometric probability is the study of probabilities involved in geometric problems. In our case, we are interested in the probability that a 'randomly' placed line in a plane will intersect a certain set of links (e.g., links whose removal would disconnect the network). It should be noted that the problem we are interested in is very similar in nature to the Buffon's Needle problem [68].

Before proceeding further, we will present some useful notation. Let C be a closed bounded convex set on the plane. Let L_C be the perimeter of C (where perimeter is the length of the boundary). Also, let $[C]$ denote the set of all lines in the plane which intersect C.

Geometric probability tells us how to assign a measure to sets of lines; let this measure be denoted by m. The rest of this section reviews results from geometric probability (see [50,65]) that are necessary for the development of this work. We first present some geometrical arguments to get intuition about what $m([C])$ should be for a given set C.

Let Z be a line in the plane, O be the origin, and H be the closest point on Z to
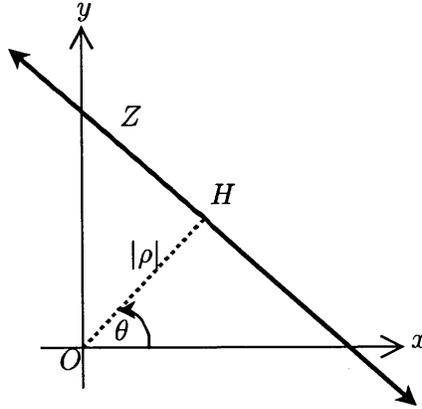
Figure 3-1: Let Z be a line in the plane, O be the origin, and H be the closest point on Z to O (see Fig. 3-1). Note that every line in the plane can be parameterized by ρ and θ where θ is the smallest non-negative angle between the x-axis and the line's normal and $|\rho|$ is the Euclidean distance between H and O such that $\rho \in \mathbb{R}$, $\theta \in [0, \pi)$, and the equation of the line is $x \cos\theta + y \sin\theta - \rho = 0$.
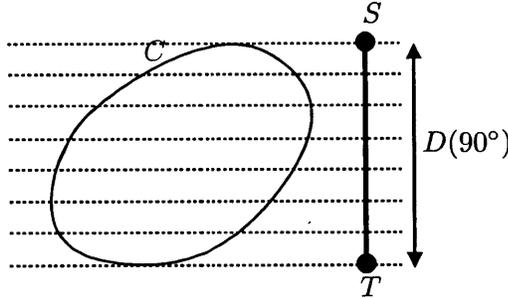
Figure 3-2: Consider a set of parallel lines ($\theta = 90°$ but variable ρ) that intersect C. Projecting C onto a line perpendicular to the parallel lines results in line segment $\overline{ST}$. Let $D(90°)$ be the length of this line segment. It seems reasonable for a measure to map the set of parallel lines in the plane with $\theta = 90°$ that intersect C to $D(90°)$. Note that $D(\theta)$ remains unchanged if C is translated.

O (see Fig. 3-1). Note that every line in the plane can be parameterized by ρ and θ where θ is the smallest non-negative angle between the x-axis and the line's normal and $|\rho|$ is the Euclidean distance between H and O such that $\rho \in \mathbb{R}$, $\theta \in [0, \pi)$, and the equation of the line is $x \cos\theta + y \sin\theta - \rho = 0$.

Let C be a bounded closed convex set in the plane. We first start by considering a set of parallel lines (fixed θ but variable ρ) that intersect C (see Fig. 3-2). Projecting C onto a line perpendicular to the parallel lines results in a line segment (see Fig. 3-2). Let $D(\theta)$ be the length of this line segment. It seems reasonable for a measure to map the set of parallel lines in the plane with angle θ that intersect C to $D(\theta)$. Note that $D(\theta)$ remains unchanged if C is translated.

Intuitively, by considering $D(\theta)$ over all angles, it seems reasonable to assign $m([C])$ as $\int_0^\pi D(\theta) d\theta$. Note that $D(\theta)$ is invariant under the translation of C and

62

since $D(\theta) = D(\theta + \pi)$ we know $\int_0^\pi D(\theta)d\theta$ is invariant under the translation and rotation of C.

Note, we have not yet properly defined the measure, we have only noted a property we would like it to have. For a proper definition of the measure see Appendix 3.A. We will now present two examples of evaluating $\int_0^\pi D(\theta)d\theta$. Consider a circle of radius 1. Since the projection of this circle onto any line is a line segment of length 2, we know that $D(\theta) = 2$ for all θ. So $\int_0^\pi D(\theta)d\theta = 2\pi$. Next consider a horizontal line segment length 1. By simple trigonometry, we know $D(\theta) = |\cos\theta|$ and thus $\int_0^\pi D(\theta)d\theta = \sin\theta|_0^{\frac{\pi}{2}} - \sin\theta|_{\frac{\pi}{2}}^\pi = 2$.

We now present an important result from geometric probability.

Lemma 11. *Let C be a bounded closed convex set and $D(\theta)$ be defined the same as above. Now,*

$$m([C]) = \int_{[C]} d\rho d\theta = \int_0^\pi D(\theta)d\theta = L_C$$

Note this is consistent with the above examples. See [50, 65] for a proof of the above statement and see Appendix 3.A for an intuitive argument.

3.2.2 Single Link Failures

Let $[Q]$ and $[C]$ be sets of lines in the plane such that $[Q] \subset [C]$. Given m, the probability a 'random' line is in the set $[Q]$ when it is known to be in the set $[C]$ is defined to be ratio of measures [65], $\frac{m([Q])}{m([C])}$. This definition appeals to intuition; $m([C])$ in some sense represents the 'weight' of lines in $[C]$ and $m([Q])$ represents the 'weight' of lines in $[Q]$. Therefore it makes sense that the probability a line in $[C]$ is also in $[Q]$ is $\frac{m([Q])}{m([C])}$.

We now present an example relating to network survivability. Consider a rectangle C with height a and width b and a line segment Q of length l inside C (see Fig. 3-3). Now we consider a random line-cut. We have:

$$\Pr(Q \text{ cut}|C \text{ cut}) = \frac{m([Q])}{m([C])} = \frac{L_Q}{L_C} = \frac{l}{a+b}$$
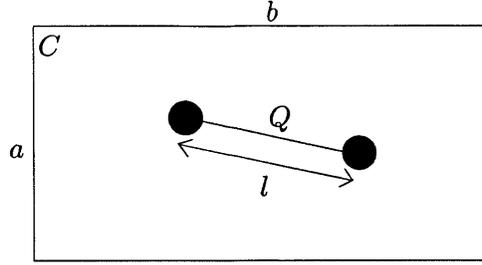
Figure 3-3: Rectangle C with link Q inside.
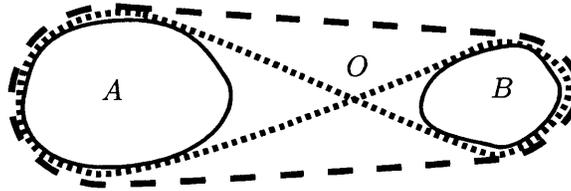
Figure 3-4: The dotted curve shows the internal cover of A and B, denoted by $I(A, B)$. The dashed curve shows the boundary of the convex hull of $A \cup B$.

3.2.3 Pairwise Link Failures

We now present a classic definition and result in geometric probability which allows us to find pairwise link failure probabilities with respect to a random line-cut.

Definition 3 (Internal Cover). *The internal cover of two bounded convex sets in the plane, A and B, denoted by $I(A, B)$ is given by the following. If $A \bigcap B = \emptyset$ then the internal cover is realized by a closed elastic string drawn about A and B and crossing over a point O placed between A and B [65] (see Fig. 3-4). If $A \bigcap B \neq \emptyset$, then the internal cover is realized by a string which is wrapped around the entire perimeter of both A and B. In this case, $L_{I(A,B)} = L_A + L_B$.*

Let $conv(A)$ denote the convex hull of set A.

Lemma 12 ([65]). *If A and B are bounded closed convex sets,*

$$m([A] \bigcap [B]) = L_{I(A,B)} - L_{conv(A \cup B)}$$
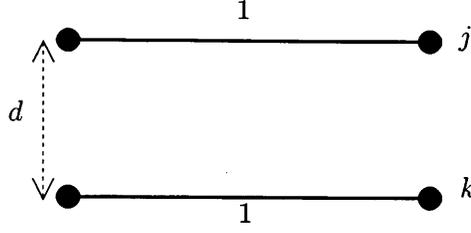
Figure 3-5: $Pr(\ j \text{ and } k \text{ cut } |\ j \text{ cut }) = \sqrt{1+d^2} - d$
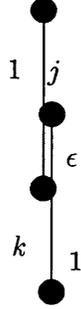
Figure 3-6: $Pr(\ j \text{ and } k \text{ cut } |\ j \text{ cut }) = \epsilon$

For a proof, see Appendix 3.B.

Given two links, j and k, by definition the probability a 'random' line is in the set $[j] \cap [k]$ when it known to be in the set $[j]$ (note that $[j] \cap [k] \subset [j]$) is the ratio of measures $\frac{m([j] \cap [k])}{m([j])}$. So, using Lemma 12, we find

$$\Pr(k \text{ cut}|j \text{ cut}) = \frac{m([k] \bigcap [j])}{m([j])} = \frac{L_{I(j,k)} - L_{conv(j \cup k)}}{L_j}$$

Examples demonstrating of the above result for pairwise link failures are given below.

Example 1: Two parallel links, j and k, of length 1 are separated by a distance d. The nodes form corners of a rectangle. See Fig. 3-5. Since the length of a diagonal is given by $\sqrt{1+d^2}$, we know $L_{I(j,k)} = 2 + 2\sqrt{1+d^2}$. Also, the perimeter of the rectangle is given by $2 + 2d$. Therefore, $Pr(\ j \text{ and } k \text{ cut } |\ j \text{ cut }) = \sqrt{1+d^2} - d$.

Example 2: Two links, j and k, of length 1 overlap as shown in Fig. 3-6 where the length of the overlap is ϵ. These links intersect, so $L_{I(j,k)} = 4$ by definition. Also,

65
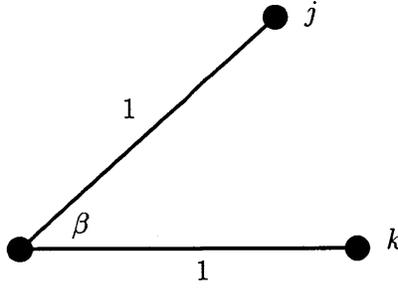
Figure 3-7: $Pr(\ j$ and k cut $\mid j$ cut $) = 1 - \sin\frac{\beta}{2}$

it is evident $L_{conv(j\cup k)} = 2(2 - \epsilon)$. Therefore, $Pr(\ j$ and k cut $\mid j$ cut $) = \epsilon$.

Example 3: Two links, j and k, of length 1 are at an angle β to each other and share a common node. See Fig. 3-7. These links intersect, so $L_{I(j,k)} = 4$. Also, the perimeter of the convex hull is given by $2 + 2\sin\frac{\beta}{2}$. Therefore, $Pr(\ j$ and k cut $\mid j$ cut $) = 1 - \sin\frac{\beta}{2}$. This result agrees with our intuition. If $\beta = 0$, then the links are on top of each other and the probability is one. If $\beta = \pi$, then only lines which intersect the shared node intersect both links and the probability is zero.

3.3 Geographically Correlated Link Failures Under a Random Line Cut

In this section we present an algorithm which calculates the measure of lines that intersect every line segment in a set of segments. This result will allow us to calculate the probability that a random line-cut intersects a certain set of links in a network (e.g. links whose removal would disconnect the network). We will then use this to efficiently calculate network performance measures with respect to random line-cuts. The details of this section may be skipped and the reader may proceed without loss of continuity to Section 3.4 on evaluating network reliability.

Assume we are given a set of line segments, Q, on a plane such that the endpoints are in general form; that is, no three endpoints are collinear[1]. Let the i^{th} line segment

---

[1]This assumption is not restrictive as we can slightly perturb the location of the endpoints to satisfy this condition.

66
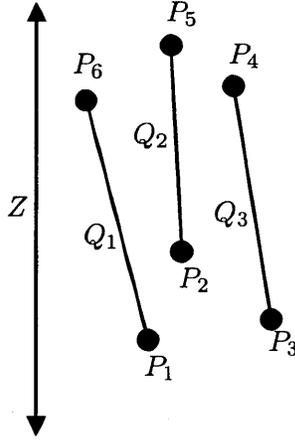
Figure 3-8: Shown above are three line segments, Q_1, Q_2, and Q_3, and a line Z which does not intersect the convex hull of P. We want to find $m([Q_1] \bigcap [Q_2] \bigcap [Q_3])$ which is equivalent to $m(Z \bigcap [Q_1] \bigcap [Q_2] \bigcap [Q_3])$.

be denoted by Q_i. Our goal is to find $m(\bigcap_{i=1}^{|Q|} [Q_i])$; that is, the measure of the set of lines that intersect all segments in Q.

Sylvester in [68] shows how to solve for $m(\bigcap_{i=1}^{|Q|} [Q_i])$. However, this approach takes exponential time in $|Q|$; this is because the perimeter of the convex hull of every subset of Q must be considered. Ambartzumian in [4] and [5] provides an algorithm to calculate $m(\bigcap_{i=1}^{|Q|} [Q_i])$ in polynomial time. The algorithm in this section presents an alternate way to solve for $m(\bigcap_{i=1}^{|Q|} [Q_i])$ in polynomial time by reducing the problem to finding pairwise link failures, as was done in Section 3.2.3.

In the following, for clarity of presentation we break down our procedure into steps for finding $m(\bigcap_{i=1}^{|Q|} [Q_i])$.

Step 1:

Let P be the set of endpoints of the line segments in Q. Let us impose an arbitrary ordering on P and denote the i^{th} point in P by P_i. Let $\overline{P_i P_j}$ be the line segment between P_i and P_j.

We start by arbitrarily placing a vertical line Z such that it does not intersect the convex hull of P (see Fig. 3-8). Note that $m(\bigcap_{i=1}^{|Q|} [Q_i]) = m((\bigcap_{i=1}^{|Q|} [Q_i]) \bigcap [Z])$ because the set of all lines which do not intersect Z has measure zero.

Step 2:

Consider all lines that intersect two points in P. Let the intersection points of

67
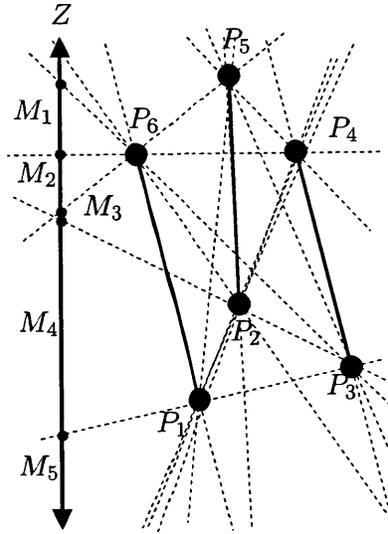
Figure 3-9: Consider all lines that intersect two points in P. Let the intersection points of these lines and Z be denoted by α (shown as dots on Z above). Let the divisions of of Z into line segments and rays by points in α be denoted by M. Let us impose an ordering on M and denote M_i to be the i^{th} segment in M.

these lines and Z be denoted by α. Let the divisions of Z into line segments and rays by points in α be denoted by M (see Fig. 3-9). Let us impose an ordering on M and denote M_i to be the i^{th} segment in M.

Step 3:

Now,

$$
\begin{aligned}
m(\bigcap_{i=1}^{|Q|}[Q_i]) &= m((\bigcap_{i=1}^{|Q|}[Q_i])\bigcap[Z]) \\
&= m((\bigcap_{i=1}^{|Q|}[Q_i])\bigcap(\bigcup_{M_j \in M}[M_j])) \\
&= m\Big(\bigcup_{M_j \in M}((\bigcap_{i=1}^{|Q|}[Q_i])\bigcap[M_j])\Big)
\end{aligned}
$$

Since every $[M_j]$ is disjoint from $[M_k]$ when $j \neq k$ up to measure zero, we have:

$$
m\Big(\bigcup_{M_j \in M}((\bigcap_{i=1}^{|Q|}[Q_i])\bigcap[M_j])\Big) = \sum_{M_j \in M}m((\bigcap_{i=1}^{|Q|}[Q_i])\bigcap[M_j])
$$

68

Our problem is now reduced to computing $m\big((\bigcap_{i=1}^{|Q|}[Q_i])\bigcap[M_j]\big)$ for every j. That is, the measure of the set of lines that intersect both M_j and each of the segments in Q. We will show that computing this is easy because it is equivalent to computing $m([\overline{P_k P_l}]\bigcap[M_j])$ for some k and l. That is, $m\big((\bigcap_{i=1}^{|Q|}[Q_i])\bigcap[M_j]\big)$ is the same as the measure of the set of lines intersecting M_j and a line segment connecting two points in P.

Step 4:

In the following steps, we assume X is a point on Z such that $X \notin \alpha$.

Definition 4 $(T(X))$. $T(X)$ *is an* ordered *set of all points in P such that when Z is rotated counter-clockwise about X, the order in which points in P are intersected is the ordering in $T(X)$.*

For an illustration of this definition see Fig. 3-11.

Lemma 13. $T(X) = T(X')$ *for every* $X \in M_j$ *and* $X' \in M_j$.

Intuitively, this lemma states that the ordering of $T(X)$ is the same for all X in M_j.

Proof. We want to show the ordering of $T(X)$ is constant for all X in any M_j. This is equivalent to showing the pairwise ordering in $T(X)$ is constant for all X in any M_j.

We will use geometry techniques to prove the pairwise ordering in $T(X)$ is constant. In order to do this we will present some notation. Consider two different points in P, P_1 and P_2. Assume the line that contains P_1 and P_2 intersects Z (the proof is trivial otherwise) and denote this intersection point by W. Let the distance between W and P_1 be given by d_1 and the distance between W and P_2 be given by d_2. Without loss of generality, assume $d_1 < d_2$. Let X be a point on Z such that $X \notin W$. Let x denote the distance between X and W. Let θ_1 be the angle Z must rotate counter-clockwise about X to intersect P_1 and θ_2 be the angle to intersect P_2. Realizing arctan is a strictly monotonically increasing function, from geometry we
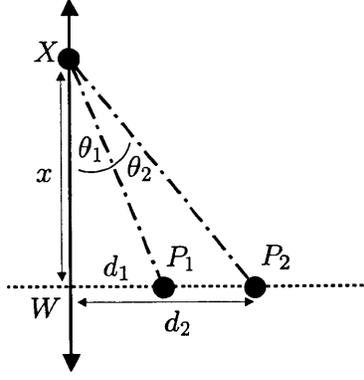
Figure 3-10: This is the setup for the proof of Lemma 13. Here X is above W, so $\theta_1 = \arctan(\frac{d_1}{x})$ and $\theta_2 = \arctan(\frac{d_2}{x})$.

have:

$$\theta_1 = \arctan(\frac{d_1}{x}) < \arctan(\frac{d_2}{x}) = \theta_2 \ \forall \ X \text{ above W}$$

$$\theta_1 = \pi - \arctan(\frac{d_1}{x}) > \pi - \arctan(\frac{d_2}{x}) = \theta_2 \ \forall \ X \text{ below W}$$

See figure 3-10 for clarification.

Now let $X' \in M_j$ for some j such that $X' \notin W$. Now, the equations above imply: P_1 comes before P_2 in both $T(X)$ and $T(X')$ or P_1 comes after P_2 in both $T(X)$ and $T(X')$. Because the above holds for any two different P_1 and P_2, this completely specifies the ordering of $T(X)$ and $T(X')$ and also implies $T(X) = T(X')$. $\square$

Step 5:

Definition 5 (A_1^X). A_1^X is the last point in $T(X)$ such that there does not exist P_k and P_l ahead of A_1^X where $\overline{P_k P_l} \in Q$.

Definition 6 (A_2^X). A_2^X is the first point in $T(X)$ such that there exists a P_k before A_2^X where $\overline{P_k A_2^X} \in Q$.

See Fig. 3-11 for an example.

Lemma 14. *If A_1^X comes before A_2^X, then $[X] \bigcap (\bigcap_{i=1}^{|Q|}[Q_i]) = [X] \bigcap [\overline{A_1^X A_2^X}]$, otherwise if A_2^X comes before A_1^X, then $[X] \bigcap (\bigcap_{i=1}^{|Q|}[Q_i]) = \emptyset$.*

Intuitively, this lemma says the set all lines which intersect X and every line segment in Q is the same as the set of all lines which intersect X and some $\overline{P_k P_l}$.
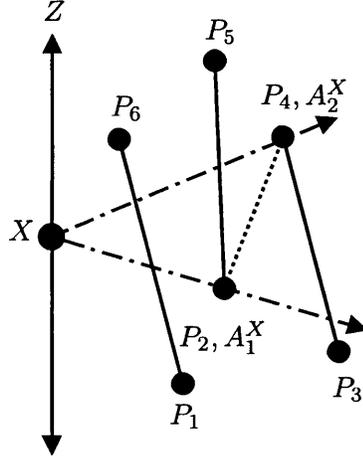
Figure 3-11: An example to demonstrate the definitions of $T(X)$, A_1^X, and A_2^X. Here $T(X)$ is the ordered set $\{P_1, P_3, P_2, P_4, P_5, P_6\}$. A_1^X is the last point in $T(X)$ such that there does not exist a P_i and P_j ahead of A_1^X where $\overline{P_iP_j} \in Q$. Here $A_1^X = P_2$. A_2^X is the first point in $T(X)$ such that there exists a P_i before A_2^X where $\overline{P_iA_2^X} \in Q$. Here $A_2^X = P_4$. Note how the dotted line segment $\overline{A_1^X A_2^X}$ has the property $[X] \cap (\cap_{i=1}^{|Q|}[Q_i]) = [X] \cap [\overline{A_1^X A_2^X}]$.

Take Fig. 3-11 as an example. The set of lines that intersect X and all three line segments is equivalent to the set of lines that intersect X and $\overline{P_2 P_4}$.

Proof. We first use the definitions of A_1^X and A_2^X to find the angles of lines which intersect X and every $Q_i \in Q$. Then conditioned on the ordering of A_1^X and A_2^X in $T(X)$, we use this set of angles to prove the lemma.

We first introduce some useful terminology. Let V_θ denote the line which is a counter-clockwise rotation of Z about X by θ degrees. Let V_{θ_1} denote the line which intersects both A_1^X and X and let V_{θ_2} denote the line which intersects both A_2^X and X. The definition of A_1^X implies θ_1 is the smallest θ such that for every $Q_i \in Q$ there exists a V_θ with $\theta \leq \theta_1$ such that Q_i is intersected. Intuitively, θ_1 is the smallest angle θ such that V_θ can intersect every $Q_i \in Q$. The definition of A_2^X implies θ_2 is the largest θ such that for every $Q_i \in Q$ there exists a V_θ with $\theta \geq \theta_2$ such that Q_i is intersected. Intuitively, θ_2 is the largest angle θ such that V_θ can intersect all $Q_i \in Q$. Since θ_1 is the smallest angle θ such that V_θ can intersect every $Q_i \in Q$ and θ_2 is the largest angle θ such that V_θ can intersect all $Q_i \in Q$, this implies V_θ intersects every $Q_i \in Q$ iff $\theta_1 \leq \theta \leq \theta_2$.

If we assume A_1^X comes before A_2^X in $T(X)$, this implies $\theta_1 \leq \theta_2$. Note from geometry, we know a line subsects angle $\angle A_1^X X A_2^X$ iff this line intersects X and

$\overline{A_1^X A_2^X}$; thus $\theta_1 \leq \theta \leq \theta_2$ iff V_θ intersects $\overline{A_1^X A_2^X}$. Since V_θ intersects every $Q_i \in Q$ iff $\theta_1 \leq \theta \leq \theta_2$, this implies $[X] \bigcap (\bigcap_{i=1}^{|Q|} [Q_i]) = [X] \bigcap \overline{[A_1^X A_2^X]}$.

If we assume A_2^X comes before A_1^X in $T(X)$, this implies $\theta_2 \leq \theta_1$. Since V_θ intersects every $Q_i \in Q$ iff $\theta_1 \leq \theta \leq \theta_2$, this implies if $\theta_2 < \theta_1$ no line intersects X and every $Q_i \in Q$. Also, if $\theta_2 = \theta_1$ only the line V_{θ_1} intersects X and every $Q_i \in Q$. In this case, since V_{θ_1} intersects A_1^X and A_2^X, V_{θ_1} is equivalent to $B_{A_1^X, A_2^X}$. $\qquad\square$

Lemma 15. *Assume $X \in M_j$. If A_1^X comes before A_2^X in $T(X)$, then $m\big([M_j] \bigcap (\bigcap_{i=1}^{|Q|} [Q_i])\big) = m([M_j] \bigcap \overline{[A_1^X A_2^X]})$, otherwise $m([M_j] \bigcap (\bigcap_{i=1}^{|Q|} [Q_i])) = 0$.*

Intuitively, this lemma says the measure of all lines which intersect M_j and all line segments in Q is the same as the measure of all lines which intersect M_j and some line segment connecting two points in P. *This reduces the problem to finding the measure of the set of lines intersecting two line segments, a problem which we already know how to solve* (see Lemma 12).

Proof. Direct result of Lemmas 13 and 14 and the fact $m([\alpha_i]) \doteq 0 \ \forall \alpha_i \in \alpha$. $\qquad\square$

Summary:

In steps 1 and 2 we place a vertical line Z and partition it into a set of line segments M. In step 3 we show $m(\bigcap_{i=1}^{|Q|} [Q_i]) = \sum_{M_j \in M} m((\bigcap_{i=1}^{|Q|} [Q_i]) \bigcap [M_j])$. Steps 4 and 5 when when combined with a lemma about rays (see appendix) show how to compute $m((\bigcap_{i=1}^{|Q|} [Q_i]) \bigcap [M_j])$ in constant time assuming we know A_1^X and A_2^X. For a given $X \in M_j$, $T(X)$ can be computed in polynomial time by sorting the angles between $\overline{XP_i}$ and Z for all i. A_1^X and A_2^X can then be found by enumerating through $T(X)$. Since $|M|$ is polynomial, this allows us to calculate $m(\bigcap_{i=1}^{|Q|} [Q_i])$ in polynomial time.

The complexity of this algorithm can be reduced by going through all M_j 'in order,' thus eliminating the need to sort P for all M_j in M.

3.4 Evaluating Network Reliability Under A Random Line

In this section we introduce and show how to evaluate some performance metrics with respect to a random line-cut. After introducing our network model, we show that every performance metric can be evaluated in polynomial time. In particular, we can evaluate average k-terminal reliability in polynomial time under 'random' line-cuts. This is in contrast to the case of independent link failures for which there exists no known polynomial time algorithm to calculate this reliability metric.

3.4.1 Network Model

We start by describing our network model. Our geometric graph model contains a set of nodes N where each node is represented by a point on the plane. We assume the node locations are in general form; that is no three are collinear. A link between two nodes is represented by a line segment with endpoints at the respective node locations. In order to assign probabilities to random line events, we assume the set which contains all nodes and links ($conv(N)$) is a subset of some bounded closed convex set C with perimeter L_C. If a 'random' line that intersects C also intersects some links, those links are disrupted. Our goal is to evaluate the performance metrics described below in Definition 7 after a single random line-cut that intersects C.

3.4.2 Performance Metrics

We first introduce some network performance metrics and then describe how to evaluate each one after the removal of the intersected links. We will use the tools developed in Section 3.3 to evaluate average values of these metrics with respect to a random line-cut.

Definition 7. *[Performance Metrics]*

- ATR - *The all terminal reliability of the network. The all terminal reliability is*

defined here as 1 if the network is connected and 0 otherwise. In order to verify connectivity of the network a Breadth First Search algorithm can be used.

- ATTR - *The average two terminal reliability of the network over all pairs of nodes. The two terminal reliability between two nodes is defined here as 1, if there is a path between them and 0, otherwise [61]. Effectively this metric is the probability a randomly chosen pair of nodes is connected. If the network is fully connected, the value of ATTR is 1. Otherwise, we have to sum over the number of node pairs in every connected component and divide it by the total number of node pairs in the network. That is, we sum the value of $k(k-1)$ over every connected component, where k is the number of nodes in each of the components, and then divide this sum by $N(N-1)$. This ratio gives the fraction of node pairs that are connected to each other. In order to verify connectivity or to count the number of nodes in each connected component a Breadth First Search algorithm can be used.*

In this chapter we only discuss the above two metrics. However, the following relevant metrics can also be evaluated with respect to random line-cuts using the results of this section.

- TC - *The total capacity of the intersected links.*

- MFST - *The maximum flow between a given pair of nodes s and t.*

- AMF - *The average value of maximum flow between all pairs of nodes.*

It is apparent from the descriptions above that evaluating each metric after the removal of intersected links takes polynomial time in $|N|$.

3.4.3   Evaluation of the Metrics

We now show how to evaluate the metrics in Definition 7 with respect to a random-line cut. The basic idea is that every line which separates the nodes in the same way removes the same set of links. Using the techniques in Section 3.3, we calculate
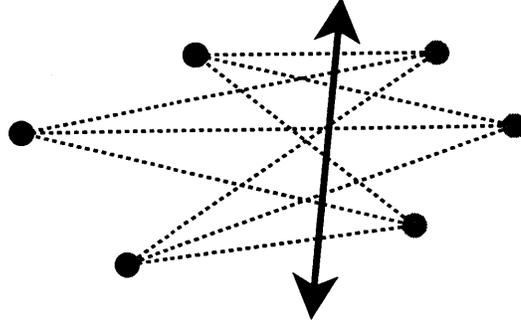
Figure 3-12: Consider a line-partition of a set of nodes, N, into two non-empty subsets in the figure above. One subset has nodes colored grey and the other has nodes colored black. A line separates N into these subsets iff it intersects every dashed line segment connecting a grey node and black node.

the measure of the set of lines that separate the nodes in this way; this allows us to calculate the weighted average of a metric over all possible cuts. We start by introducing some useful terminology.

Definition 8 (Line-partition). *A line-partition is a partition of a set of nodes into two subsets which are separated by a line. It is important to notice that not all partitions of N are line-partitions.*

Let P be the set of all line-partitions created by lines that intersects $conv(N)$. For each line-partition p in P, let $[p]$ be the set of all lines which form the line-partition p. For a particular p, let the set of all line segments connecting a node in one subset to a node in the other subset be given by Q_p.

Lemma 16. $m([p]) = m(\cap_{q \in Q_p}[q])$ *for every $p \in P$*

Proof. If a line intersects every line segment in Q_p, then it separates the endpoints of the line segments in Q_p into subsets that form p or it intersects a node. On the other hand, if a line forms a line partition p, then it separates nodes into two subsets and thus will intersect every line segment that has endpoints in both subsets (this is precisely Q_p). See Fig. 3-12. Thus $[p] = \cap_{q \in Q_p}[q]$ except for a set of lines which intersect nodes. Since the set of lines which intersect nodes has zero measure (points have zero perimeter), the result follows. $\square$

Now, let $[r]$ be the set of lines that intersect C but not $conv(N)$. That is, $[r] = [C]\setminus[conv(N)]$. Thus, $m([r]) = m([C]) - m([conv(N)])$ by countable additivity of

measures.

Note that $(\cup_{p \in P}[p]) \cup [r] = [C]$ up to a set of measure zero. Now, since every line which forms the same line-partition removes the same links, evaluating the performance measures to a random line-cut becomes a weighted average over each partition. Let $Y(p)$ be the performance metric on the network when links that intersect a line in $[p]$ are removed. Since $\frac{m([p])}{L_C}$ is the probability a random line-cut will create a partition p, the performance metric to a random line-cut can be expressed as

$$\frac{m([r])}{L_C}Y(r) + \sum_{p \in P} \frac{m([p])}{L_C}Y(p) =$$
$$\frac{L_C - L_{conv(N)}}{L_C}Y(r) + \sum_{p \in P} \frac{m(\bigcap_{q \in Q_p}[q])}{L_C}Y(p) \qquad (3.1)$$

Section 3.3 shows how to calculate $m(\bigcap_{q \in Q_p}[q])$ in polynomial time. The performance metrics in Definition 7 can be calculated in polynomial time as discussed above. In the following, we will show that $|P|$ is $O(|N|^2)$.

Lemma 17 ([37]). *There are $O(|N|^2)$ line-partitions of a set of $|N|$ nodes.*

Proof. [37] shows there are $\binom{|N|}{2} + 1$ line-partitions of a set of $|N|$ points, no three of which are collinear. $\qquad \square$

We will now provide some intuition behind the above result. Consider a line that forms a line-partition in which neither subset of nodes is empty (the line intersects $conv(N)$). Now rotate this line clockwise until nodes prevent any further clockwise movement (imagine that the line cannot pass through the nodes). There will be two points stopping the line from moving any further. Now, these two points specify this partition, and since there are $\binom{|N|}{2}$ ways to pick two nodes, there are $\binom{|N|}{2}$ partitions (see Fig. 3-13). The additional partition comes from the case when the line does not intersect $conv(N)$.

Theorem 4. *Evaluating any performance metric in Definition 7 with respect to a random line cut takes polynomial time in N.*

76
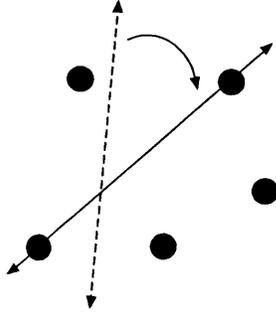
Figure 3-13: Consider a line-cut resulting in a line-partition in which neither subset of nodes is empty. Now rotate this line-cut clockwise until nodes prevent any further clockwise movement (imagine that the line cannot pass through the nodes).

Proof. Since $|P|$ is polynomial in $|N|$ and evaluating $m(\bigcap_{q \in Q_p}[q])$ and $Y(p)$ takes polynomial time, Equation 3.1 can be evaluated for any performance metric in polynomial time. $\qquad\square$

This is particularly interesting for the case of $ATTR$ because there is no known polynomial algorithm to find $ATTR$ assuming independent link failures. This is a consequence of the fact not all partitions of N are line-partitions.

3.5 Numerical Results to Random Line-cuts

We first present an example that demonstrates the significance of geometry on the survivability of the network. We then find $ATTR$ of a real-world network to a random line-cut.

3.5.1 An Example to Demonstrate the Importance of Geometry

In this example, every link has a length of one, so every link is intersected by a random line-cut with equal probability. We consider different geometries of the same network and evaluate ATR to random line-cuts. For comparison, we also evaluate ATR assuming independent links failures.

Consider a network of $|N|$ nodes connected in serial by line segments of length 1. We consider two different cases of geometries for this network. In case (i) the

77
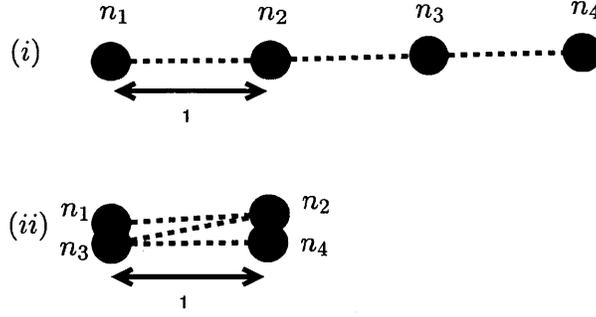
Figure 3-14: Two networks of 4 nodes connected in serial by line segments of length 1. The network in case (i) resembles a line segment of length 3 and the network in case (ii) resembles a line segment of length 1.
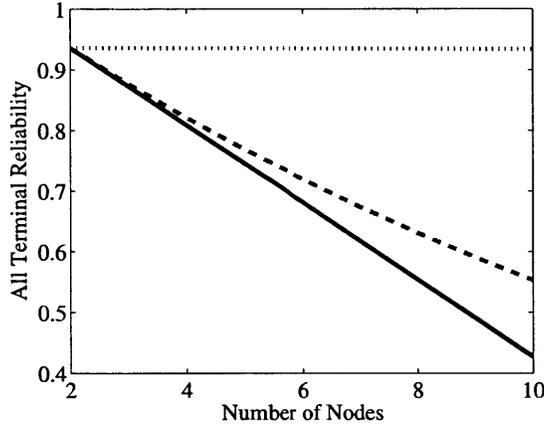
Figure 3-15: Assuming $L_C = 10\pi$, the figure shows ATR versus number of nodes for different network configurations. The dotted line represents the network which resembles a line segment of length 1, the solid line represents the network which resembles a line segment of length $|N| - 1$, and the dashed line represents independent link failures.

network resembles a line segment of length $|N| - 1$, and in case (ii) the network resembles a line segment of length 1 (see Fig. 3-14). Assuming $L_C = 10\pi$ and letting $|N|$ vary (assuming $|N| \leq 10$), we calculate ATR to random line-cuts in both cases using methods described in Section 3.4. Also, since any particular link of length 1 fails with probability $\frac{2}{10\pi}$ with respect to a random line-cut, we evaluate ATR when links fail independently with probability $\frac{2}{10\pi}$.

Fig. 3-15 shows the results. In case (i), ATR is approximately $1 - \frac{2(|N|-1)}{10\pi}$ since ATR is 1 if any link is intersected and 0 otherwise. In case (ii), ATR is approximately $1 - \frac{2}{10\pi}$ for all $|N|$ (again, since ATR is 1 if any link is intersected and 0 otherwise). When links fail independently with probability $\frac{2}{10\pi}$, ATR is given by $Pr(\text{no links fail}) = (1 - \frac{2}{10\pi})^{|N|-1}$. Note this value lies between the results for the two geometric networks.

78
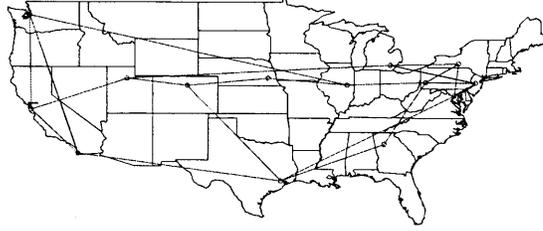
Figure 3-16: This figure shows NSFNET from 1991 [51].
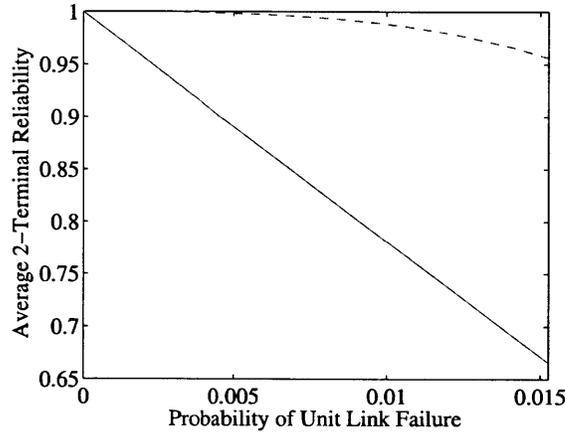
Figure 3-17: The solid line shows $ATTR$ in NSFNET versus the probability a unit (latitude/longitude) of fiber is cut by a random line. The dashed line shows $ATTR$ in NSFNET assuming links fail *independently* such that links fail with the same probability as in the random line-cut case.

These results agree with intuition. In case (i) when the network resembles a line segment of length $|N| - 1$, the network is spread out and the probability of any link being intersected is larger when $|N|$ is larger. In case (ii) when the network resembles a line segment of length 1, the network has a perimeter of 2 and thus the probability of any link being intersected is small. This example highlights the importance of node location on the survivability of the network (see Fig. 3-15).

3.5.2   A Real-World Example

In this example we consider NSFNET as found in 1991 [51] (see Fig. 3-16). This network has 14 nodes and connects major universities across the U.S.. We assume the network is contained within a convex set with perimeter L_C.

Using the results of Section 3.4, we calculate $ATTR$ of NSFNET to random-line cuts while L_C varies. Given the length of a particular link, the probability that link is

79

cut is proportional to $\frac{2}{L_C}$. So we can plot $ATTR$ versus $\frac{2}{L_C}$. Note $\frac{2}{L_C}$ is the probability a unit of fiber is cut (since a unit of fiber has a perimeter of 2). See Fig. 3-17 for results. Note the linear form of the result in the figure; this agrees with Equation 3.1.

Next, we calculated $ATTR$ of NSFNET assuming independent link failures such that links fail with the same probability as in the random line-cut case. Thus the probability a link fails is still proportional to its length, however links fail independently. Since the total number of links is small, calculating $ATTR$ by enumerating all possible failures is still feasible (possible failures are exponential in number of links). See Fig. 3-17 for results. Note that $ATTR$ in the independent failure model is greater than in the case of random line-cuts. Perhaps this is because at least two links must fail independently to disconnect the network, however a line which intersects the network is guaranteed to disconnect it. Since most backbone networks are likely to be well connected, we expect a random line-cut to lead to lower $ATTR$ than independent link failures in the real-world setting.

3.6 Network Design Under Random Line-cuts

In this section we present some network design problems in the context of random line-cuts. In all the proposed problems the location of every node is fixed; the problem is to find a set of links most robust to some metric under some constraints.

(i) Let N be a set of nodes fixed on the plane. As before, assume all links are represented as line segments between the points. A reasonable goal is to design a connected network with the least expected number of links cut by a random line. By linearity of expectation, the expected number of links cut is proportional to $\sum_{i \in Q} l_i$ where Q is the set of links and l_i is the length of link i. So, this problem reduces to minimizing the total length of links in the network while ensuring the graph is connected. This is equivalent to finding a Euclidean minimum spanning tree of N which can be done in polynomial time. Note however that the resulting network is not robust because a single link failure will disconnect it.

(ii) We next consider ring networks. Let N be a set of nodes fixed on the plane.

A reasonable goal is to design a connected ring network with the least expected number of links cut by a random line. As before, the expected number of links cut is proportional to $\sum_{i \in Q} l_i$ where Q is the set of links and l_i is the length of link i. So this problem reduces to finding the minimum length cycle that visits every node exactly once and returns to the starting node. This is equivalent to the Euclidean traveling salesman problem which is hard to compute [43].

(iii) The final problem considers how to connect two nodes such that the path between them is robust to random-line disasters. More precisely, let N be a set of nodes fixed on the plane. Let $S \in N$ and $T \in N$ such that $S \neq T$. Let Q be a set of links. The problem is to find a path from S to T consisting of links from Q that has the minimum probability of being cut. This may correspond to finding the most robust path between two cities along preexisting conduits. Since the probability a random line will intersect the path is proportional to the perimeter of the convex hull of the path (a line intersects the path iff it intersects the convex hull of the path), we want to find a path such that the perimeter of the convex hull of the path is minimized. That is, we want to find a path with minimum perimeter convex hull such that the path starts at S and ends at T, its edges belong to Q, and it contains no repeated vertices. The authors do not know a polynomial time algorithm to solve this problem (except in a trivialized setting).

3.7 Modeling Random Circular Cuts

We now shift focus from random lines to randomly located disks of a particular radius. The circular form of the attack model may better model the effect of storms or large bombs. After introducing some basic definitions from geometric probability, we review classical results which allow us to find single link failure probabilities. These results are requisite for Section 3.8 where we show how to find joint link failure probabilities to random disk-cuts.
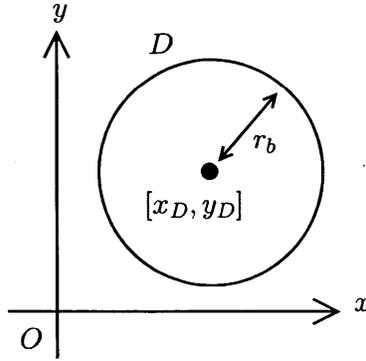
Figure 3-18: Every disk in the plane of radius r_b can be parameterized by the location of its center. Denote the center of disk D as $[x_D, y_D]$.

3.7.1 Geometric Probability

Geometric probability is the study of probabilities involved in geometric problems. In our case, we are interested in the probability that a 'randomly' placed disk (of a particular radius) in the plane will only intersect a certain set of links (e.g., links whose removal would disconnect the network). We model a disaster event in the network as a *single* randomly located disk of a radius r_b.

Before proceeding further, we will present some useful notation. Denote the perimeter of a set of points in the plane C by L_C and its area by R_C. Given a set in the plane, let $\langle \cdot \rangle$ denote the set of all disks in the plane of radius r_b that intersect it.

Geometric probability tells us how to assign a measure to sets of disks; let this measure be denoted by μ. The rest of this section reviews results from geometric probability (see [50, 65]) that are necessary for the development of this work.

Note that every disk in the plane of radius r_b can be parameterized by the location of its center. Denote the center of disk D as $[x_D, y_D]$ (see Fig. 3-18). Let D_O be the disk of radius r_b centered at the origin.

We now present the definition of the measure μ.

Definition 9 (Measure of a set of disks). *The measure μ of a set of disks G is defined as the integral*

$$\mu(G) = \int_G dx dy$$

Note we use G to denote both a set of disks and the set of centers of these disks.
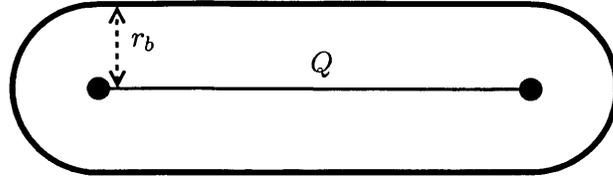
Figure 3-19: The dotted shape above represents $Q \oplus D_O$. This shape is known as a hippodrome and it represents the set of all points whose shortest distance to Q is less than or equal to r_b [32]. Denote the hippodrome corresponding to a link Q and radius r_b by $H(Q, r_b)$. Note that a disk D of radius r_b intersects Q iff $[x_D, y_D] \in H(Q, r_b)$.

This integral is the area of G in the (x, y) plane and will be denoted by $area(G)$. This definition appeals to intuition; in the same way the 'size' of a set of points in the plane is its area, the 'size' of a set of disks is the area of the disks centers.

Definition 10 (Minkowski Sum). *The Minkowski sum of two sets in the plane A and B in Euclidean space, denoted by $A \oplus B$, is given by*

$$A \oplus B = \{a + b | a \in A, b \in B\}$$

Intuitively, every point in the Minkowski sum $C \oplus D_O$ represents a center of a disk of radius r_b that intersects C. We will now discuss an important example. Let Q be a line segment link; consider $Q \oplus D_O$ (see Fig. 3-19). This shape is known as a hippodrome and it represents the set of all points whose distance to Q is less than or equal to r_b [32]. Denote the hippodrome corresponding to a link Q and radius r_b by $H(Q, r_b)$. Note that a disk D of radius r_b intersects Q iff $[x_D, y_D] \in H(Q, r_b)$.

Lemma 18 ([65]). *Let C be a bounded closed convex set of points in the plane, then*

$$\mu(\langle C \rangle) = \int_{\langle C \rangle} dxdy = area(C \oplus D_O) = R_C + L_C r_b + \pi r_b{}^2$$

Intuitively, every point in the Minkowski sum $C \oplus D_O$ represents a center of a unique disk of radius r_b that intersects C. Integrating over the set of centers of these disks yields the measure of $\langle C \rangle$. For example, consider a line segment link Q of length d. Now the measure of the set of disks of radius r_b that intersect Q is

$$\mu(\langle Q \rangle) = area(Q \oplus D_O) = area(H(Q, r_b)) = 2dr_b + \pi r_b{}^2$$
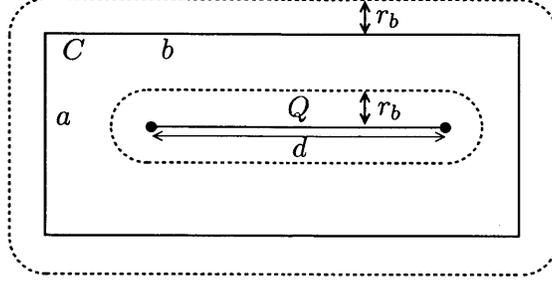
Figure 3-20: The rectangle C and line segment link Q are shown in solid. The center of every disaster that intersects Q is given by the dotted hippodrome $H(Q, r_b)$. The center of every disaster that intersects C is given by $C \oplus D_O$ which is shown as the larger dotted shape.

3.7.2 Single Link Failures

Let $\langle Q \rangle$ and $\langle C \rangle$ be sets of disks of radius r_b in the plane such that $\langle Q \rangle \subset \langle C \rangle$. Given μ, the probability that a 'random' disk is in the set $\langle Q \rangle$ given it is in the set $\langle C \rangle$ is given by the ratio $\frac{\mu(\langle Q \rangle)}{\mu(\langle C \rangle)}$ [65]. Note that C contains the centers of all possible disk failures and is required for normalization purposes.

We now present an example relating to network survivability. Consider a rectangle C with height a and width b and a line segment Q of length d inside C (see Fig. 3-20). Now we consider a random disk-cut. We have:

$$
\begin{aligned}
\Pr(Q \text{ cut} | C \text{ cut}) &= \frac{\mu(\langle Q \rangle)}{\mu(\langle C \rangle)} = \frac{area(Q \oplus D_O)}{area(C \oplus D_O)} \\
&= \frac{2dr_b + \pi r_b{}^2}{ab + 2(a+b)r_b + \pi r_b{}^2}
\end{aligned}
\tag{3.2}
$$

3.8 Geographically Correlated Link Failures Under Circular Cuts

In this section we present an algorithm that calculates the measure of disks of radius r_b intersecting only a particular set of links. This result will allow us to calculate the probability that a random disk-cut intersects a certain set of links in a network (e.g. links whose removal would disconnect the network). We will then use this to efficiently calculate network performance measures with respect to random disk-cuts.
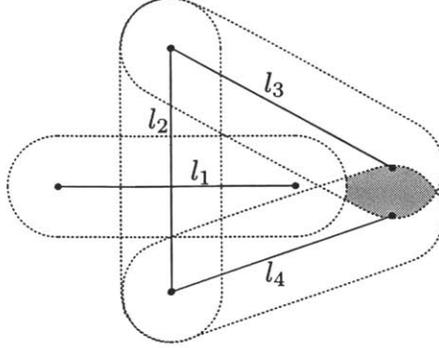
Figure 3-21: The measure of disks that intersect C, l_3, and l_4 but neither l_1 or l_2 is given by the area of the shaded region above (for ease of presentation we do not picture C). This measure can be written as $\mu(\langle C\rangle \cap \langle l_3\rangle \cap \langle l_4\rangle - \langle l_1\rangle \cup \langle l_2\rangle)$ as well as $area((C \oplus D_O) \cap H(l_3, r_b) \cap H(l_4, r_b) - H(l_1, r_b) \cup H(l_2, r_b))$.

Let L be the set of all line segment links in the network and C be a convex polygon that contains L. Consider some set of links $K \subset L$. We wish to find the measure of all disks of radius r_b that intersect C and every link in K but intersect no links in $L - K$. See Fig. 3-21 for an example. This measure is given by $\mu\left(\langle C\rangle \cap \left(\cap_{k\in K}\langle k\rangle\right) - \cup_{q\in(L-K)}\langle q\rangle\right)$. It is clear that a disk D belongs to this measured set iff i) $[x_D, y_D] \in C \oplus D_O$, ii) $[x_D, y_D] \in H(k, r_b)$ $\forall k \in K$, and iii) $[x_D, y_D] \notin H(q, r_b)$ $\forall q \in (L - K)$. So, this measure can also be written as $area\left((C \oplus D_O) \cap \left(\cap_{k\in K} H(k, r_b)\right) - \cup_{q\in(L-K)} H(q, r_b)\right)$. For ease of presentation we abuse notation and denote this measure by $area(K)$.

Definition 11 ($area(K)$). *Let $area(K)$ be given by the measure of all disks of radius r_b that intersect C and every link in K but intersect no links in $L - K$.*

3.8.1 Approximation

We note that finding $area(K)$ seems difficult because it requires finding the area of intersections and unions of hippodromes. In the following we describe a method for approximating $area(K)$ which is based on approximating hippodromes by polygons for which there are known methods to calculate intersections, unions, and area. We approximate $H(l, r_b)$ by the inscribing polygon $\hat{H}_n(l, r_b)$ such that $H(l, r_b)$ shares the line segment portion of its boundary with $\hat{H}_n(l, r_b)$ and each end of $\hat{H}_n(l, r_b)$ forms half of a regular $2n$-sided polygon (see Fig. 3-22). Let $\widehat{area}_n(K)$ be defined the same

85
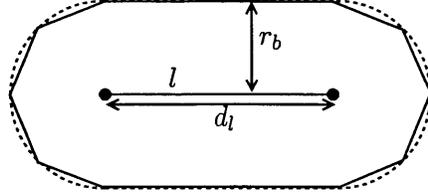
Figure 3-22: We approximate $H(l, r_b)$, shown as a dashed hippodrome above, by the inscribing polygon $\hat{H}_n(l, r_b)$ such that $H(l, r_b)$ shares the line segment portion of its boundary with $\hat{H}_n(l, r_b)$ and each end of $\hat{H}_n(l, r_b)$ forms half of a regular $2n$-sided polygon. The solid polygon above is $\hat{H}_4(l, r_b)$.

as $area(K)$ except that every hippodrome is replaced by its polygon approximation. Using techniques for finding the intersection, union, and area of polygons [55], we can find $\widehat{area}_n(K)$ in polynomial time.

We now present a lemma that shows $\widehat{area}_n(K)$ is a good approximation for $area(K)$ for large enough n. A proof may be found in the appendix.

Lemma 19. $\lim_{n \to \infty} \widehat{area}_n(K) = area(K) \ \forall \ K \subset L$

3.9 Evaluating Network Reliability Metrics

In this section we introduce and show how to evaluate some performance metrics with respect to a random disk-cut.

3.9.1 Network Model and Metrics

We start by describing our network model. Let N be the set of nodes in our geographical graph where each node is represented by a point on the plane. A link between two nodes is represented by a line segment with endpoints at the respective node locations. Let L be the set of all links in the network; we assume every link has a capacity associated with it. Let C denote the convex polygon containing the network (C is required in order to assign probabilities to random disk events). If a 'random' disk that intersects C also intersects some links, those links are disrupted. Our goal is to evaluate the performance metrics described below after a single random disk-cut that intersects C.

We now introduce two important network performance metrics which are evaluated after the removal of the intersected links. We will use tools from geometric
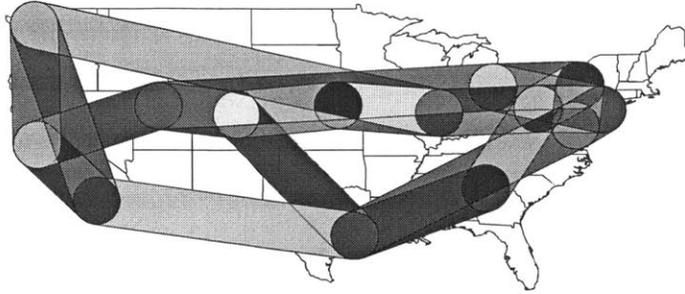
Figure 3-23: This figure shows the hippodromes and related intersections with respect to NSFNET [51] and a circular cut of $r_b = 2$.

probability to evaluate 'average' values of these metrics with respect to a random disk-cut. The first metric we consider is the total capacity of intersected links. Let this metric be denoted by TC. The other metric we consider is the fraction of node pairs that remain connected to each other. Let this metric be denoted by $ATTR$.

3.9.2 Evaluation of the Metrics Under Random Circular Cuts

We now show how to evaluate the metrics above with respect to a random-disk cut or radius r_b. The basic idea is that the center of all disks of radius r_b that intersect a particular set of links (and no other links) is some set in the plane. By showing the number of these sets we need to consider grows polynomially in N and by evaluating the area of each set, we can evaluate a 'weighted average' of a metric over all possible cuts.

Let P be the set of all subsets of L that can be intersected by exactly one disk of radius r_b. Evaluating performance metrics to a random disk-cut is a weighted average over every $K \in P$. Let $Y(K)$ be a reliability metric evaluated after the removal of every link in K. Since $\frac{area(K)}{area(C \oplus D_O)}$ is the probability a random disk of radius r_b that intersects C also intersects every link in K and no links in $(L - K)$, the performance metric to a random disk-cut can be expressed as:

$$\sum_{K \in P} \frac{area(K)}{area(C \oplus D_O)} Y(K) \qquad (3.3)$$

Section 3.8 shows how to approximate $area(K)$ in polynomial time. $Y(K)$ for the

performance metrics can also be calculated in polynomial time. In the following, we apply the theory of arrangements to show that the size of P grows polynomially with respect to N. For technical reasons this theory cannot be directly applied to this setting and requires modification.

Let ∂ denote the boundary of a set. Consider the set of curves $R = \partial C \cup \{\partial H(l, r_b) | l \in L\}$. These curves partition C, the set containing the network, into maximally connected regions called faces that are bounded by the curves in R. By enumerating these faces, we can enumerate every element in P (since every disk in a particular face intersects the same links). Arrangements, a computational geometry tool, allow us to enumerate the faces of a set of curves in $\mathbb{R}^2$ in polynomial time. However, the theory requires that every pair of curves intersect in a finite number of locations [31] which does not hold in our setting. Nonetheless, the theory can be applied with a minor perturbation to the geometry.

Since enumerating P, evaluating $Y(K)$, and approximating $area(K)$ all take polynomial time, the network performance metrics can be approximated in polynomial time under a random disk failure.

3.10 Numerical Results

In this section we evaluate some network metrics using the results of the previous section. We consider NSFNET as found in 1991 [51] and the ARCOS-1 ring network [29]. The NSFNET network we consider has 14 nodes and connects major universities across the U.S. (see Fig. 3-23). ARCOS-1 has 24 nodes and connects regions on the Dominican Republic, Florida, Mexico, Panama, and Venezuela (see Fig. 3-24). All distance units mentioned in this section are in longitude and latitude coordinates (one unit is approximately 60 miles) and for simplicity we assume latitude and longitude coordinates are projected directly to [x, y] pairs on the plane. We assume that all the link capacities are equal to 1. We also assume each network is contained within a rectangular set C.
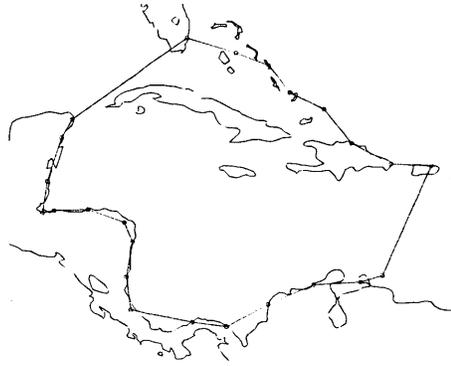
Figure 3-24: ARCOS-1 network circa 2009 [29].
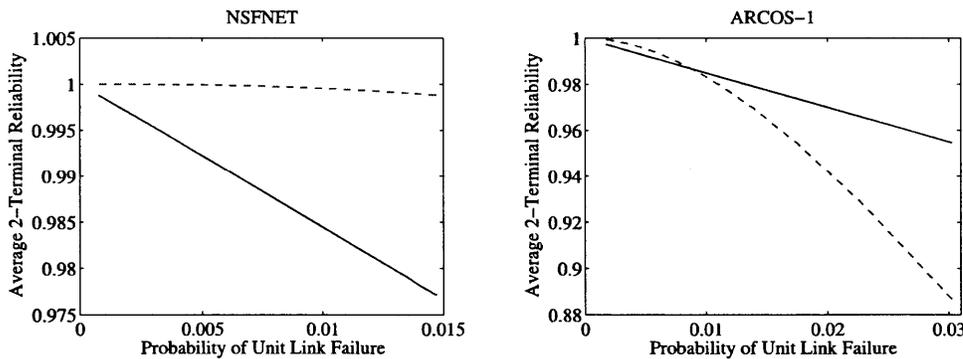
Figure 3-25: The solid line shows $ATTR$ versus the probability a unit (latitude/longitude) of fiber is cut by a random disk of $r_b = 2$. The dashed line shows $ATTR$ assuming links fail *independently* such that links fail with the same probability as in the random disk-cut case.

3.10.1 Independent Versus Correlated Failures

Using the results of Section 3.9, we calculate $ATTR$ of NSFNET and ARCOS-1 to random-disk cuts of $r_b = 2$ while the size of C varies. The size of C is varied to change the probability a unit of fiber is cut. So we can plot $ATTR$ versus the probability a unit of fiber is cut. See Fig. 3-25 for results. Note the linear form of the result in the figure; this agrees with Equation 3.3 since $1/area(C \oplus D)$ is proportional to the probability a unit of fiber is cut.

Next, we calculate $ATTR$ of the networks assuming independent link failures such that links fail with the same probability as in the random disk-cut case. Thus the probability a link fails is still a function of its length, however links fail independently. Since the total number of links is small in each network, calculating $ATTR$ by enumerating all possible failures is still feasible (possible failures are exponential in number of links). Note the total expected number of removed links is the same for

89

both the independent and geographically failure models. See Fig. 3-25 for results.

Notice that in NSFNET *ATTR* under independent failures is greater than in the case of random disk-cuts. Perhaps this is because in most cases at least three links must fail independently to disconnect the network; however a disk that intersects a node is guaranteed to disconnect the network. Since most backbone networks are likely to be well connected, we expect a random disk-cut to lead to lower *ATTR* than independent link failures in this type of mesh network setting. We also note that similar results were found for the random line-cut setting.

Looking at the results for the ARCOS-1 network we see the opposite tendency; *ATTR* under independent failures is typically less than the case of random disk-cuts. Perhaps this is because a single disk that intersects ARCOS-1 usually only removes two adjacent links creating components of size 1 and $|N| - 1$ (where $|N|$ is the number of nodes) whereas just two independent link failures on opposite sides of the ring create components of size $|N|/2$ and $|N|/2$ (which results in lower ATTR).

3.10.2 Multiple Disk Failures

We calculate the TC metric of the NSFNET and ARCOS-1 networks under sequential disk failures, both intentional and random. We assume every additional random failure is located independently of the previous failures. We first describe how to evaluate metrics after sequential failures, then we present some numerical results.

To calculate a network metric after two randomly located sequential disk failures, we simply evaluate the weighted average of the metric over each pair of possible areas (each area represents the set of centers of disks that remove exactly the same links). Equation 3.3 then becomes $\sum_{K' \in P} \sum_{K \in P} \frac{area(K') \, area(K)}{area(C \oplus D_O)^2} Y(K \cup K')$. For n failures, Equation 3.3 becomes

$$\sum_{K_1 \in P} \cdots \sum_{K_n \in P} \left(\prod_{i=1}^{n} \frac{area(K_i)}{area(C \oplus D_O)} \right) Y(\bigcup_{i=1}^{n} K_i).$$

In chapter 2 we propose an algorithm to evaluate network reliability metrics after an *intentional* disk failure. To calculate a network metric after sequential intentional

NSFNET
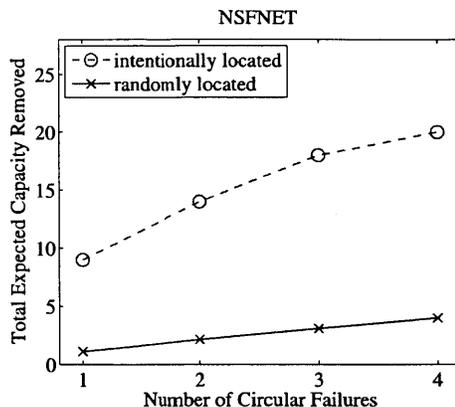
Figure 3-26: The solid line shows the total expected capacity removed versus the number of *randomly* located circular disasters of $r_b = 2$. Using the algorithm in chapter 2, the dashed line shows the total capacity removed versus the number of *intentionally* located circular disasters of $r_b = 2$.

failures, we simply apply the algorithm found in chapter 2 iteratively.

Fig. 3-26 shows the results for multiple failures, both intentional and random for NSFNET (similar results for ARCOS-1 are not shown). As expected, the plots are sub-linear since each additional failure is being placed on a smaller network. Note that random failures result in much less disruption than intentional failures.

3.11 Network Design Under Random Circular Cuts

In this section we discuss some network design problems in the context of random disk-cuts. In all the proposed problems the location of every node is fixed; the problem is to find a set of links most robust to some metric under some constraints. In the following, let N be a set of nodes fixed on the plane and assume all links are represented as line segments between the nodes.

(i) A reasonable goal is to design a connected network with the least expected number of links cut by a random disk of radius r_b. By Eqn. 3.2 and linearity of expectation, the expected number of links cut is proportional to $\pi r_b^2(|N| - 1) + 2r_b \sum_{l \in U} d_l$ where U is the set of links chosen and d_l is the length of link l. So, this problem reduces to minimizing the total length of links in the network while ensuring the graph is connected. This is equivalent to finding a Euclidean minimum spanning tree of N which can be done in polynomial time. Note however that the resulting

91

network is not robust because a single link failure will disconnect it. We also note this is the same result we get for optimizing for random line-cuts. If fact, it can be shown that this result will hold for any convex-shaped cut.

(ii) We next consider ring networks. A reasonable goal is to design a connected ring network with the least expected number of links cut by a random disk of radius r_b. As before, the expected number of links cut is affine in the total length of the links. So this problem reduces to finding the minimum length Hamiltonian cycle. This is equivalent to the Euclidean traveling salesman problem which is hard to compute [43]. We again note this is the same result we get for optimizing for random line-cuts. If fact, it can be shown that this result will hold for any convex-shaped cut.

(iii) The final problem considers how to connect two nodes such that the path between them is robust to a random-disk disaster of radius r_b. Let S and T be a pair of nodes in N and let U be a set of links. The problem is to find a ST path consisting of links from U that has the minimum probability of being cut. This may correspond to finding the most robust path between two cities along preexisting conduits. Since a disk intersects a path iff it intersects a hippodrome corresponding to a link in a path, we want to find a ST path whose edges belong to U and whose area of the union of corresponding hippodromes is minimized. The authors do not know a polynomial time algorithm to solve the above problem except in a trivialized setting.

An interesting example is given in Fig. 3-27. The uppermost path here gives the most robust path to a random disk failure; however, the bottom path is the shortest. This shows the shortest path is not necessarily the most robust to failure.

3.12   Conclusions

Motivated by applications in the area of network robustness and survivability, we focused on the problem of geographically correlated network failures. Namely, we focused on randomly located geographical attacks on the network which can model the 'random' nature of a natural disaster or collateral damage. In particular, we focused on random line and disk cuts. Using tools from geometric-probability we demon-

92
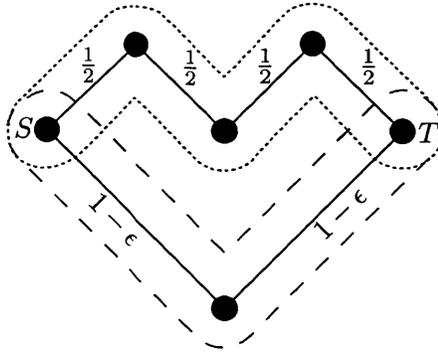
Figure 3-27: The bottom path is shorter than the top path by 2ϵ. However the top path is more robust to a random circular failure as compared to the bottom path since the blue dotted shape has less area than the red dashed shape. This shows the shortest path is not necessarily the most robust to failure.

strated how to compute failure probabilities and showed how to calculate $ATTR$ and other network performance metrics in polynomial time under these failure models. This result is significant because calculating this metric assuming independent link failures in known to be NP-hard [10]. We then presented some numerical results to demonstrate the significance of geometry on the survivability of the network and also discussed network design problems in the context of random line and disk failures.

Our approach provides a fundamentally new way to look at network survivability that takes into account the geographical correlation between links. Some future research directions include the consideration of multiple line-cuts (instead of a single line failure), convex cuts (e.g., oval cuts), and robust network design in the face of geographical failures.

3.A   Definition of the Measure m and Intuition Behind Lemma 11

We now present the definition of the measure.

Definition 12 (Measure of a set of lines). *The measure of a set of lines G is defined as the integral*

$$m(G) = \int_G d\rho d\theta$$

Note we use G to denote both a set of lines and its equivalent set of points in the (ρ, θ) plane. In some sense, this integral is the area of G in the (ρ, θ) plane.
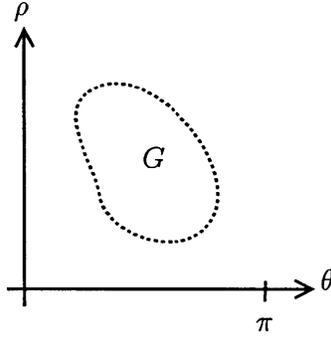
Figure 3-28: Since every line in the plane can be parameterized by ρ and θ, we can represent a line in G as a point in the (ρ,θ) plane. Integrating over the set of all these points allows us to assign a measure to the set G.

For a visualization of the measure, see Fig. 3-28. Since every line in the plane can be parameterized by ρ and θ, we can represent each line in G as a point (ρ,θ) in $\mathbb{R} \times [0, \pi)$.

Note that given this definition we can show the measure with respect to any line segment is twice its length. This fact is used below.

Now we present an intuitive argument to why $m([C]) = L_C$. Given a convex set C, let S be an ordered set of line segments such that the union of all the line segments in S form a closed convex curve that approximates the boundary of C (see Fig. 3-29). Let s_i be the i^{th} element of S and let d_i be the length of line segment s_i. Let $l(\rho, \theta)$ be the line parameterized by ρ and θ. Now let

$$N(\rho, \theta) = \sum_{j=1}^{|S|} \mathbf{1}_{l(\rho,\theta) \cap s_j \neq \emptyset}$$

That is, $N(\rho, \theta)$ is the total number of $s_i \in S$ that are intersected by the line parameterized by ρ and θ.

94
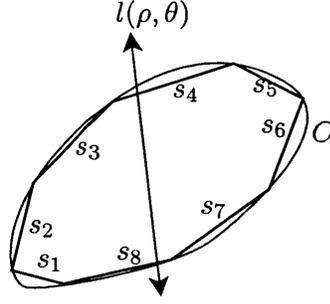
Figure 3-29: The blue curve in the figure above is the boundary of a convex set C. $\cup_{i=1}^{8} s_i$ is a closed convex curve which approximates the boundary of C. $S = \{s_i | i \in \mathbf{N}$ and $1 \le i \le 8\}$. $N(\rho, \theta) = \sum_{j=1}^{8} 1_{l(\rho,\theta) \cap s_j \ne \emptyset}$, that is $N(\rho, \theta)$ is the number of $s_i \in S$ that are intersected by the line parameterized by ρ and θ.

Now consider,

$$
\begin{aligned}
\int_{[\cup_{i=1}^{|S|} s_i]} N(\rho, \theta) d\rho d\theta &= \int_{[\cup_{i=1}^{|S|} s_i]} \sum_{j=1}^{|S|} 1_{l(\rho,\theta) \cap s_j \ne \emptyset} d\rho d\theta \\
&= \sum_{j=1}^{|S|} \int_{[\cup_{i=1}^{|S|} s_i]} 1_{l(\rho,\theta) \cap s_j \ne \emptyset} d\rho d\theta \\
&= \sum_{j=1}^{|S|} \int_{[s_j]} d\rho d\theta \\
&= \sum_{j=1}^{|S|} 2 d_j
\end{aligned}
$$

The key is to note that since $\cup_{i=1}^{|S|} s_i$ forms a closed convex curve, $N(\rho, \theta) = 2$ if $l(\rho, \theta) \cap (\cup_{i=1}^{|S|} s_i) \ne \emptyset$ except for a set of lines with measure zero. So,

$$
\int_{[\cup_{i=1}^{|S|} s_i]} N(\rho, \theta) d\rho d\theta = \int_{[\cup_{i=1}^{|S|} s_i]} 2 d\rho d\theta = 2m([\cup_{i=1}^{|S|} s_i])
$$

Thus,

$$
m([\cup_{i=1}^{|S|} s_i]) = \sum_{j=1}^{|S|} d_j
$$

Since $\cup_{i=1}^{|S|} s_i$ approximates the boundary of C, it seems reasonable that $m([C]) = L_C$.

3.B Proof of Lemma 12

This proof can be found in [50] and is only presented here for completeness and convenience.

Proof. Let A and B be bounded closed convex sets and let S^C denote the complement of set S. First note that $[A \cup B] = [A] \cup [B]$. This is because a line intersects $A \cup B$ iff it intersects A or it intersects B. Now note that $[A] = ([A] \cap [B]^C) \cup ([A] \cap [B])$ and $[B]$ and $[B]^C$ are disjoint, so

$$m([A]) = m([A] \cap [B]^C) + m([A] \cap [B]) = L_A$$

Similarly note that $[B] = ([B] \cap [A]^C) \cup ([B] \cap [A])$ and since $[A]$ and $[A]^C$ are disjoint, so

$$m([B]) = m([B] \cap [A]^C) + m([B] \cap [A]) = L_B$$

We will now show that the theorem holds true in the case where $A \cap B \neq \emptyset$ and then show it also holds when $A \cap B = \emptyset$.

First assume $A \cap B \neq \emptyset$. Since a line intersects $A \cup B$ iff the line intersects $conv(A \cup B)$, we know $[A \cup B] = [conv(A \cup B)]$. So,

$$m([A] \cup [B]) = m([A \cup B]) = L_{conv(A \cup B)}$$

Note that

$$[A] \cup [B] = ([A] \cap [B]^C) \cup ([A]^C \cap [B]) \cup ([A] \cap [B])$$

and so

$$\begin{aligned} m([A] \cup [B]) &= m([A] \cap [B]^C) + m([A]^C \cap [B]) \\ &\quad + m([A] \cap [B]) \\ &= L_{conv(A \cup B)} \end{aligned}$$
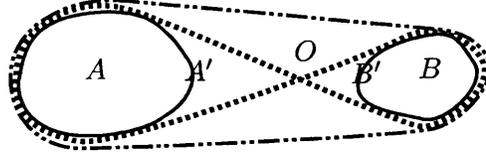
Figure 3-30: The sets with solid boundary are two disjoint convex sets A and B. The sets with dashed boundary are two convex sets A' and B' such that $A' \cap B' = O$ and $A' \cup B' = I(A,B)$. Note that a line intersects A and B and not O iff the line intersects A' and B' and not O. The set with the dash-dot boundary is $conv(A \cup B)$ and note that $conv(A' \cup B') = conv(A \cup B)$.

Combining this equation with the ones above, we get

$$
\begin{aligned}
m([A] \cap [B]) &= m([A]) + m([B]) - m([A] \cup [B]) \\
&= L_A + L_B - L_{conv(A \cup B)}
\end{aligned}
$$

which is consistent with the theorem (since $L_{I(A,B)} = L_A + L_B$)

Now assume $A \cap B = \emptyset$. Let A' and B' be supersets of A and B respectively, such that $A' \cap B'$ is a single point, O, and $A' \cup B' = I(A, B)$. For an example illustration, see Fig. 3-30. Note that $conv(A \cup B) = conv(A' \cup B')$. Also note that a line intersects A and B and not O iff it intersects A' and B' and not O. This implies $[A] \cap [B] = [A'] \cap [B']$ up to a set of measure zero and thus $m([A] \cap [B]) = m([A'] \cap [B'])$. Since $A' \cap B' \neq \emptyset$, we can apply the previous result and attain

$$
\begin{aligned}
m([A] \cap [B]) &= m([A'] \cap [B']) \\
&= m([A']) + m([B']) - m([A'] \cup [B']) \\
&= L_{A'} + L_{B'} - L_{conv(A' \cup B')} \\
&= L_{I(A,B)} - L_{conv(A \cup B)}
\end{aligned}
$$

$\square$

3.C   A Lemma About Rays

Lemma 20. *Let* $d([x_a, y_a], [x_b, y_b]) = \sqrt{(x_b - x_a)^2 + (y_b - y_a)^2}$. *Let* A *be a vertical ray with endpoint* $[x_A, y_A]$ *and* B *be a line segment with endpoints* $[x_B^1, y_B^1]$ *and* $[x_B^2, y_B^2]$

such that the line containing B does not intersect A except possibly at $[x_A, y_A]$ and the angle between A and the line segment connecting $[x_A, y_A]$ and $[x_B^1, y_B^1]$ is greater or equal than the angle between A and the line segment connecting $[x_A, y_A]$ and $[x_B^2, y_B^2]$. Then,

$$
\begin{aligned}
m([A] \cap [B]) \;=\;& d([x_A, y_A], [x_B^2, y_B^2]) - d([x_A, y_A], [x_B^1, y_B^1]) \\
& + (y_B^2 - y_B^1) Z_{up}
\end{aligned}
$$

where $Z_{up} = 1$ if A extends infinitely upwards and $Z_{up} = -1$ otherwise.

We first present some intuition behind this lemma. We cannot use Lemma 12 directly to get an expression for $m([A] \cap [B])$ because A is unbounded (it is a ray); however, the intuition remains the same. In some sense, $L_{I(A,B)} - L_{conv(A \cup B)}$ is $d([x_A, y_A], [x_B^2, y_B^2]) - d([x_A, y_A], [x_B^1, y_B^1]) +$ 'length' of vertical ray from $[x_B^1, y_B^1] -$ 'length' of vertical ray from $[x_B^2, y_B^2]$. Again, in some sense, the difference in the 'length' of the vertical rays is $\pm(y_B^2 - y_B^1)$ (the sign depends on the orientation of A) which is the desired result. In order to obtain this rigorously, we use the fact that measures are continuous from below.

Proof. Assume ray A extends infinitely in the upwards direction. Let A_i be the line segment of length i that is a subset of the A and has an endpoint at $[x_A, y_A]$ (Note: $\bigcup_{i=1}^{\infty} A_i = A$). See Fig. 3-31.

We know by Lemma 12 that

$$
\begin{aligned}
m([A_i] \cap [B]) &= L_{I(A_i, B)} - L_{conv(A_i \cup B)} \\
&= d([x_A, y_A], [x_B^2, y_B^2]) + d([x_A, y_A + i], [x_B^1, y_B^1]) \\
&\quad - d([x_A, y_A], [x_B^1, y_B^1]) - d([x_A, y_A + i], [x_B^2, y_B^2])
\end{aligned}
$$

Now note,

$$
\begin{aligned}
m([A] \cap [B]) \;=\;& m([\bigcup_{i=1}^{\infty} A_i] \cap [B]) \\
=\;& m(\bigcup_{i=1}^{\infty} ([A_i] \cap [B]))
\end{aligned}
$$

98
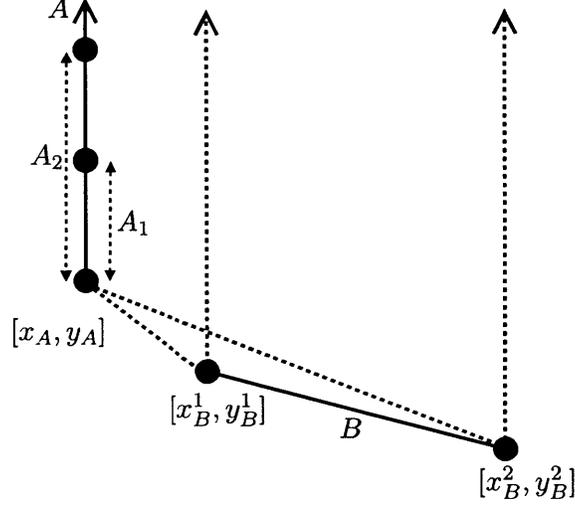
Figure 3-31: We cannot use Lemma 12 directly to get an expression for $m([A] \cap [B])$ because A is unbounded; however, the intuition remains the same. In some sense, $L_{I(A,B)} - L_{conv(A \cup B)}$ is $d([x_A, y_A], [x_B^2, y_B^2]) - d([x_A, y_A], [x_B^1, y_B^1]) +$ 'length' of vertical ray from $[x_B^1, y_B^1] -$ 'length' of vertical ray from $[x_B^2, y_B^2]$. Again, in some sense the difference in the 'length' of the vertical rays is $\pm(y_B^2 - y_B^1)$ which is the desired result. In order to obtain this rigorously, we must use the fact that measures are continuous from below.

Since $A_i \subset A_{i+1}$, this implies $[A_i] \cap [B] \subset [A_{i+1}] \cap [B]$, and since measures are continuous from below we have,

$$m([A] \cap [B]) = m(\cup_{i=1}^\infty ([A_i] \cap [B])) = \lim_{i \to \infty} m([A_i] \cap [B])$$

$$= \lim_{i \to \infty} d([x_A, y_A], [x_B^2, y_B^2]) + d([x_A, y_A + i], [x_B^1, y_B^1])$$

$$- d([x_A, y_A + i], [x_B^2, y_B^2]) - d([x_A, y_A], [x_B^1, y_B^1])$$

$$= d([x_A, y_A], [x_B^2, y_B^2]) - d([x_A, y_A], [x_B^1, y_B^1]) + y_B^2 - y_B^1$$

This is because

$$\lim_{i \to \infty} \sqrt{(x_A - x_B^1)^2 + (y_A + i - y_B^1)^2} - \sqrt{(x_A - x_B^2)^2 + (y_A + i - y_B^2)^2}$$

$$= \lim_{i \to \infty} \frac{(x_A - x_B^1)^2 - 2(y_A + i)y_B^1 + (y_B^1)^2 - (x_A - x_B^2)^2 + 2(y_A + i)y_B^2 - (y_B^2)^2}{\sqrt{(x_A - x_B^1)^2 + (y_A + i - y_B^1)^2} + \sqrt{(x_A - x_B^2)^2 + (y_A + i - y_B^2)^2}}$$

$$= y_B^2 - y_B^1$$

If we assume ray A extends infinitely in the downwards direction, an analogous result follows.

$\square$

3.D   Proof of Lemma 19

We first develop Lemmas 21 and 22 which allow us to prove Lemma 19.

Let $A_l(n) = area \left(H(l, r_b) - \hat{H}_n(l, r_b) \right)$.

Lemma 21. $\lim_{n \to \infty} A_l(n) = 0 \; \forall l \in L$

Proof. Let d_l denote the length of link l. Since $\hat{H}_n(l, r_b) \subset H(l, r_b)$ for all $l \in L$ we know

$$
\begin{aligned}
A_l(n) &= area \left(H(l, r_b) - \hat{H}_n(l, r_b) \right) \\
&= area \left(H(l, r_b) \right) - area \left(\hat{H}_n(l, r_b) \right) \\
&= (2d_l r_b + \pi r_b^2) - (2d_l r_b + n \sin(\pi/n) r_b^2)
\end{aligned}
$$

So,

$$
\lim_{n \to \infty} A_l(n) = (2d_l r_b + \pi r_b^2) - (2d_l r_b + \pi r_b^2) = 0 \; \forall l \in L
$$

$\square$

Lemma 22. *Let $K \subset L$. Now,*

$$
area(K) - |K| A_l(n) \leq \widehat{area}_n(K) \leq area(K) + |L - K| A_{l'}(n) \quad \text{for some } l \in L
$$

and some $l' \in L$.

Proof. We first prove the right hand side.

$\widehat{area}(K)$

$$
= area \left((C \oplus D) \cap (\bigcap_{k \in K} \hat{H}_k) - \bigcup_{q \in (L-K)} \hat{H}_q \right)
$$

$$
\leq area \left((C \oplus D) \cap (\bigcap_{k \in K} H_k) - \bigcup_{q \in (L-K)} \hat{H}_q \right) \quad \text{since } \hat{H}_k \subset H_k
$$

$$
= area \left((C \oplus D) \cap (\bigcap_{k \in K} H_k) - \left(\bigcup_{q \in (L-K)} H_q - \left(\bigcup_{q \in (L-K)} H_q - \bigcup_{q \in (L-K)} \hat{H}_q \right) \right) \right)
$$

since if $B \subset A$ then $B = A - (A - B)$

$$= area \left(\left(\left((C \oplus D) \cap (\bigcap_{k \in K} H_k) - \bigcup_{q \in (L-K)} H_q \right) \cup \right. \right.$$

$$\left. \left. \left((C \oplus D) \cap (\bigcap_{k \in K} H_k) \cap \left(\bigcup_{q \in (L-K)} H_q - \bigcup_{q \in (L-K)} \hat{H}_q \right) \right) \right) \right)$$

since $A - (B - C) = (A - B) \cup (A \cap C)$

$$\leq area \left((C \oplus D) \cap (\bigcap_{k \in K} H_k) - \bigcup_{q \in (L-K)} H_q \right) + area \left(\bigcup_{q \in (L-K)} H_q - \bigcup_{q \in (L-K)} \hat{H}_q \right)$$

by subadditivity

$$\leq area\,(K) + area \left(\bigcup_{q \in (L-K)} (H_q - \hat{H}_q) \right)$$

$$= area\,(K) + area \left(\bigcup_{q \in (L-K)} \left(\left(H_q - \hat{H}_q \right) - \hat{H}_q \right) \right) \quad \text{by subadditivity}$$

$$\leq area\,(K) + \sum_{q \in (L-K)} area \left(H_q - \hat{H}_q \right) \quad \text{by subadditivity}$$

$$\leq area\,(K) + |L - K| A_q(n) \text{ for some } q \in (L - K)$$

We now prove the left hand side.

$\widehat{area}(K)$

$$= area \left((C \oplus D) \cap (\bigcap_{k \in K} \hat{H}_k) - \bigcup_{q \in (L-K)} \hat{H}_q) \right)$$

$$\geq area \left((C \oplus D) \cap (\bigcap_{k \in K} \hat{H}_k) - \bigcup_{q \in (L-K)} H_q) \right) \quad \text{since } \hat{H}_k \subset H_k$$

$$= area \left((C \oplus D) \cap \left(\bigcup_{k \in K} H_k - \left(\bigcup_{k \in K} H_k - \bigcup_{k \in K} \hat{H}_k \right) \right) - \bigcup_{q \in (L-K)} H_q \right)$$

since if $B \subset A$ then $B = A - (A - B)$

101

$$= area\left(\left(\left((C \oplus D) \cap (\bigcup_{k \in K} H_k)\right) - \left(\bigcup_{k \in K} H_k - \bigcup_{k \in K} \hat{H}_k\right)\right) - \bigcup_{q \in (L-K)} H_q\right)$$

since $A \cap (B - C) = (A \cap B) - C$

$$= area\left(\left((C \oplus D) \cap (\bigcup_{k \in K} H_k) - \bigcup_{q \in (L-K)} H_q\right) - \left(\bigcup_{k \in K} H_k - \bigcup_{k \in K} \hat{H}_k\right)\right)$$

since $(A - B) - C = (A - C) - B$

$$\geq area\left((C \oplus D) \cap (\bigcup_{k \in K} H_k) - \bigcup_{q \in (L-K)} H_q\right) - area\left(\bigcup_{k \in K} H_k - \bigcup_{k \in K} \hat{H}_k\right)$$

by subadditivity

$$\geq area\left((C \oplus D) \cap (\bigcup_{k \in K} H_k) - \bigcup_{q \in (L-K)} H_q\right) - area\left(\bigcup_{k \in K} \left(H_k - \hat{H}_k\right) - \bigcup_{k \in K} \hat{H}_k\right)$$

$$\geq area\,(K) - area\left(\bigcup_{k \in K} (H_k - \hat{H}_k)\right) \quad \text{by subadditivity}$$

$$\geq area\,(K) - \sum_{k \in K} area\left(H_k - \hat{H}_k\right) \quad \text{by subadditivity}$$

$$= area\,(K) - |K| A_k(n) \text{ for some } k \in K$$

$\square$

We now prove Lemma 2. Note

$$\lim_{n \to \infty} (area(K) - |K| A_l(n)) \leq \lim_{n \to \infty} \widehat{area}_n(K) \text{ for some } l \in L$$

So,

$$area(K) - |K| \lim_{n \to \infty} A_l(n) \leq \lim_{n \to \infty} \widehat{area}_n(K) \text{ for some } l \in L$$

Thus by Lemma 21,

$$area(K) \leq \lim_{n \to \infty} \widehat{area}_n(K)$$

Similarly, we have

$$\lim_{n \to \infty} \widehat{area}_n(K) \leq \lim_{n \to \infty} \left(area(K) + |L - K| A_{l'}(n) \right) \text{ for some } l' \in L$$

So,

$$\lim_{n \to \infty} \widehat{area}_n(K) \leq area(K) + |L - K| \lim_{n \to \infty} A_{l'}(n) \text{ for some } l' \in L$$

Thus by Lemma 21,

$$\lim_{n \to \infty} \widehat{area}_n(K) \leq area(K)$$

Chapter 4

Geographic Min-Cut and Max-Flow

In chapter 2 we considered the problem of finding the worst-case location for a failure in a geographic network with respect to certain network connectivity measures. This models the scenario where the network is attacked once with the intention to reduce its capacity or connectivity. In chapter 3 we considered the impact of a randomly located disaster on network connectivity; the random location of the disaster can model failures resulting from a natural disaster or collateral (non-targeted) damage in an attack. In this chapter we consider the problem of finding the minimum number of failures, modeled as circular disks, to disconnect two nodes and the maximum number of failure disjoint paths between two nodes. This approach provides a way to look at network survivability in the face of *multiple* disasters or attacks that takes into account the geographical correlation between links.

4.1   Introduction

We first consider a geographical variant of the min-cut problem. Given a set of points on the plane, each of which represents a node, and non-crossing line segments between these points representing links, what is the minimum number of circular attacks such that two nodes, S and T, are disconnected from each other. If applied to the national
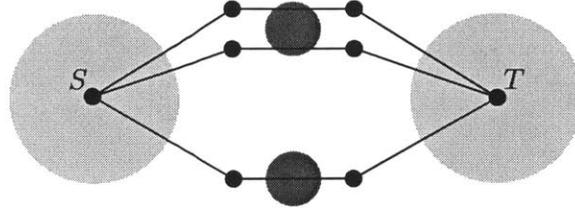
Figure 4-1: The light gray area (yellow area in online color version) above represents the protected zone that no circular failure may be centered. The gray disks (red disks in online color version) represent disasters that remove links (of unit capacity) they intersect. Two disasters are required to disconnect the two nodes S and T (shown above), so the geographic min-cut is two. Also, since the top pair of paths can be intersected by the same failure, geographic max-flow is two; two failure disjoint paths are given by the topmost and bottommost path. In contrast, the standard min-cut and max-flow is three.

fiber plant, the solution to this problem is the number of geographic failures required to disconnect two cities. If we do not restrict the locations of potential failure sites, the geographic min-cut will be at most one because nodes S or T can trivially be eliminated with a single failure. In order to make the problem more relevant and realistic we restrict potential failure locations (see Fig. 4-1). This can represent fibers that have been hardened against EMP attacks or a well defended city. We show that we only need to consider a polynomial number of possible failure sites, thus reducing the geographic min-cut to a discrete problem. Then applying the methods from [18], we show how to find a solution in polynomial time. We obtain numerical results for a specific backbone network [46], thereby demonstrating the applicability of our min-cut algorithm to a real-world network.

Next, in the context of geographic attacks and path-protection algorithms we study a geographic max-flow problem: the largest set of paths between nodes S and T such that no two paths can be intersected by the same failure. The solution to this problem gives the maximum number of paths that are geographically disjoint with respect to disasters of a particular radius (i.e. the maximum number of failure disjoint backup paths). See Fig. 4-1 for an example. Again, to avoid triviality we restrict the locations of potential failure sites so that nodes S or T cannot simply be eliminated with a single failure. We then develop an ILP formulation, an exact algorithm, and a heuristic algorithm for this geographic max-flow problem.

Finally, we explore the analogue to the min-cut max-flow theorem in the geographic setting. In particular, we show that the cardinality of the solutions to these

geographic min-cut and max-flow problems are not the same. Supported by simulation results, we conjecture this difference is no greater than one, i.e. max-flow $\leq$ min-cut $\leq$ max-flow +1.

4.2 Related work

The traditional min-cut and max-flow problem has been extensively studied in the literature, however most attention has been focused on the single layer setting. Under this model, min-cut is the minimum number of links to disconnect S and T and max-flow is the maximum number of link disjoint paths. In [45] the authors generalized this min-cut and max-flow concept to a cross layer fiber network setting where a single physical link failure may disconnect several logical links. In the same vein, our work may be seen as a min-cut max-flow problem in another cross layer setting where a single physical disaster may remove several physical links in the vicinity of one another.

Min-cut and max-flow problems similar to the ones presented here have also received some attention in the literature. Recently [66] considered the problem of a geographic max-flow and min-cut in a wireless network setting. In [52] the problem of finding the maximum number of geographically disjoint paths with total minimum cost is discussed in the continuous setting where paths may be placed anywhere within a polygonal domain. Finally, [18] considers a related problem to the geographic max-flow and min-cut, where failures of nearly arbitrary shape[1] can occur in a finite set of locations. Here we take the geography into account by allowing failures to take place at any location, yet restricting the shape of a failure to a geometric disk.

4.3 Geographic Min-Cut

We start by formulating the geographic min-cut problem and presenting an algorithm to solve this problem in polynomial time. We then use this algorithm to analyze the

---

[1]In particular, every disaster in [18] must have a shape that in homeomorphic to the unit disk.

vulnerability of a real-world network.

4.3.1   Network Model and Problem Formulation

Let N be an ordered set of points in the plane representing nodes. Assume the points representing the nodes are in general position, that is no three points are collinear. Let $n \in N$; the cartesian pair $[x_n, y_n]$ denotes the location of n. Denote an undirected link from node i to node j by (i, j). For simplicity of presentation we denote by (i, j) as a link represented as a line segment in the plane with endpoints at node i and node j. Let the set of undirected links be given by E. We assume that the graph is simple (contains no self-loops or multiple edges) and connected, and links do not intersect each other except at node locations.

We model a disaster as a closed disk of radius r_b centered at some point $[x, y]$. We denote this disk as $\text{hole}_{r_b}([x, y])$. A hole removes all links that intersect it. Every hole in this chapter is assumed to have radius r_b. Let D be an ordered set of holes where d_i is the i^{th} hole. Let $[x_{d_i}, y_{d_i}]$ denote the cartesian pair that corresponds to the center of hole d_i. In the problem defined below, we assume a hole may be centered anywhere in the plane, except inside a protective disk of radius r_p centered at nodes S and T.

We now define the following problem and demonstrate its formulation.

Geographical Min-Cut By Circular Disasters (GMCCD) Problem: *Given a graph drawn in the plane $G = \{N, E\}$, two distinct nodes S and T, hole radius r_b, and protection radius r_p, find a minimum cardinality set of holes that disconnect S and T.*

Let $E^D = \{e \in E \mid \forall d_i \in D, e \cap d_i = \emptyset\}$ and $G^D = \{N, E^D\}$. Intuitively, G^D is the graph G after the links intersecting any hole in D have been removed. The solution

108

to the GMCCD optimization problem below is a geographical min-cut.

$$\min \ |D|$$

$$\text{such that } \forall d_i \in D, \quad \sqrt{(x_{d_i} - x_S)^2 + (y_{d_i} - y_S)^2} \geq r_p \qquad (4.1)$$

$$\forall d_i \in D, \quad \sqrt{(x_{d_i} - x_T)^2 + (y_{d_i} - y_T)^2} \geq r_p \qquad (4.2)$$

$$S \text{ and } T \text{ are in different components of } G^D$$

Denote C by the cardinality of a solution to the GMCCD problem.

In the above we seek the minimum cardinality set of holes such that S and T are disconnected after their removal. Constraints (4.1) and (4.2) ensure that hole locations are not in the protected zone.

4.3.2 Algorithm to Solve GMCCD Problem

Here we describe an algorithm that finds a solution to the GMCCD problem. For clarity of presentation we break down the algorithm into steps. We initially note that holes may be centered anywhere not inside the protective disks; thus there are an infinite uncountable number of holes to consider in general. The first step (step 1) of the algorithm reduces this infinitely sized set of potential holes to a polynomial sized set by extending the methods in chapter 2. Once this set of holes is enumerated, we can apply a simplified algorithm for computing geographic min-cut based on [18]. We do this by first creating a dual-like graph (step 2) and then running an algorithm based on shortest closed walks on this new dual-like graph to solve the GMCCD problem (step 3).

Step 1: There are an infinite number of hole locations centered outside the protective disks; in this step we find a polynomial sized set of holes from which we can construct a solution to the GMCCD problem.

Before proceeding, we introduce some notation. Let $H(e, r_b)$ be the set of points whose shortest distance to line segment e is less than or equal to r_b. Such a shape is known as a hippodrome [32]. Note that a hole of radius r_b is centered in $H(e, r_b)$ if
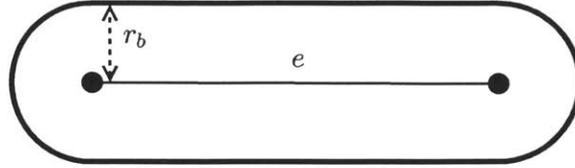
Figure 4-2: Let $H(e, r_b)$ be the set of points whose distance to link l is less than or equal to r_b. Such a shape (shown above) is known as a *hippodrome* [32]. A hole of radius r_b is centered in $H(e, r_b)$ if and only if the hole intersects l.
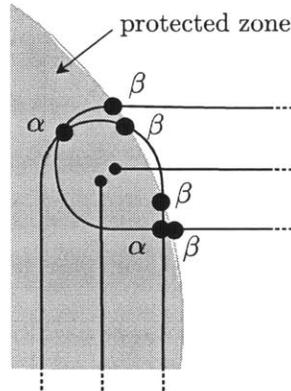
Figure 4-3: The light gray area (yellow area in the online color version) above represents the protected zone. α represent centers of some holes given by the algorithm in chapter 2. These holes intersect both links above, however they are centered in the protected zone. We consider additional holes centered at the points labeled β. Note two of these points correspond to holes that intersect both links and are not centered within the protected zone.

and only if the hole intersects e (see Fig. 4-2).

In chapter 2 we considered the same failure model without the protected zone. Under this model we found a polynomial size set of hole locations such that every hole in the plane can be represented by one of these locations and intersects at least the same set of links. For example, any hole centered in the intersection of the two hippodromes in Fig. 4-3 can be represented by a hole centered on one of the two points labeled α. Holes centered at one of these locations will intersect the same links as any hole centered in the intersection of the hippodromes.

The polynomially sized set of potential failure locations found by the algorithm in chapter 2 cannot be used for the GMCCD problem because of the restrictions that holes cannot be placed inside the protected zones. For example, the set of holes found using this method would have us consider the holes marked with α in Fig. 4-3. However, these holes are centered inside the protected zone and cannot be considered. If we consider additional holes that are centered at the intersection of the boundaries of the protected zones and hippodromes (shown by points labeled by β in Fig. 4-3),

110
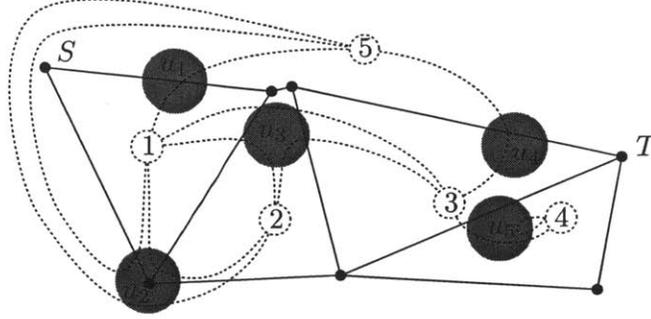
Figure 4-4: The dual-like graph is shown by the dotted potion of the figure above. The solid dots and line segments represent the original network G. For ease of presentation, we take the set of gray disks (red disks in the online color version) above to be U. G has five faces; each of these faces represents a node in K (shown as dashed circles). There exists a link between two nodes in K for each hole in U that intersects the faces they represent. Note, there exist two holes intersecting face one and face five, u_1 and u_2. So there exist two links between node one and node five in K; one corresponding to u_1 and the other corresponding to u_2. Also, for presentation purposes the only self-loop in K shown is located at node 4 and corresponds to u_5; there are more self-loops in K (see Fig. 4-6).

we can show that this expanded set of potential failure locations is sufficient. We omit the details here, which can be found in Appendix 4.A. Let this polynomially sized set of potential hole locations for the GMCCD problem be given by U.

Step 2: We construct an undirected dual-like graph from G, the original graph, and U, the polynomially sized set of potential hole locations. Let this dual-like graph be denoted by K.

We first introduce some notation. The drawing of G in the plane partitions the rest of the plane into connected regions called faces (even the outer, infinitely large region). For example, the graph in Fig. 4-4 divides the plane into five faces, four bounded faces and one infinitely large face.

We now describe the dual-like graph K. Every node in the dual-like graph K corresponds to a face in G. For example, in Fig. 4-4 G has five faces; each of these faces represents a node in K (shown as dashed circles). There exists a link between two nodes in K for each hole $u \in U$ that intersects the faces they represent. For example, in Fig. 4-4 there exist two holes intersecting face one and face five, u_1 and u_2. So there exist two links between node one and node five in K; one corresponding to u_1 and the other corresponding to u_2. Note, because every link in K is associated with a hole, there exist more than one edge between two nodes in K if more than one hole intersects their corresponding faces.
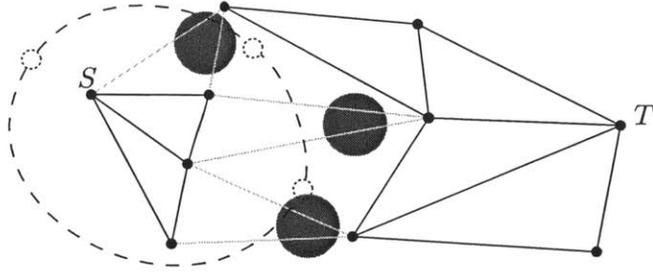
Figure 4-5: The dashed links above represents a closed walk in K such that the corresponding holes (shown as disks) remove links which disconnect S and T. By searching over a set of closed walks in K, we will be able to find a solution to the GMCCD problem.

Step 3: The final step finds a solution to the GMCCD problem by considering a set of closed walks in K and then from this set finds the shortest walk whose corresponding holes disconnect S and T (see Fig. 4-5). This is similar to a known algorithm to find the min-cut in a planar graph (in the standard sense); the algorithm finds the shortest closed walk in the dual graph that disconnects S from T [62].

We now describe the algorithm. First, for all nodes in the dual-like graph run Dijkstra's algorithm [19]. This gives a shortest path tree rooted at n. Denote links in this tree by C_n. Notice that when a set of links is removed from the graph new faces are created. Intuitively a shortest path in K between two nodes gives the minimum number of disasters such that the faces corresponding to these nodes will be contained in a larger face after the disaster. It is worth emphasizing that this face is not necessarily the outer face of the new graph.

Next, for every link e in K consider the closed walk in $C_n \cup e$ which contains node n and link e. A solution to the GMCCD problem is given by finding the closed walk in $C_n \cup e$ for all n nodes and links e in K and then searching over these walks for the shortest one whose corresponding holes disconnect S and T.

For example, consider Fig. 4-6. Let the link from node n_i to n_j associated with hole u be given by $\{(n_i, n_j), u\}$. The solid links are the links in the shortest path tree rooted at node 2, C_2. Consider the link $\{(1,5), u_2\}$. Now $C_2 \cup \{(1,5), u_2\}$ contains a closed walk given by $\left\{\{(1,2), u_2\}, \{(2,5), u_2\}, \{(1,5), u_2\}\right\}$. Since hole u_2 does not disconnect S and T (every hole in this cycle is marked with u_2), $\{u_2\}$
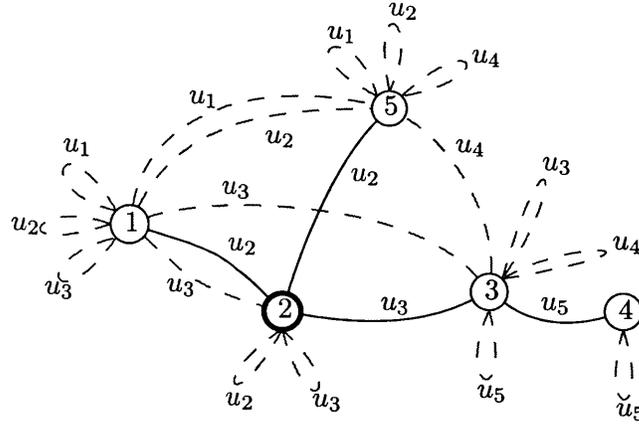
Figure 4-6: The graph shown above is K from Fig. 4-4 complete with self-loops. Every link is marked with its respective hole.

is not a candidate solution. Now consider the link $\{(1,5),u_1\}$. The resulting closed walk is given by $\Big\{\{(1,2),u_2\},\{(2,5),u_2\},\{(1,5),u_1\}\Big\}$. Since disasters u_1 and u_2 disconnect S and T, $\{u_1,u_2\}$ is a candidate solution. Enumerating over all nodes and edges in K and finding the minimum cardinality candidate solution solves the GMCCD problem (in this example, a solution is given by $\{u_1,u_2\}$).

Theorem 5. *The algorithm described in steps 1-3 finds a solution to the GMCCD problem.*

Proof. In step 1 we identify a polynomial sized set of locations such that we can find a geographic min-cut considering only holes placed at these locations. Once these locations have been identified the correctness of steps 2 and 3 follow from [18]. □

Let M be the set of nodes in K. As a result of Euler's formula $|N|-|E|+|M|=2$ we know $|M|$ is polynomial in $|N|$. Since the algorithm considers a closed walk for every node-link pair in K, we know the algorithm runs in polynomial time in $|N|$.

4.3.3 Numerical Results

We used the algorithm presented in the previous section to solve the GMCCD problem for a major network provider [46]. We replace every link intersection with a node in this network which allows our algorithm to be applied. All distance units mentioned here are in longitude and latitude coordinates (one unit is approximately 60 miles)

113
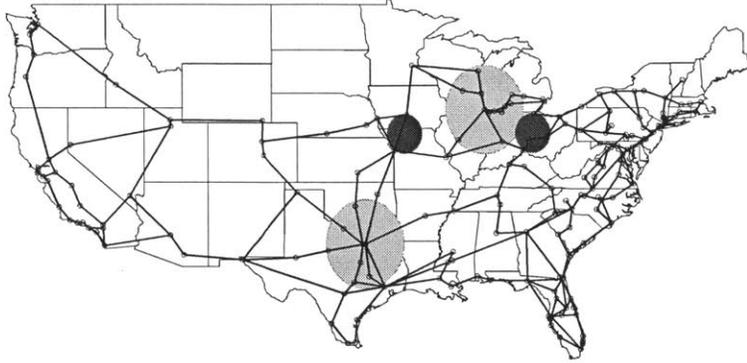
Figure 4-7: A solution to the GMCCD problem when $r_b = 1.3$, $r_p = 3.0$, $S = $ Dallas, and $T = $ Chicago. The gray disks (red disks in the online color version) represent the hole locations and the light gray disks (yellow disks in the online color version) represent the protected zones. Only two disasters, located at 'choke' points to the east and west of Chicago, are required to disconnect these cities. So, a network designer trying to increases robustness of the Chicago-Dallas connection may consider laying additional fiber down from St. Louis to western Tennessee.
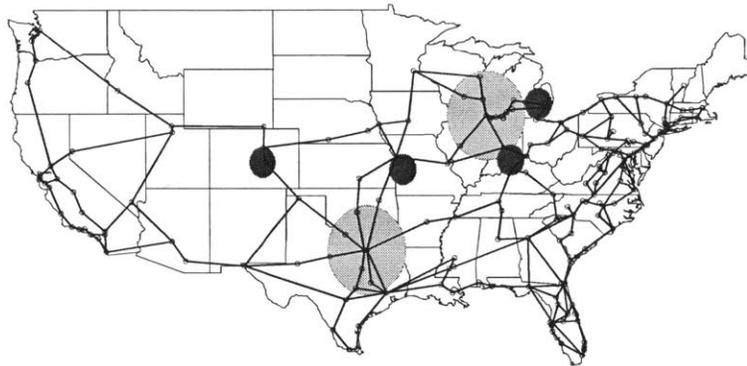
Figure 4-8: A solution to the GMCCD problem when $r_b = 1$, $r_p = 3.0$, $S = $ Dallas, and $T = $ Chicago. The gray disks (red disks in the online color version) represent the hole locations and the light gray disks (yellow disks in the online color version) represent the protected zones. Note four disasters with $r_b = 1$ are required to disconnect the two cities, whereas only two disasters are required with $r_b = 1.3$ (see Fig. 4-7).

and for simplicity we assume latitude and longitude coordinates are projected directly to $[x, y]$ pairs on the plane.

Fig. 4-7 shows a solution to the GMCCD problem when $r_b = 1.3$, $r_p = 3$, $S = $ Dallas, and $T = $ Chicago. Only two disasters, located at 'choke' points to the east and west of Chicago, are required to disconnect these cities. So, a network designer trying to increases robustness of the Chicago-Dallas connection may consider laying additional fiber down from St. Louis to western Tennessee. Fig. 4-8 shows that when r_b is reduced slightly to 1 a total of four disasters are required to disconnect the two cities.

Fig. 4-9 shows a solution to the GMCCD problem when $r_b = 1.3$, $r_p = 3.0$, $S = $ Los Angeles, and $T = $ New York. Only two disasters, located in Utah and New

Figure 4-9: A solution to the GMCCD problem when $r_b = 1.3$, $r_p = 3.0$, S = Los Angeles, and T = New York. The gray disks (red disks in the online color version) represent the hole locations and the light gray disks (yellow disks in the online color version) represent the protected zones. Only two disasters, located in Utah and New Mexico, are required to disconnect these cities. Note that in this particular fiber network, these same two failures disconnect every east cost city from every west coast city.

Mexico, are required to disconnect these cities. Note that in this particular fiber network, these same two failures disconnect every east cost city from every west coast city.

4.4   Geographic Max-Flow

In the context of geographic attacks and path-protection algorithms we consider the geographic max-flow problem: the maximum number of paths between nodes S and T such that no two paths can be disconnected by the same hole. The solution to this problem gives the maximum number of paths which are geographically disjoint with respect to disks of a particular radius. In other words, the solution gives the maximum number of paths between a pair of nodes such that a hole of a particular radius intersecting one of the paths does not affect the connection of the other paths. Similar to the GMCCD problem, to avoid triviality we restrict the locations of potential failure sites so that S and T cannot be eliminated with a single failure. This is analogous to the maximum node disjoint path problem where two paths are considered disjoint if they have no nodes in common except for S and T.

In this section we formulate the geographic max-flow problem and present an ILP to find its solution. We then discuss the fundamental differences between our problem and the work in [18] and present previous results that apply to our setting. Finally,

115

we develop an exact algorithm and a heuristic algorithm for the GMFCD problem and present numerical results based on real-world networks.

4.4.1 Problem Formulation

We use the network and disaster model from the last section. In the problem defined below, we assume a hole may be centered anywhere in the plane, except for inside protective disks of radius r_p centered at S and T.

Geographical Max-Flow By Circular Disasters (GMFCD) Problem: *Given a graph drawn in the plane $G = \{N, E\}$, two distinct nodes S and T, hole radius r_b, and protection radius r_p, find the maximum cardinality set of paths connecting S and T such that no hole intersects a pair of these paths.*

Let P be a set of paths from S to T. Let H be the set of all holes in the plane centered outside the open disks of radius r_p centered at S and T (centered outside the protected zone). The solution to the GMFCD optimization problem below is a geographical max-flow.

$$\max \ |P|$$

such that $\nexists h \in H$ where

$$p_i \cap h \neq \emptyset \text{ and } p_j \cap h \neq \emptyset \ \substack{\forall \, p_i \in P \\ \forall \, p_j \in P}, \ i \neq j$$

Let F denote the cardinality of a solution to the GMFCD problem.

4.4.2 ILP Formulation of GMFCD Problem

We will now present an ILP formulation of the GMFCD problem with a polynomial number of constraints. The idea for this formulation is to find paths, each with a different 'label', such that each one of these paths obeys some flow constraints and every pair of these paths is failure disjoint.

Recall that C is the cardinality of the solution to the GMCCD problem. Note that $F \leq C$ since every path in a GMFCD solution must be intersected by a hole

116

in order to disconnect the network and there exists no hole that intersects a pair of paths in a GMCCD solution. Let $A = \{1, 2, \ldots, C\}$. We use this set to limit the number of variables in the ILP formulation.

Define the following $\{0,1\}$ variables for all links $(i,j) \in E$ and for all $a \in A$:

$$x_{ij}^a = \begin{cases} 1 & \text{if } (i,j) \text{ has label } a \\ 0 & \text{otherwise} \end{cases}$$

We call link (i,j) *active with label a* if $x_{ij}^a = 1$. In the ILP constraints below we ensure that sets of active links with the same label obey flow conservation constraints.

Define the following $\{0,1\}$ variables for all nodes $i \in N$ and for all $a \in A$:

$$y_i^a = \begin{cases} 1 & \text{if there exists a node } j \text{ such that } x_{ij}^a = 1 \\ 0 & \text{otherwise} \end{cases}$$

Intuitively, y_i^a is 1 if any active links with label a have an endpoint at i. This variable allows us to write the flow constraints in (4.5) below.

Define the following $\{0,1\}$ constants for all links $(i,j) \in E$ and for all links $(k,l) \in E$:

$$\beta_{kl}^{ij} = \begin{cases} 1 & \text{if } \exists h \in H \text{ that intersects both } (i,j) \text{ and } (k,l) \\ 0 & \text{otherwise} \end{cases}$$

Effectively, $\beta_{ij}^{kl} = 1$ if links (i,j) and (k,l) intersect the same hole. If a pair of active links is intersected by the same hole then constraint (4.4) below ensures they cannot have different labels, and thus the paths they form are failure disjoint.

The solution to the ILP below is a solution to the GMFCD problem.

$$\max \quad \sum_{a \in A} \sum_{(S,j) \in E} x^a_{Sj} \tag{4.3}$$

such that

$$x^a_{ij} + x^{a'}_{kl} \le 1 \ \forall \begin{subarray}{l} a \in A \\ (i,j) \in E \end{subarray} \text{ and } \forall \begin{subarray}{l} a' \in A \\ (k,l) \in E \end{subarray} \tag{4.4}$$

$$\text{where } \beta^{ij}_{kl} = 1 \text{ and } a \ne a'$$

$$\sum_{j:(i,j) \in E} x^a_{ij} \le \begin{cases} y^a_i & \text{if } i = S \text{ or } T \\ 2y^a_i & \text{otherwise} \end{cases} \ \forall \begin{subarray}{l} i \in N \\ a \in A \end{subarray} \tag{4.5}$$

$$x^a_{ij}, y^a_i \in \{0,1\} \ \forall i \in N, \forall (i,j) \in E, \text{ and } \forall a \in A$$

Constraint (4.4) above ensures that a pair of active links with differing labels cannot be intersected by the same hole. Note that constraint (4.4) is generated only for x^a_{ij} and $x^{a'}_{ij}$ pairs where $\beta^{kl}_{ij} = 1$ and $a \ne a'$ (this reduces the total number of constraints), so there must be some offline computation done to find β^{kl}_{ij}.

Constraint (4.5) consists of flow conservation equations that ensure the total number of active links with a particular label and endpoints is either 0 or 2 except for nodes S and T (0 or 1 for nodes S and T). This ensures active links with a particular label form an ST path (or a cycle not including S or T).

The objective function in the above ILP maximizes the total number of active label-link pairs that have an endpoint at S. Since the flow constraint (4.5) above ensures an active link with label a and endpoint at S must be part of an ST path consisting of links active with label a and constraint (4.4) ensures differently labeled links do not interfere, this ILP will give us the maximum number of failure disjoint paths (i.e. a solution to the GMFCD problem). In Section 4.4.6 we obtain numerical results using this ILP and its relaxations, and compare these results to heuristic algorithms.
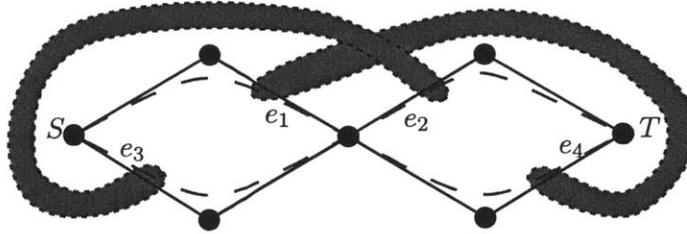
Figure 4-10: A graph and holes in the context of [18]. There are two holes shown in gray above (red in online color version). One hole intersects e_2 and e_3 and the other intersects e_1 and e_4. So the two dashed paths in above constitute a geographic max-flow in this setting. Note however that in our context, there exists a hole centered at the middle node that intersects the middle four links so $F = 1$. This example makes clear the key difference between the two settings; in our setting geographic max-flow paths must be node disjoint except perhaps for nodes located inside the protected zone.

4.4.3 Bounds on C and F

In this section we discuss fundamental differences between the GMFCD and GMCCD problems and the problems in [18] and we also present bounds on F as a function of C. We first discuss the relationship between our problems and the ones found in [18]. The max-flow and min-cut problem in [18] differs from the GMFCD and GMCCD problem in two key aspects. First, a hole in [18] need not be a disk; the only requirement is that every hole be homeomorphic to the unit disk. See Fig. 4-10 for an example of these holes. Second, in [18] holes may only be placed in a finite number of set locations (as opposed to our case where there exists an infinite number of holes outside the protected zones). This is a crucial difference because under the model of [18] some nodes or links may not be intersected by a hole. This means that it is possible for a pair of geographic max-flow paths to intersect each other. In contrast, in the context of our geometric problems, since holes can be centered anywhere on the plane outside the protected zone, we know that a pair of geographic max-flow paths must be node disjoint outside the protected zone (see Fig. 4-10).

We now present a few bounds on C and F. We first note that $C \neq F$. A simple example demonstrating this is given in Fig. 4-12 (a similar example can be found in [66] in the context of wireless networks). Note in this example $C = 2$ and $F = 1$; a geographic min-cut is given by $\{u_1, u_2\}$ and the max-flow is given by the path corresponding to the dashed curve. This is interesting as it shows the analogue to the max-flow min-cut theorem [19] does not hold in our setting. Also we know that

119

$F \leq C$ because every geographic max-flow path must be intersected by a hole in a geographic min-cut or otherwise there would remain a path from S to T after the removal of holes on the min-cut.

Once we establish the fact that we need only to consider a polynomial number of hole locations, as discussed in step 1 in Section 4.3, it follows that the GMFCD and GMCCD problems are special cases of the geographic max-flow and min-cut problems described in [18]. Thus, some of the results presented in [18] can be applied to our setting.

For example, in the special case where S and T share a common face (that is, S and T are both nodes on the same face) [18] shows that $C \leq F + 1$. Moreover, in our setting this bound is tight (i.e., can be met with equality) as demonstrated by the example in Fig. 4-12.

We now describe a family of graphs for which [18] shows that $C = F$. In order to describe these graphs we introduce some notation. Assume S and T are on a same face, denoted by B. Consider the two paths between S and T that form the boundary of B. Denote them by q and r respectively. We now define a *bad* hole with respect to face B.

Definition 13. *A bad hole with respect to face B is a hole that intersects both q and r but does not contain a curve with endpoints on q and r that only intersects faces other than B.*

For an example of some bad holes, see the holes in Fig. 4-11.

Lemma 23. *[18] If there does not exist a bad hole with respect to a common ST face, then $C = F$.*

In fact, when there is no bad hole a simple greedy algorithm is optimal. The greedy algorithm starts with path q (or r) and removes all links within r_b of path q (outside the protected zone). The common face will now be a subset of a larger face for which a new q and r are defined. We remove all links within r_b of this new path q and repeat until S and T are no longer connected.

120
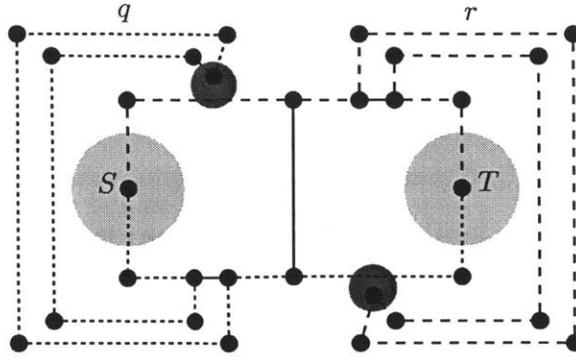
Figure 4-11: The gray disks (red disks in the online color version) above are examples of *bad* holes on this graph. Note S and T share the outer infinitely large face. Consider the paths that form the boundary of this common face, shown as the dotted path labeled q and the dashed path labeled r. Each of these holes intersects both q and r but does not contain a curve with endpoints on q and r that only intersects the inside faces, thus they are bad holes. We also note that the greedy algorithm outputs a single path, q, whereas the optimal solution is given by the two center paths which form a rectangle.
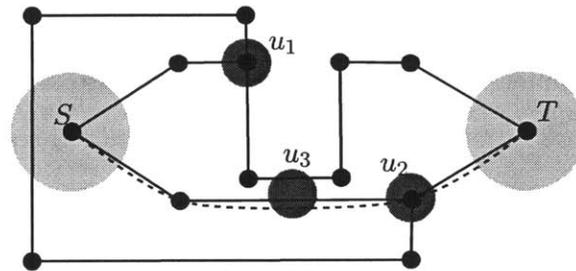
Figure 4-12: A simple network where S and T lie in the same face (a version of this example may be found in [66]). The protected zone is shown as the light gray disk (yellow disk in the online color version). All relevant holes in U are shown above in gray (red in online color version); others holes can effectively be ignored. Note $C = 2$ and $F = 1$ (a geographic min-cut is given by $\{u_1, u_2\}$ and max-flow given by path corresponding to the blue dotted curve). This shows the analogue to the max-flow min-cut theorem [19] does not hold in our setting. Also, it shows that the bound $C \leq F + 1$ (shown in [18]) is tight for our problem when S and T lie in the same face.

Lemma 24. *[18] If there does not exist a bad hole with respect to a common ST face, then greedy algorithm returns a solution to the GMFCD problem.*

It is interesting to note that the greedy algorithm is not always guaranteed to give an optimal solution when there exist bad holes (in contrast the analogous greedy algorithm always works in the non-geographic setting [21]). Fig. 4-11 shows an example of the greedy approach failing. The greedy algorithm outputs just one path whereas the optimal solution is given by the two paths that form a rectangle.

4.4.4 Exact Algorithm

Next we present an algorithm to solve the GMFCD problem exactly that works by applying a greedy routine to every ST path. We first give a brief overview of the algo-

rithm. Let p be a ST path in G. We remove every link that is not hole disjoint with p (effectively, every link outside the protected zone that intersects a 'worm' around p is removed). Denote the resulting graph by G' and let F' denote the cardinality of the geographic max-flow for G'. S and T now share a common face on G' (with a caveat described in Appendix 4.B). We will show that the greedy algorithm on G' finds the geographic max-flow for G'. Additionally, if p belongs to a solution to the GMFCD problem, we will show $F = F' + 1$ and thus p combined with the set of paths found by the greedy approach is an optimal solution. By enumerating over all ST paths, the algorithm will eventually consider a path belonging to a solution and thus solve the GMFCD problem. See Algorithm 4 for an explicit description.

Algorithm 4 Exact Algorithm to Solve GMFCD Problem

1: maxDisjointPaths $\leftarrow \emptyset$
2: **for** every ST path p **do**
3:     **call greedyRoutine**(p)
4: **return** maxDisjointPaths
Procedure greedyRoutine(p)
5: disjointPaths $\leftarrow p$
6: $G' \leftarrow G$ except for links that intersect a hole that intersects p
7: **while** S and T in same component of G' **do**
8:     **call removeQ**
9: **if** |disjointPaths| > |maxDisjointPaths| **then**
10:     maxDisjointPaths $\leftarrow$ disjointPaths
Procedure removeQ
11: $\{q, r\} \leftarrow ST$ paths that form the boundary of the new face
12: disjointPaths $\leftarrow$ disjointPaths $\cup \, q$
13: $G' \leftarrow G'$ except for links that intersect a hole that intersects q

---

We now present a few lemmas which help prove that Algorithm 4 solves the GMFCD problem.

Lemma 25. *G' contains no bad holes with respect to the new common face S and T share.*

A proof can be found in Appendix 4.C.

Lemma 26. *If p is a path in a solution to the GMFCD problem, then $F' + 1 = F$.*

Proof. G' necessarily contains all the other disjoint paths in the solution to the GM-FCD problem because only links not hole disjoint from p were removed from G. Since p is a single path, we have $F' + 1 = F$. □

Theorem 6. *Algorithm 4 finds an exact solution for the GMFCD problem.*

Proof. A path in a solution to the GMFCD problem will be considered by Algorithm 4 since every ST path is enumerated. Let p denote one of these paths. By Lemma 25 we know G' contains no bad holes with respect to the new common face S and T share. So, by Lemma 24 the greedy algorithm obtains a geographic max-flow for G'. Since p is assumed to be in the solution, by Lemma 26 we know $F' + 1 = F$. Thus, path p combined with the result of the greedy algorithm on G' is a solution to the GMFCD problem. □

This algorithm may not be practical since typically the number of ST paths grows exponentially with the size of a graph, however it gives insight to the development of a good heuristic algorithm.

4.4.5 Heuristics

The basis of the heuristic algorithm presented here is to try to identify the paths that are likely to be in the geographic max-flow. The algorithm works similarly to the exact algorithm above except we apply the greedy routine to a subset of paths, instead of every ST path. In particular, the subset of paths considered consists of those found by a standard (node disjoint) max-flow algorithm on the original topology. We apply the greedy routine on every one of these paths and return the largest set of disjoint paths found. In the next section we provide some numerical results using this heuristic. See Algorithm 5 for an explicit description.

4.4.6 Numerical Results

Similar to Section 4.3.3, we discuss the results of our developed algorithms for the GMFCD problem when applied to a major network provider [46].

Algorithm 5 Heuristic Algorithm For GMFCD Problem

1: maxDisjointPaths ← ∅
2: P ← max-flow ST paths (non-geographic)
3: **for** every path $p \in P$ **do**
4:     **call greedyRoutine(**p**)**
5: **return** maxDisjointPaths
Procedure greedyRoutine(p**)**
6: disjointPaths ← p
7: G' ← G except for links that intersect a hole that intersects p
8: **while** S and T in same component of G' **do**
9:     **call removeQ**
10: **if** |disjointPaths| > |maxDisjointPaths| **then**
11:     maxDisjointPaths ← disjointPaths
Procedure removeQ
12: $\{q, r\}$ ← ST paths that form the boundary of the new face
13: disjointPaths ← disjointPaths ∪ q
14: G' ← G' except for links that intersect a hole that intersects q

---

Fig. 4-13 shows a result of the GMFCD heuristic algorithm. The four disks represent hole locations in a geographic min-cut. The four 'worms' correspond to hole disjoint paths found using the GMFCD heuristic algorithm. Since the cardinality of the geographic max-flow and min-cut solutions is the same and $F \leq C$, we know the heuristic has found an optimal solution to the GMFCD problem.

If there is no restriction on S and T belonging to the same face, it is known $C \leq 2F + 2$ [18]. We believe the disjointness of geographic max-flow paths in our setting allows for this bound to be tightened. We conjecture that $C \leq F + 1$ in our setting. Using the algorithm in Section 4.3 and running CPLEX on the the ILP in Section 4.4.2, we solve the GMCCD and GMFCD problems exactly for 1000 randomly generated graphs consisting of 13 nodes. We found $C = F$ for 99% of the instances and $C = F + 1$ for the remaining 1%. There was not a single example where C exceeded F by more than 1, thus supporting our conjecture.

4.4.7 Complexity of the GMFCD Problem

The max-flow problem in [18] is shown to be NP-hard, however, the proof does not directly transfer to our setting since in our setting geographic max-flow paths cannot intersect outside protected zones. We believe a polynomial time solution may be
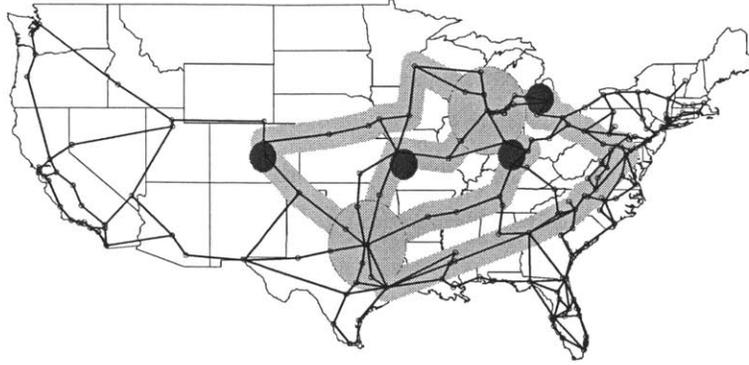
Figure 4-13: Result of GMFCD heuristic algorithm when $r_b = 1.0$, $r_p = 3.0$, $S = $ Dallas, and $T = $ Chicago. The four gray disks (red disks in the online color version) represent the hole locations in a geographic min-cut and the light gray disks (yellow disks in the online color version) represent the protected zones. The four light gray 'worms' (teal 'worms' in the online color version) correspond to hole disjoint paths found using the heuristic algorithm. Since the cardinality of the max-flow and min-cut solutions is the same and $F \leq C$, we know the heuristic has found an optimal solution to the GMFCD problem.

possible and this is a subject of future work.

4.5   Conclusions and Future Work

Motivated by applications in the area of network survivability, in this chapter we present a geographic max-flow and min-cut problem where failures, modeled as disks, may be placed anywhere in the graph except for certain protected zones. We show these problems can be reduced to discrete ones and present a polynomial time algorithm for the GMCCD problem based on ideas from [18] and chapter 2. We then develop an ILP formulation, an exact algorithm, and a heuristic algorithm for the GMFCD problem. Using these algorithms, we obtain numerical results for an example backbone network, thereby demonstrating the applicability of our algorithms to a real-world network.

Our approach provides a way to look at network survivability in the face of multiple disasters or attacks that takes into account the geographical correlation between links. Some future directions include application of this approach to the electric power transmission network, finding a tight bound on the difference between geographic min-cut and max-flow (i.e. the analog to the max-flow min-cut theorem), and the development of network design tools (e.g. how to build a network under some constraints such that geographic min-cut is maximized).

125

4.A   Details of Step 1

There are an infinite number of hole locations centered outside the protective disks; in this step we find a polynomial sized set of candidate holes for the GMCCD problem. We first make a note about holes. Let h and h' be holes such that h' intersects every link h does in addition to possibly other links. We note that if h belongs to a set of holes that disconnects S and T, then replacing h with h' will still result in S and T being disconnected (if S and T are disconnected, removing additional links will also leave them disconnected).

We now describe how to find a set of potential failure locations for the GMCCD problem. First, we apply an algorithm from chapter 2 where the graph and disaster model is the same except that there exists no protected zones. The algorithm creates a polynomial sized set of holes, denoted by A, such that for every hole in the plane there exists a hole $a \in A$ that intersects at least the same set of links. Note however, one of these holes may be centered inside one of the protective disks around S or T (see Fig. 4-3), and so we must consider additional holes to solve the GMBCD problem. Let A' be the set of disks in A not centered in the protected zone.

Let h be a hole not centered in the protected zone and let a be a hole in A that intersects at least the same links as hole h. We will show there exists a polynomially sized set of holes centered outside the protected zone, denoted by M, such that if hole a is in the protected zone, then a hole $m \in M$ intersects at least the same links as hole h. So, for every hole not centered in the protected zone there exists a hole in $A' \cup M$ that intersects at least the same set of links. Thus, $A' \cup M$ gives us a polynomial sized set of candidate holes for the GMBCD problem.

In the following we present notation that allows us to describe the locations of the holes in M. Let O^S denote the circle centered at S with radius r_p and let O^T denote the circle centered at T with radius r_p. Let e denote a link and let $\overleftrightarrow{e}$ denote the line that contains e. Let $\partial H(e, r_b)$ denote the boundary of $H(e, r_b)$.

Lemma 27. *Let M be the set of all holes of radius r_b centered on at least one of the following sets:*

1. $\bigcup_{e \in E}\{O^S \cap \partial H(e, r_b) | e$ does not have an endpoint at S or $r_p \neq r_b\}$,

2. $\bigcup_{e \in E}\{O^T \cap \partial H(e, r_b) | e$ does not have an endpoint at T or $r_p \neq r_b\}$,

3. $\bigcup_{e \in E}\{O^S \cap \overleftrightarrow{e} | e$ has an endpoint at S and $r_p = r_b\}$,

4. $\bigcup_{e \in E}\{O^T \cap \overleftrightarrow{e} | e$ has an endpoint at T and $r_p = r_b\}$.

Assume there exists a polynomially sized set of holes A such that for every hole in the plane there exists a hole $a \in A$ that intersects at least the the same links. Given a hole h not centered in the protected zone, if a hole $a \in A$ is in the protected zone and intersects at least the same set of links as h, then there exists a hole $m \in M$ that intersects at least the same links as h.

Proof. Let β denote the center of hole h and α denote the center of a. Let Z denote the set of all links that intersect h. Note that $x \in \cap_{z \in Z} H(z, r_b)$ iff x is the center of a hole that intersects every $z \in Z$. So $\beta \in \cap_{z \in Z} H(z, r_b)$ and $\alpha \in \cap_{z \in Z} H(z, r_b)$. Also note that $\cap_{z \in Z} H(z, r_b)$ is convex and thus path-connected, so there exists a path p in $\cap_{z \in Z} H(z, r_b)$ from α to β that necessarily intersects O^S or O^T. Let y be this intersection point. The hole centered at y must necessarily intersect at least the same links as h since $y \in \cap_{z \in Z} H(z, r_b)$.

W.l.o.g. assume y lies on O^S. Note every point in the sets (i) and (iii) above lies on O^S. Let y' be the first point in (i) or (iii) in the clockwise direction from y on O^S. Now, $y' \in \cap_{z \in Z} H(z, r_b)$ since every hippodrome y intersects is also intersected by y'. Therefore a hole centered on (i) or (iii) must intersect every link that intersects h. $\qquad \square$

Since for each link we consider at most eight holes ($|O^S \cap \partial H| \leq 4$ and $|O^T \cap \partial H| \leq 4$ under conditions above), M is of polynomial size. Since $|A'|$ is polynomial, the set of potential holes for the GMCCD problem, $A' \cup M$, is of polynomial size. For the remainder of the section let $U = A' \cup M$.
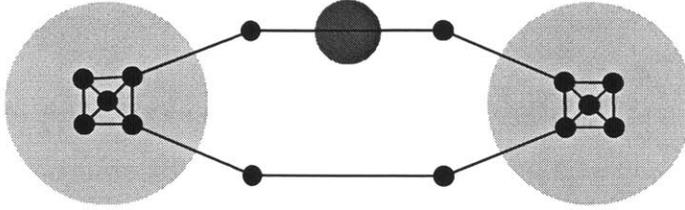
Figure 4-14: No hole intersects the nodes located in the protected zone, so removing links not hole disjoint with some path will never result in a graph where S and T are on the same face.

4.B   Modifying G

Here we discuss a caveat for applying Algorithms 4 and 5. We first note that removing every link not hole disjoint with some ST path does not ensure S and T lie on the same face. See Fig. 4-14 for an example. Here we note how to modify G so that when links around a ST path are removed S and T are guaranteed to share a common face. Let O be a circle of radius $r_p - r_b$ centered at S (assume $r_p - r_b \geq 0$). Note no hole may intersect a link anywhere within this circle. Place nodes everywhere O intersects a link. Consider the links forming paths that lie entirely within O that have endpoints at S and these new nodes. Replace these links such that there exists a path inside O from S to each of these new nodes such that these paths do not intersect each other except at S. Repeat this process for T. Now when links within r_b of a particular path are removed (outside the protected zone) S and T are guaranteed to be on the same face. Since the removed links do not intersect any hole and since the connectivity of G is unchanged outside the protected zone, the solution to the GMFCD problem remains the same. We assume the exact algorithm is applied after this modification.

4.C   Proof of Lemma 26

Proof. Let B denote the new common face S and T share after the removal of links not hole disjoint with path p. Let q and r be two ST paths that form the boundary of face B (see Fig. 4-15).

We will use proof by contradiction. Assume there exists a bad hole with respect to B, denoted by h. Hole h must intersect both q and r but not p (because all links intersecting a hole that intersects p are removed). Thus, h must contain a curve with
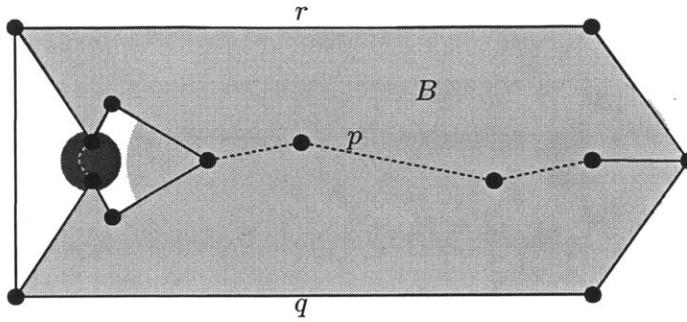
Figure 4-15: The dashed links represent path p. Every link not hole disjoint with path p is removed. The grey face (teal face in the online color version) B above represents the new common face S and T share. q and r are the two ST paths that form the boundary of B. If a hole intersects both q and r, it must not intersect p (because all links intersecting a hole that intersects p are removed). This hole must then contain a curve with endpoints on q and r that only intersects faces other than B (shown as the grey dashed curve above).

endpoints on q and r that only intersects faces other than B (see Fig. 4-15). So, h is not a bad hole, a contradiction. □

Chapter 5

Power Network Reliability Problems

In this chapter, we use our developed tools to consider the effect of geographically correlated failures on power transmission networks. Similar to fiber infrastructures, power transmission networks are vulnerable to large-scale natural disasters or attacks, such as hurricanes or geomagnetic storms [3, 22, 59]. The geographical layout of the network affects the impact of such real-world disasters since they occur in specific geographic locations. For example, a hurricane or earthquake can cause failure of electric power lines that directly transmit power to a large city, thereby likely causing significant disruptions to power services. However, the damage to the power network infrastructure is not necessarily limited to these initial failures; power networks are also vulnerable to *cascading failures*. Cascading failures occur when an initial failure in the network changes power flows, which must obey physical law constraints, such that additional lines overload and fail. This in turn causes the power flows to change again; this process will continue until some stability is reached. A well known example of a cascading failure is the 2003 blackout where a significant area of the northeastern U.S. lost power [6]. In this chapter we consider two failure models. The first model considers power networks with respect to a randomly located geographic disaster and subsequent cascading failures. The second model builds on the first; we describe a dependency between power and data networks and consider the connectivity of data

131

networks in this context. For each model, we present numerical results based on real-world networks.

5.1   Overview of Models and Related Work

We initially consider a two-stage failure model for power networks. The first stage removes power lines that intersect a randomly located disk. The second stage then calculates the cascading failure that occurs due to the removal of the initial links. By using the tools developed in chapter 3 and the cascading failure model presented in [16], we are able to calculate the effect of this type of failure in power networks. To the best of our knowledge, [14] is the only other work to look at the effect of geographically correlated failures on power networks.

Then motivated by the effects of power loss on data networks [30], in the final part of chapter 5 we consider the survivability of data networks with respect to power networks. We assume data nodes rely on the operation of the closest power load nodes to function. We present numerical results that show data network connectivity is significantly lower when power network dependency is considered; this implies power network effects have a significant impact on the survivability of real-world data networks.

We now discuss some related work. Power network resilience has been considered in the past [9, 17], however so far only [14] has considered the effects of a targeted geographic failure model. In this chapter we consider the effect of non-targeted geographic attacks on the power network. Some recent work has modeled the interdependence between data and power networks and demonstrated asymptotic percolation results [23]; however they did not consider power flows or geography in their models. Additionally, [63] considered a geographic dependence model but did not consider failures which were geographically correlated.

5.2 Assessing Power Network Reliability

We now consider a geographic failure model for power networks where a disaster is modeled as a 'randomly' located disk. This can describe the effect of some natural disasters such as geomagnetic storms [3, 22, 59] or hurricanes, in addition to collateral (non-targeted) damage from attacks on other continental networks (e.g. an attack on the communication or transportation networks). Our goal is to be able to understand and quantify the effect of large-scale non-targeted disasters and their resulting cascade effects on the power network. We first describe the network and failure model and then propose metrics to be evaluated on a real-world network.

5.2.1 Network and Failure Model

The network model remains the same as in the previous chapters; we consider a network such that nodes are represented by points on the plane and links are represented by line segments.

The failure model consists of two stages; the first stage is link failures caused by the random circular disaster and the next stage is the resulting cascading failures. We first describe the initial failures caused by the random circular disk (which is the same as the failure model presented in chapter 3). We model a disaster event in the network as a *single* randomly located disk of a radius r_b centered within an area of interest C (i.e. C is a set of points in the plane where the disaster may be centered). If the randomly located disk intersects some power lines, we assume those lines are destroyed.

Geometric probability [65] allows us to assign a measure to sets of disks. As in chapter 3, this measure is simply defined as the Lebesgue integral over the set of disk centers. Using this measure and tools from computational geometry, we can find the probability a randomly located disk that intersects C also intersects some set of links. See Fig. 5-1 for a simple example and Fig. 5-2 for an example with respect to the Italian high-voltage electrical transmission network (HVIET).

After this initial failure, due to power flow constraints, a cascading failure may
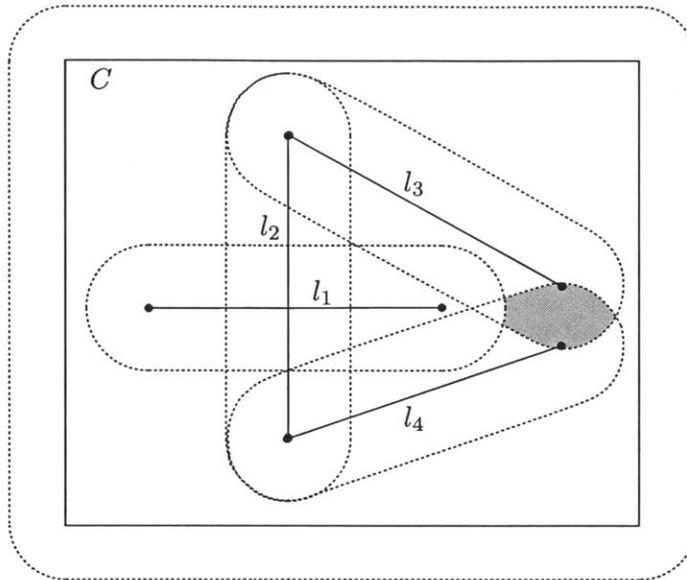
Figure 5-1: The probability that a randomly located disk centered in C intersects only l_3 and l_4 is given by the ratio of the area of the shaded region to the area of the large rounded rectangle.
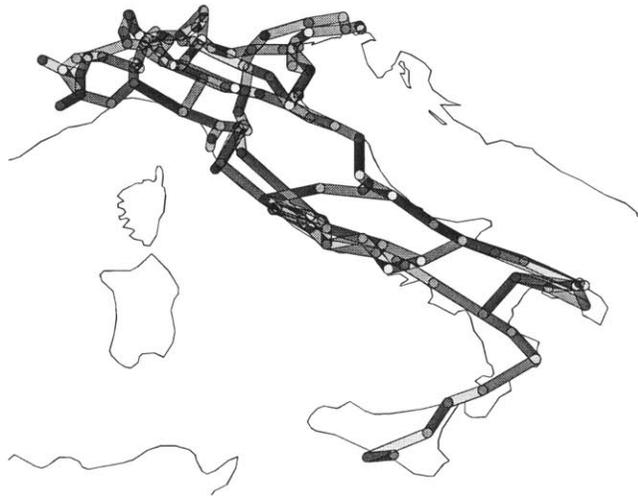
Figure 5-2: Every shaded region above represents a set of disk centers whose radius is ≈ 8 kilometers and only intersects a particular set of power lines should a failure be centered within that region. The network being represented is the Italian high-voltage electrical transmission network (HVIET) [63, 64].

occur. We will use the same power flow and cascading failure model described in [16]. Thus these geometric probability tools along with the cascading failure model allow us to analyze the effects of large scale randomly located disasters on the power network.

We now present our failure model for power flows and cascading failures in power networks. We use the same models as found in [14,16] and even borrow some notation. The details of the DC power flow and cascading model may be skipped and the reader may proceed to Section 5.2.2 without loss of continuity.

DC Power Flow Model

We now describe the DC power flow model which is a linearized version of the more complicated AC power flow model. We use the DC model because it is more tractable and easier to find solutions for power flows.

Let β_i represent the amount of power injected at node i. If $\beta_i > 0$ then node i is a source of power and may represent a generator where power is injected into the system. If $\beta_i < 0$ then node i is a sink of power and may represent demand at this node. We call these type of nodes power demand nodes. If $\beta_i = 0$ then power is neither injected or removed at node i and may represent a power bus. Let N be the set of nodes in the network.

Let (i, j) denote the power line from node i to node j and let E denote the set of all these lines. Let x_{ij} denote the reactance of (i, j) and let u_{ij} denote the capacity of (i, j).

A DC power flow can be described by the amount of power to flow from node i to node j on (i, j), denoted by f_{ij}, and the phase angle at node i, denoted by θ_i. A DC power flow must obey the following constraints.

$$\sum_{j:(i,j)\in E} f_{ij} = \beta_i \ \ \forall i \in N \tag{5.1}$$

$$\theta_i - \theta_j = x_{ij} f_{ij} \ \ \forall (i, j) \tag{5.2}$$

Equation (5.1) constrains the total power out of a node to be equal to the amount

of power injected at that node (power conservation). For example, if a node is a generator then the net power flow out must be the amount of power generated at that node. Equation (5.2) is the analogue to Ohm's law; the amount of power through a power line is proportional to the difference in phase angles θ_i and θ_j.

It should be noted that the power flow has a feasible solution as long as $\sum_{i \in K} \beta_i = 0$ for every connected component K in the network (that is, aggregate supply equals aggregate demand for that component) [16]. Additionally, the values of the power flows are unique [16].

Cascading Failure Model

We now describe the cascading failure model. Again, this model can be found in [14, 16], but is presented here for completeness.

Before any failures occur, we assume the network is connected and that $\sum_{i \in N} \beta_i = 0$. In other words, we assume aggregate demand is equal to aggregate supply.

We now describe the cascading failure model in steps.

1. Set $\tilde{f}_{ij}$ to be the absolute value of the power flow on (i, j) before any failure occurs.

2. Consider some subset of power lines to be initially removed from the network.

3. In order to calculate DC power flows for this modified network, aggregate supply and demand must match in each component. Hence, we proportionately reduce supply (or demand) at nodes in each component until this condition is met. This may model load shedding or a ramping down of generators.

4. Power flows f_{ij} are then calculated for the remaining lines.

5. Let $\tilde{f}_{ij} = \alpha |f_{ij}| + (1 - \alpha) \tilde{f}_{ij}$.

   $\tilde{f}_{ij}$ represents some 'moving average' of flow through the power line (i, j) and can be thought of as modeling of some thermal effects. α is a parameter in this moving average set to a value between 0 and 1. If α is small, then the line will take more time steps to 'heat up'; if $\alpha = 1$ then the line can be thought of

as feelings the effects of the new flow instantaneously. In this work we assume $\alpha = 0.5$.

6. We then remove all lines for which $\tilde{f}_{ij} > u_{ij}$. This may cause an additional change in the power flows (hence the cascade); we go back to step 3 and the process repeats until no flow is above capacity.

It should be noted that we were not able to attain the capacities of power lines for real power networks. Hence, in order to approximate the capacities on a power network we calculate the initial power flows on each line and then set u_{ij} proportional to $|f_{ij}|$ before any failures occur. This proportion is called the Factor of Safety (FoS) and relates to the amount of 'spare capacity' on the power lines. In other words $u_{ij} = |f_{ij}| \times FoS$ before any failures occur. For real power grids, it is believed that a good approximation for FoS is 1.2 [14]. Hence, for the majority of this work, we assume $FoS = 1.2$.

5.2.2 Performance Metrics and Numerical Results

Our goal is to analyze the effect of a randomly located circular disk failure in conjunction with cascading failures on power networks. Let the yield be the fraction of demand satisfied after the disaster and resulting cascade. By calculating the probabilities of relevant joint link failures using the tools and equations in section 3.9 and considering the resulting cascading effects, one can evaluate the expected value as well as the distribution of the yield to a randomly located disk failure event.

We now discuss some numerical results based on the HVIET network[1]. Fig. 5-3 shows the cumulative distribution function (CDF) of the average yield on the HVIET network with disaster radius of 50 kilometers. Addressing the effect of Factor of Safety, Fig. 5-4 shows how average yield changes as the factor of safety (FoS) is changed (Factor of Safety relates to the amount of 'spare capacity' on power lines). Note when $FoS = 1$, then there is no spare capacity allocated on the power lines, so when a failure event occurs the resulting cascading failure brings down most of

---

[1]We would like to thank the authors of [63,64] for sharing their data.

CDF of Yield on HVIET network with radius≈50km, FoS=1.2, and α=0.5, Average Yield = 0.78336
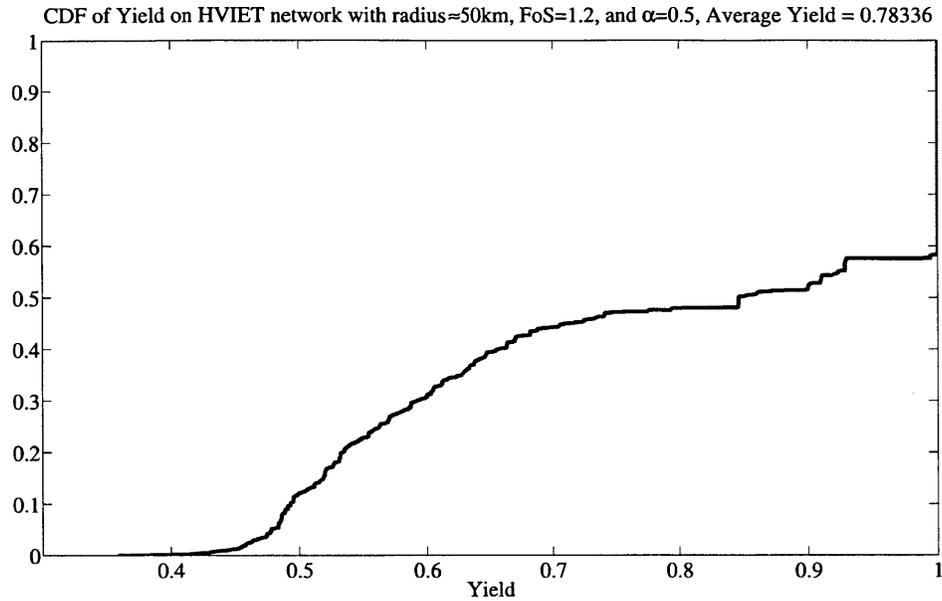

Figure 5-3: CDF of the average yield on the HVIET network with disaster radius of approximately 50 kilometers. We assume that the region of interest is given by the convex hull of the network. Note that there is a significant probability the yield is 1; this is mainly caused by disks centered within the region of interest that do not intersect the network.

the network. As FoS increases, the amount of spare capacity on the power lines increase, so the average yield increases as well, as one would expect. For example, when $FoS = 2$ the failure event will not have much effect on the yield. Addressing the effect of the radius of the disaster, Fig. 5-5 shows as the radius of the initial disaster increases, the average yield in the network decreases.

We now compare the effect of independent random link failures to the effect of a randomly located circular failure. We initially calculate the average yield of HVIET to a circular disaster while the size of the region of interest C varies. The size of C is varied to change the probability a unit of fiber is cut. So we can plot average yield versus the probability a unit of fiber is cut. See Fig. 5-6 for results.

Next, we calculate average yield assuming independent link failures such that links fail with the same probability as in the random disk-cut case. Thus the probability a link fails is still a function of its length, however links fail independently. Since the total number of power lines is not small, calculating average yield by enumerating all possible failures is not feasible (possible failures are exponential in number of links). Instead we use a Monte Carlo approach, using 4000 samples for each particular
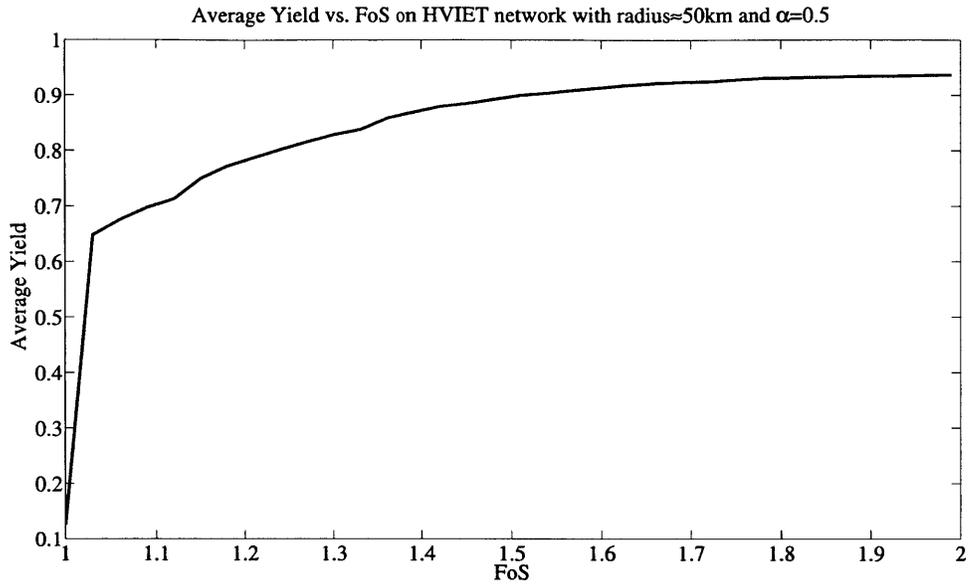
Figure 5-4: Average yield vs. FoS on the HVIET network with disaster radius of approximately 50 kilometers. When the $FoS = 1$, then there is no spare capacity allocated on the power lines, so when a failure event occurs the resulting cascading failure brings down most of the network. As FoS increases, the average yield increases as well, as one would expect. Note when $FoS = 2$, then the failure event will not have much smaller effect on the yield.
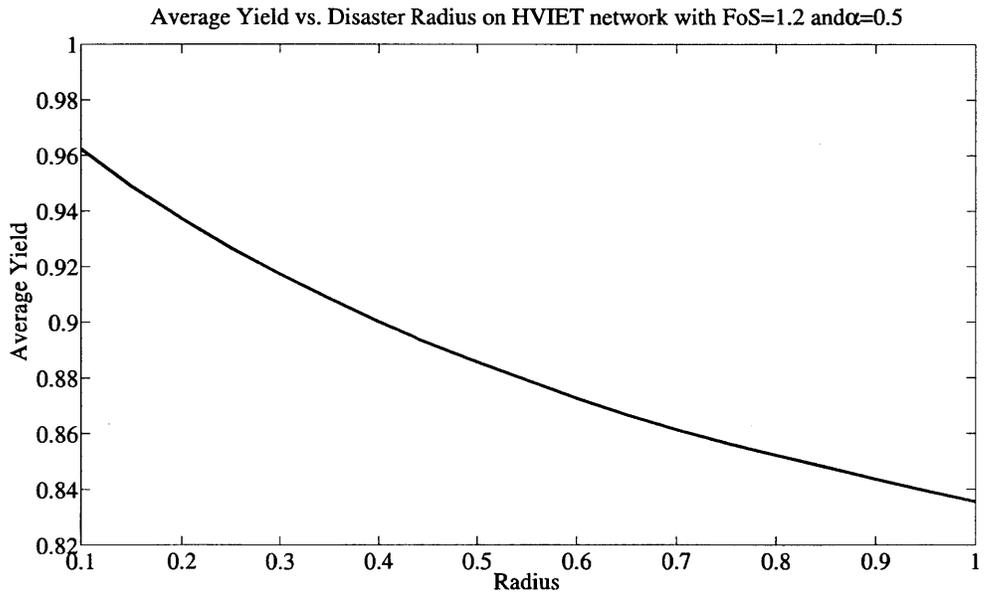
Figure 5-5: Average yield vs. radius (in terms of degrees of latitude/longitude) on the HVIET network. As the radius of the initial disaster increases, the average yield in the network decreases.

139

Average Yield vs. Probability of Unit Link Failure for radius≈8km, FoS=1.2, and α=0.5
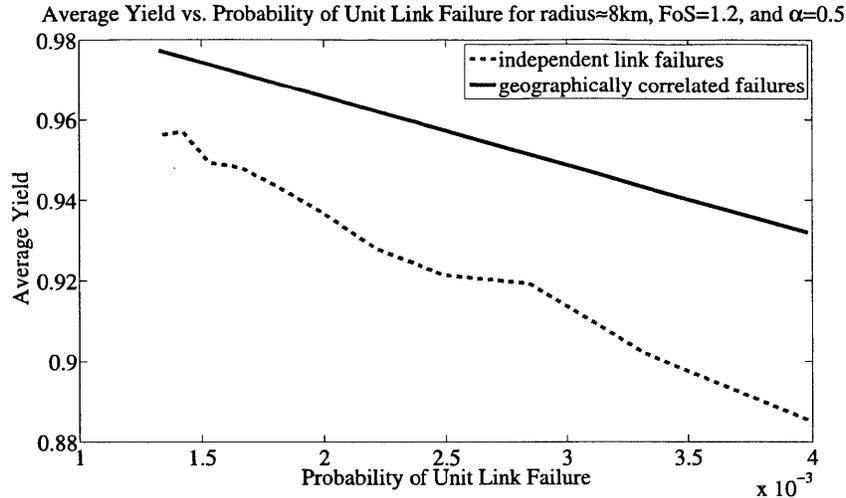
Figure 5-6: The solid line shows average yield in HVIET versus the probability a unit (latitude/longitude) of power line is cut by a random disk of radius approximately 8km. The dashed line shows average yield in HVIET assuming power lines fail *independently* such that lines fail with the same probability as in the random disk case.

probability of unit link failure sample point. See Fig. 5-6 for results.

Notice that average yield under independent failures is less than in the case of random disk-cuts. This result demonstrates geographic disasters on power networks have key differences from independent smaller scale failures (e.g. power line failure due to brush growth). Perhaps this is because some power supply nodes and power demand nodes are near each other and so a random disk may be more likely to effectively remove both these nodes simultaneously which may reduce the chances of a large cascading failure (since power loads will remain balanced). Also note the contrast to the result in Fig. 3-25 for the NSFNET data network where independent failures have less impact than in the case of random disk-cuts; this highlights a fundamental difference in the survivability between power and data networks.

5.2.3 Possible Extensions

In the context of random geographic failures and power networks, the following problems are potential extensions for future work:

1. **Other metrics beyond yield**

   Consider other metrics beyond yield such as the distribution of number of lines destroyed or the distribution of connected components. These distributions will

allow us to better understand the impact of a random geographical disaster on the survivability of the power grid.

2. **Computationally efficient algorithms**

   Development of efficient algorithms to calculate the yield in general networks that scale well with network size. Analyzing the running time of our current algorithms and developing faster methods will allow us to obtain numerical results on larger and more detailed real-world power networks.

3. **Extending the probabilistic failure model**

   Currently, our model assumes that every power line intersected by a circular attack is removed from the network. However, power lines within a disaster region may not always fail (e.g. shielded power lines near a hurricane may remain operational). So the attack may have a probabilistic effect on the lines. It would be interesting to capture this doubly random effect; we model a disaster as a randomly located disk that also has a non-deterministic effect on the intersected power lines.

4. **AC Power Flow Model**

   A more realistic power flow model can be considered. Currently, many papers on power networks assumes a DC power flow model [17] [14]; this type of model is very simple and ignores certain effects that may occur during a cascade. The AC power flow model is a more realistic flow model, though it is harder to solve for the flow equations [56]. We can alter our failure model to incorporate the more realistic AC power flow model and study the impact of the cascading model on yield and other performance metrics.

5. **Robust Design**

   In addition to the above items, we can study some power network design issues. One goal may be to increase the average yield in the network under a random circular disk disaster. To this end, we may consider how to add additional

power lines or increase capacities of certain power lines in order to increase average yield. For example, we may consider what Factor of Safety is required to guarantee the expected yield above a certain threshold.

5.3 Design of Infrastructure Robust To Power Failures

Many systems and networks depend on reliable delivery of power from the electric grid. For example, power is required to operate street lights for transportation networks in cities. Another example are fiber networks; power is needed at backbone routers and amplifiers (on fiber links) or else those components will fail. Since cascading power outages can be widespread, their effect on dependent systems can be devastating. In particular, due to the widespread nature of blackouts, continental fiber networks may become disconnected if the power failure affects a large area that includes the networks physical components. For example, the blackouts of 2003 had a significant effect on the connectivity of the Internet [30].

Motivated by the dependencies of many networks and systems on the power network, we consider the design of robust infrastructures with respect to cascading power failures caused by a randomly located geographic attack. We first describe a model for the dependence of a network on the power network. We then present our failure model and compare data network reliability with and without power network dependency. We then close by proposing an example problem formulation for the design of robust data networks with respect to power outages.

5.3.1 Dependence on Power Network

As described above, many networks and systems require power to operate properly; that is, failure to provide power to systems can cause failure in those dependent systems. Although these systems typically have backup power supplies, backup generators are often unreliable. We assume, as in the previous section, that the power
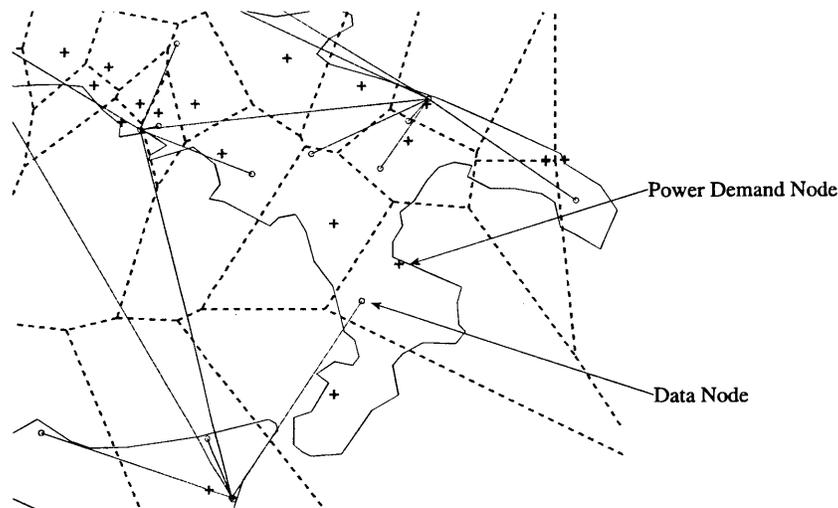
Figure 5-7: Part of the backbone of the Italian research network (GARR) [63, 64] is shown above by solid line segments representing links and circles representing nodes. The dashed segments represent the Voronoi cells based on the locations of power demand nodes, shown by crosses above, in the Italian high-voltage electrical transmission network (HVIET) [63, 64]. Our model assumes that data nodes extract power from the closest power demand node; when a demand node fails, data nodes located within its Voronoi cell are assumed to fail as well.

network is represented by points and line segments in the plane. Similarly, we assume the dependent network is also modeled by points and line segments. A dependent node is likely to draw its power from a nearby substation. So, we let a dependent node be operational if the closest (in a Euclidean sense) power demand node is still delivering power (that is $\beta_i < 0$ for node i). Thus, based on the locations of demand nodes in the power network, we can construct a Voronoi diagram; a dependent node in a particular Voronoi cell will depend on the operation of the supply node corresponding to that cell. See Fig. 5-7 for an example.

5.3.2 Failure Model

We use the same failure model for the power grid presented in the previous section augmented with data-power network dependency. This failure model consists of three stages; the first stage is link failures caused by the random circular disaster and the next stage is the resulting cascading failures in the power network. Then, the effects on the dependent network (based on geographical proximity to supply nodes) are

considered once the cascading failures have occurred.

5.3.3 Metrics for Dependent Network Robustness

Our goal is to design a network robust to failures in the power grid. In the context of a random geographic failure on the power grid and the resulting impact on dependent networks, we propose to consider the following metrics:

- **Connectivity** - In many networks, especially data networks, we are concerned with connectivity; i.e. does the network remain connected. For example, we would like for all major U.S. cities to be able to communicate with each other, therefore it is reasonable to consider the connectivity of the continental fiber network. Thus, we can consider the probability that the dependent network remains connected after a random attack on the power grid.

- *ATTR* - If a connected network cannot be guaranteed after a failure or full connectivity is not critically important, it may be useful to consider the *ATTR* metric. This is given by the probability a randomly chosen pair of nodes in the dependent network remain connected after a random attack on the power grid. In the following, we consider the effect of random disasters on real-world dependent networks using this metric.

5.3.4 Numerical Results

Using the failure model just described, we present some numerical results based on the Italian research network (GARR) and the Italian high-voltage electrical transmission network (HVIET) [63,64]. Consider Fig. 5-8. Via a Monte Carlo simulation, this figure shows how *ATTR* is significantly lower when power network dependency is considered; this implies power network effects have a significant impact on the survivability of real-world data networks. Fig. 5-9 shows a similar result when the connectivity metric is considered although the difference is not as significant. Perhaps this is because removing certain power demand nodes from the network causes the connectivity metric to be zero regardless if a cascading failure occurs.

ATTR vs. Radius for GARR Network
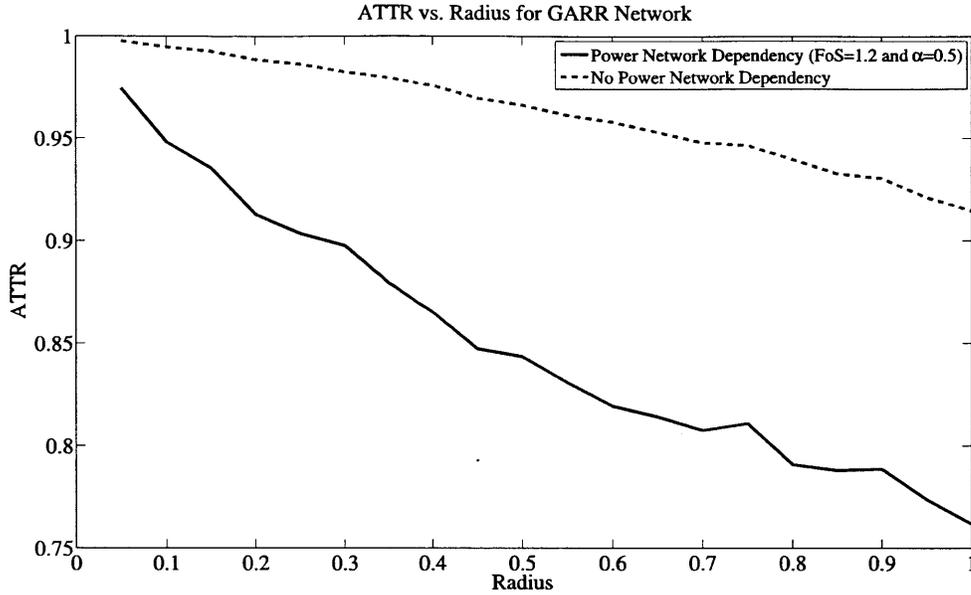
Figure 5-8: The red dashed curve shows *ATTR* for the Italian research network (GARR) as a function of the radius (in latitude/longitude coordinates) of a randomly located circular disaster *when no power networks are considered* (using tools and models from chapter 3). The blue solid curve shows *ATTR* for the GARR network when the dependency effects of Italian high-voltage electrical transmission network (HVIET) are considered. For every radius considered a Monte Carlo approach with 4000 samples was used. We note that *ATTR* is significantly lower when power network dependency is considered; this implies power network effects have a significant impact on the survivability of real-world data networks.

Connectivity vs. Radius for FoS=1.2 and α=0.5
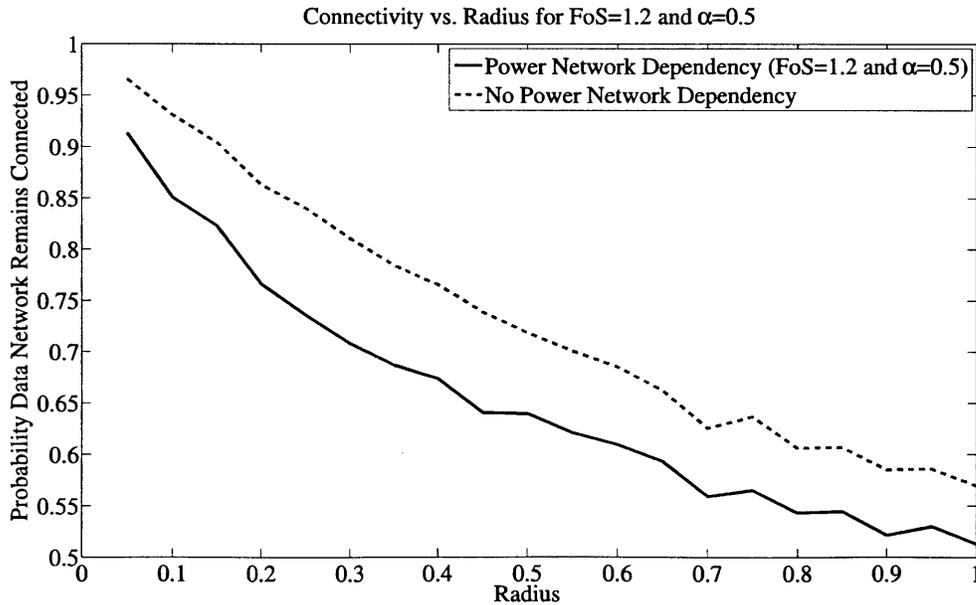
Figure 5-9: The red dashed curve shows the probability the data network remains connected for the Italian research network (GARR) as a function of the radius (in latitude/longitude coordinates) of a randomly located circular disaster *when no power networks are considered* (using tools and models from chapter 3). The blue solid curve shows the probability the data network remains connected for the GARR network when the dependency effects of Italian high-voltage electrical transmission network (HVIET) are considered. For every radius considered a Monte Carlo approach with 4000 samples was used.

5.3.5 Possible Extensions

In the context of a random geographic attack on the power grid and its effect on dependent networks, one can consider to study some network design problems. One goal may be to increase the connectivity or $ATTR$ metric in the dependent network. To this end, we may consider how to add additional power lines or increase capacities of certain power lines in order to decrease the effect of cascading failures in the power grid thereby reducing the effect on dependent networks. For example, we may consider what Factor of Safety is required to guarantee the $ATTR$ metric remains above a certain threshold in the dependent network. Alternatively, we can consider how to augment the existing dependent network so that it becomes more robust to cascading power failures. An interesting future direction would be to study the joint design of the power grid and dependent network as well as explore the tradeoffs between strengthening the power network and the dependent network.

We now discuss a design problem with respect to data networks. Suppose we wish to strengthen the connection of the data network of two major American cities under the context of random power failures caused by an attack. One problem would be to consider a maximally blackout disjoint path problem: how to find a pair of data paths with common source and destination that has the minimum probability of being affected by a blackout. The solution to this problem gives the most survivable pair of paths with respect to power blackouts. See Fig. 5-10 for an example of blackout disjoint paths.

More specifically, we assume the same failure model as before; a single randomly located disk of fixed radius removes power lines it intersects. Let R be the set of all possible initial failure events caused by the random disk. Let $r_i \in R$ and let p_i correspond to the probability of r_i (the probability a random disk causes the failure event r_i). Now assuming our two-stage failure and dependent network effects model, the problem becomes how to choose two data paths between S and T (ST paths),

Figure 5-10: Four light grey (teal in online color version) Chicago-Dallas paths are highlighted above for a particular backbone network [46]. The grey closed curve (red closed curve in online color version) surrounding eastern states represents the region of a power blackout. Note the two rightmost paths are not disjoint with respect to this blackout since those pair of data paths both intersect the blackout zone.

path$_1$ and path$_2$, such that the following expression is maximized

$$\max_{\text{path}_1,\text{path}_2\in\text{ST paths}} \sum_{r_i\in R} p_i \max\left(\mathbb{1}_{\text{path}_1 \text{ survives } r_i}, \mathbb{1}_{\text{path}_2 \text{ survives } r_i}\right).$$

This problem fits naturally into the Shared Risk Link Group (SRLG) framework presented in [44] and may provide a good starting point for analysis and heuristics.

Chapter 6

Conclusion and Future Directions

In this thesis we developed the necessary theory to evaluate network performance metrics under several geographic failure models. This allows us to begin developing some network design tools that can mitigate the effects of regional attacks such as electromagnetic pulse attacks and natural disasters. Our approach provides a fundamentally new way to look at network survivability under disasters or attacks that takes into account the geographical correlation between links. We conclude with an overview of our results and potential future directions to our work.

6.1 Conclusions and Extensions

Motivated by applications in the area of network robustness and survivability we initially focused on the problem of geographically correlated network failures. Namely, we considered graph models in which nodes and links are geographically located on a plane, and model the disaster event as a geometric object. We first looked at the properties and impact of geographical line segment cuts which are located with the intention to reduce network capacity and connectivity. We considered a simple bipartite graph that abstracts the fiber links between the east and west coasts in the U.S. or transatlantic/pacific links. Then, we considered a general graph model in which nodes are located on the Euclidian plane and studied two related problems in which cuts are modeled as line segments or as circular disks. For these cases,

149

we developed polynomial-time algorithms for finding worst-case cuts. We then used these algorithms to obtain numerical results for various performance measures. Some future research directions include the analytical consideration of arbitrarily shaped cuts, the use of computational geometric tools for the design of efficient algorithms, and how to *design* a network to mitigate the effects of these worst-case attacks.

In the next part of the thesis we focus on random geographical attacks on the network which can model the result of a natural disaster or collateral damage. In particular, we focused on random line and circular cuts. Using tools from geometric probability we demonstrate how to compute failure probabilities and show how to calculate *ATTR* and other network performance metrics in polynomial time under these failure models. This is a significant contribution because calculating this metric assuming independent link failures in known to be NP-hard [10]. We then present some numerical results to demonstrate the significance of geometry on the survivability of the network and also discussed network design problems in the context of these randomly located failures. Some future research directions include the consideration of multiple randomly located cuts (instead of a single randomly located failure), convex cuts (e.g., oval cuts), and robust network design in the face of geographical failures.

We then presented a geographic max-flow and min-cut problem where failures, modeled as disks, may be placed anywhere in the graph except for certain protected zones. We show these problems can be reduced to discrete ones and present a polynomial time algorithm for the geographic min-cut problem. We developed an ILP formulation, an exact algorithm, and a heuristic algorithm for the geographic max-flow problem. Using these algorithms, we obtain numerical results for a specific backbone network, thereby demonstrating the applicability of our algorithms to a real-world network. This approach provides a way to look at network survivability in the face of multiple disasters or attacks that takes into account the geographical correlation between links. Some future directions include application of this approach to the electric power transmission network, finding a tight bound on the difference between geographic min-cut and max-flow (i.e. the analog to the max-flow min-cut

theorem), and the development of network design tools (e.g. how to build a network under some constraints such that geographic min-cut is maximized).

Motivated by the effects of natural disasters such as geomagnetic storms [59] and cascading failures, we considered a two-stage failure model for power networks. The first stage removes power lines that intersect a randomly located disk and the second stage calculates the cascading failure that occurs due to the removal of the initial links. We used the tools developed for randomly located circular cuts and a cascading failure model to calculate the effect of this type of failure in power networks. Then motivated by the effects of power loss on data networks [30], we considered the survivability of data networks with respect to power networks. We assumed data nodes rely on the operation of the closest power load nodes to function. Through numerical results, we were able to show power network effects have a significant impact on the survivability of real-world data networks. As discussed in section 5.3.5, a natural extension is the problem of how to design backup data paths such that they are maximally 'blackout disjoint.' That is, how to find data paths such that a large power blackout would have a small probability of disconnecting all the paths, thus making the data network more robust to large scale power failures. Another way to extend the work is the consideration of other power system models and metrics (see section 5.2.3).

Another direction to explore with respect to geographically correlated failures is the connection with Shared Risk Link Groups (SRLGs). SRLGs are sets of links in a network which can fail simultaneously. Consider an analog to our work; every set of links which can be simultaneously intersected by a disk belongs to an SRLG under the random disk failure model. However, since a disaster may not necessarily destroy a fiber it seems natural to look at Probabilistic SRLGs (PSRLGs) in which every link in a SRLG fails with some probability [44]. For example, some links in a SRLG may be removed with probability $\frac{1}{2}$ and others may be removed with probability $\frac{1}{3}$. This model can capture the different vulnerabilities of fiber to real-world failures, since, for example, fiber can be above or below ground (it should be noted a model similar to this has recently been considered in [1]). Evaluating performance metrics and considering network design under this more general model can help us find locations

where fiber needs to be better protected (e.g., through shielding).

Bibliography

[1] P. Agarwal, A. Efrat, S. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, "The resilience of wdm networks to probabilistic geographical failures," in *Proc. IEEE INFOCOM'11*, Apr. 2011, to be published.

[2] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows*. Prentice Hall, 1993.

[3] V. Albertson, B. Bozoki, W. Feero, J. Kappenman, E. Larsen, D. Nordell, J. Ponder, F. Prabhakara, K. Thompson, and R. Walling, "Geomagnetic disturbance effects on power systems," *IEEE transactions on power delivery*, vol. 8, no. 3, pp. 1206–1216, 1993.

[4] R. Ambartzumian, "Combinatorial solution of the Buffon Sylvester problem," *Probability Theory and Related Fields*, vol. 29, no. 1, pp. 25–31, 1974.

[5] R. Ambartzumian and A. Baddeley, *Combinatorial Integral Geometry: With Applications to Mathematical Stereology*. John Wiley & Sons, 1982.

[6] G. Andersson, P. Donalek, R. Farmer, N. Hatziargyriou, I. Kamwa, P. Kundur, N. Martins, J. Paserba, P. Pourbeik, J. Sanchez-Gasca *et al.*, "Causes of the 2003 major grid blackouts in north america and europe, and recommended means to improve system dynamic performance," *Power Systems, IEEE Transactions on*, vol. 20, no. 4, pp. 1922–1928, 2005.

[7] A. Antony, L. Cittadini, D. Karrenberg, R. Kisteleki, T. Refice, T. Vest, and R. Wilhelm, "Mediterranean fiber cable cut (January-February 2008) analysis of network dynamics," Dept. of Computer Science and Automation, University of Roma Tre, Tech. Rep. RT-DIA-124-2008, 2008.

[8] K. Atkins, J. Chen, V. S. Anil Kumar, and A. Marathe, "The structure of electrical networks: a graph theory-based analysis," *Int. J. Critical Infrastructures*, vol. 5, no. 3, pp. 265–284, 2009.

[9] R. Baldick, B. Chowdhury, I. Dobson, Z. Dong, B. Gou, D. Hawkins, Z. Huang, M. Joung, J. Kim, D. Kirschen *et al.*, "Vulnerability assessment for cascading failures in electric power systems," in *Power Systems Conference and Exposition, 2009. PSCE'09. IEEE/PES*. IEEE, 2009, pp. 1–9.

[10] M. Ball, "Computational complexity of network reliability analysis: An overview," *Reliability, IEEE Transactions on*, vol. 35, no. 3, pp. 230–239, 1986.

[11] S. Bar, M. Gonen, and A. Wool, "A geographic directed preferential Internet topology model," in *Proc. IEEE MASCOTS'05*, Sept. 2005.

[12] A. L. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, October 1999.

[13] J. Bentley and T. Ottmann, "Algorithms for reporting and counting geometric intersections," *IEEE Trans. Comput.*, vol. 28, no. 9, pp. 643–647, 1979.

[14] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman, "Power grid vulnerability to geographically correlated failures - analysis and control implications," no. EE Technical Report 2011-05-06, 2011.

[15] R. Bhandari, *Survivable networks: algorithms for diverse routing*. Kluwer, 1999.

[16] D. Bienstock, "Optimal control of cascading power grid failures," in *PES General Meeting and submitted to IEEE CDC-ECC11*, 2011.

[17] D. Bienstock and S. Mattia, "Using mixed-integer programming to solve power grid blackout problems," *Discrete Optimization*, vol. 4, no. 1, pp. 115–141, 2007.

[18] D. Bienstock, "Some generalized max-flow min-cut problems in the plane," *Math. Oper. Res.*, vol. 16, no. 2, pp. 310–333, 1991.

[19] J. Bondy and U. Murty, "R. Graph theory," 2008.

[20] J. Borland, "Analyzing the Internet collapse," *MIT Technology Review*, Feb. 2008. Editorial. [Online]. Available: http://www.technologyreview.com/Infotech/20152/?a=f

[21] G. Borradaile and P. Klein, "An $O(n \log n)$ algorithm for maximum st-flow in a directed planar graph," *Journal of the ACM (JACM)*, vol. 56, no. 2, p. 9, 2009.

[22] D. Boteler, R. Pirjola, and H. Nevanlinna, "The effects of geomagnetic disturbances on electrical systems at the earth's surface," *Advances in Space Research*, vol. 22, no. 1, pp. 17–27, 1998.

[23] S. Buldyrev, R. Parshani, G. Paul, H. Stanley, and S. Havlin, "Catastrophic cascade of failures in interdependent networks," *Nature*, vol. 464, no. 7291, pp. 1025–1028, 2010.

[24] C. Burch, R. Carr, S. Krumke, M. Marathe, C. Phillips, and E. Sundberg, "A decomposition-based pseudoapproximation algorithm for network flow inhibition," in *Network Interdiction and Stochastic Integer Programming*, D. L. Woodruff, Ed. Kluwer, 2003, ch. 3, pp. 51–68.

[25] T. M. Chan, "A simple trapezoid sweep algorithm for reporting red/blue segment intersections," in *Proc. 6th Canad. Conf. Comp. Geom. (CCCG)*, 1994.

[26] R. L. Church, M. P. Scaparra, and R. S. Middleton, "Identifying critical infrastructure: the median and covering facility interdiction problems," *Ann. Assoc. Amer. Geographers*, vol. 94, no. 3, pp. 491–502, 2004.

[27] R. Cohen, K. Erez, D. Ben-Avraham, and S. Havlin, "Resilience of the Internet to random breakdowns," *Phys. Rev. Lett.*, vol. 85, pp. 4626–4628, Nov. 2000.

[28] R. Cohen, K. Erez, D. Ben-Avraham, and S. Havlin, "Breakdown of the Internet under intentional attack," *Phys. Rev. Lett.*, vol. 86, no. 16, pp. 3682–3685, Apr 2001.

[29] Columbus Networks, Network Map. [Online]. Available: http://www.columbus-networks.com/LandingPoints.html

[30] J. Cowie, A. Ogielski, B. Premore, E. Smith, and T. Underwood, "Impact of the 2003 blackouts on internet communications," *Preliminary Report, Renesys Corporation (updated March 1, 2004)*, 2003.

[31] H. Edelsbrunner, L. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir, "Arrangements of curves in the planetopology, combinatorics, and algorithms," *Automata, Languages and Programming*, pp. 214–229, 1988.

[32] A. Efrat and M. Sharir, "A near-linear algorithm for the planar segment-center problem," *Discrete and Computational Geometry*, vol. 16, no. 3, pp. 239–257, 1996.

[33] P. Francis, S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang, "IDMaps: A global internet host distance estimation service," *IEEE/ACM Trans. Netw.*, vol. 9, no. 5, pp. 525–540, October 2001.

[34] L. K. Gallos, R. Cohen, P. Argyrakis, A. Bunde, and S. Havlin, "Stability and topology of scale-free networks under attack and defense strategies," *Phys. Rev. Lett.*, vol. 94, no. 18, 2005.

[35] O. Gerstel and R. Ramaswami, "Optical layer survivability: a services perspective," *IEEE Commun.*, vol. 38, no. 3, pp. 104–113, Mar. 2000.

[36] A. F. Hansen, A. Kvalbein, T. Cicic, and S. Gjessing, "Resilient routing layers for network disaster planning," in *Proc. Networking - ICN 2005, LNCS, Vol. 3421*. Springer-Verlag, Apr. 2005.

[37] E. Harding, "The number of partitions of a set of n points in k dimensions induced by hyperplanes," *Proceedings of the Edinburgh Mathematical Society, Series II*, vol. 15, no. 285-289, p. 1967, 1966.

[38] B. Heezen and W. Ewing, "Turbidity currents and submarine slumps, and the 1929 grand banks [newfoundland] earthquake," *American Journal of Science*, vol. 250, no. 12, p. 849, 1952.

[39] IETF Internet Working Group, "Inference of shared risk link groups," Nov. 2001, Internet Draft. [Online]. Available: http://tools.ietf.org/html/draft-many-inference-srlg-02

[40] J. S. Foster Jr. et al., "Report of the commission to assess the threat to the United States from electromagnetic pulse (EMP) attack, Vol. I: Executive report," Apr. 2004.

[41] A. Lakhina, J. Byers, M. Crovella, and I. Matta, "On the geographic location of Internet resources," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 6, pp. 934–948, 2003.

[42] S. LaPerrière, "Taiwan earthquake fiber cuts: a service provider view." [Online]. Available: http://www.nanog.org/mtg-0702/presentations/laperriere.pdf

[43] E. Lawler, A. Rinnooy-Kan, J. Lenstra, and D. Shmoys, *The traveling salesman problem: a guided tour of combinatorial optimization.* John Wiley & Sons Inc, 1985.

[44] H. Lee, E. Modiano, and K. Lee, "Diverse routing in networks with probabilistic failures," *Networking, IEEE/ACM Transactions on*, vol. 18, no. 6, pp. 1895–1907, 2010.

[45] K. Lee, E. Modiano, and H. Lee, "Cross-layer survivability in wdm-based networks," *IEEE/ACM Transactions on Networking*, no. 99, pp. 1–1, 2009.

[46] Level 3 Communications. (2008, Sept.) Network map. [Online]. Available: http://www.level3.com/Resource-Library/Maps/Level-3-Network-Map.aspx

[47] R. P. Loukas Lazos and J. A. Ritcey, "Detection of mobile targets on the plane and in space using heterogeneous sensor networks," *Wireless Networks*.

[48] D. Magoni, "Tearing down the Internet," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 6, pp. 949–960, Aug. 2003.

[49] J. Manchester, D. Saha, and S. K. Tripathi (eds.), "Protection, restoration, and disoster recovery," *IEEE Network, Special issue*, vol. 18, no. 2, Mar.–Apr. 2004.

[50] A. Mathai, *An introduction to geometrical probability.* CRC Press, 1999.

[51] Merit Networks, Network Map. [Online]. Available: http://www.merit.edu/networkresearch/projecthistory/nsfnet/nsfnet\%5Fmaps.php

[52] J. Mitchell and V. Polishchuk, "Thick non-crossing paths and minimum-cost flows in polygonal domains," in *Proceedings 23rd ACM Symposium on Computational Geometry.* Citeseer, 2007, pp. 56–65.

[53] E. Modiano and A. Narula-Tam, "Survivable lightpath routing: a new approach to the design of WDM-based networks," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 4, pp. 800–809, May 2002.

[54] A. Narula-Tam, E. Modiano, and A. Brzezinski, "Physical topology design for survivable routing of logical rings in WDM-based networks," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 8, pp. 1525–1538, Oct. 2004.

[55] J. O'Rourke, *Computational Geometry in C.* Cambridge University Press, 1994.

[56] T. Overbye, X. Cheng, and Y. Sun, "A comparison of the ac and dc power flow models for lmp calculations," in *System Sciences, 2004. Proceedings of the 37th Annual Hawaii International Conference on.* IEEE, 2004, pp. 9–pp.

[57] C. A. Phillips, "The network inhibition problem," in *Proc. ACM STOC'93*, 1993.

[58] A. Pinar, Y. Fogel, and B. Lesieutre, "The inhibiting bisection problem," in *Proc. ACM SPAA'07*, 2007.

[59] R. Pirjola, "Geomagnetically induced currents during magnetic storms," *Plasma Science, IEEE Transactions on*, vol. 28, no. 6, pp. 1867–1873, 2000.

[60] W. Radasky, "High-altitude electromagnetic pulse (HEMP): A threat to our way of life," *IEEE-USA Today's Engineer*, Sep. 2007. [Online]. Available: http://www.todaysengineer.org/2007/Sep/HEMP.asp

[61] S. Rai and D. P. Agrawal, *Distributed Computing Network Reliability.* Los Alamitos, CA, USA: IEEE Computer Society Press, 1990.

[62] J. Reif, "Minimum st cut of a planar undirected network in $O(n \log^2(n))$ time," *Automata, Languages and Programming*, pp. 56–67, 1981.

[63] V. Rosato, L. Issacharoff, F. Tiriticco, S. Meloni, S. Porcellinis, and R. Setola, "Modelling interdependent infrastructures using interacting dynamical models," *International Journal of Critical Infrastructures*, vol. 4, no. 1, pp. 63–79, 2008.

[64] V. Rosato and S. C.-K. Chau, personal communication, 2011.

[65] L. Santalo, *Integral Geometry and Geometric Probability.* Cambridge University Press, 2004.

[66] A. Sen, S. Murthy, and S. Banerjee, "Region-based connectivity: a new paradigm for design of fault-tolerant networks," in *Proceedings of the 15th international conference on High Performance Switching and Routing.* IEEE Press, 2009, pp. 94–100.

[67] A. Sen, B. Shen, L. Zhou, and B. Hao, "Fault-tolerance in sensor networks: A new evaluation metric," in *Proc. IEEE INFOCOM'06*, 2006.

[68] J. Sylvester, "On a funicular solution of Buffon's problem of the needle in its most general form," *Acta Mathematica*, vol. 14, no. 1, pp. 185–205, 1890.

[69] C. P. Warren, L. M. Sander, and I. M. Sokolov, "Geography in a scale-free network model," *Phys. Rev. E*, vol. 66, no. 5, Nov. 2002.

[70] C. Wilson, "High altitude electromagnetic pulse (HEMP) and high power microwave (HPM) devices: Threat assessments," CRS Report for Congress, Aug. 2004. [Online]. Available: http://www.fas.org/man/crs/RL32544.pdf

[71] W. Wu, B. Moran, J. Manton, and M. Zukerman, "Topology design of undersea cables considering survivability under major disasters," in *Proc. WAINA '09*, May 2009.

[72] S. Yook, H. Jeong, and A. Barabasi, "Modeling the Internet's large-scale topology," *Proc. Natl. Acad. Sci.*, vol. 99, no. 21, pp. 13 382–13 386, 2002.

[73] H. Yukio and J. Matsukubo, "A review of recent studies of geographical scale-free networks," *Trans. Inf. Process. Society of Japan*, vol. 47, no. 3, pp. 776–785, 2006.

[74] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE Network*, vol. 14, no. 6, pp. 16–23, Nov.-Dec. 2000.