

Assessing the Impact of Geographically Correlated Network Failures

Sebastian Neumayer
Electrical Engineering
MIT
Cambridge, MA 02139
bastian@mit.edu

Gil Zussman
Electrical Engineering
Columbia University
New York, NY 10027
gil@ee.columbia.edu

Reuven Cohen
Mathematics
Bar-Ilan University
Ramat Gan, Israel 52900
reuven@math.biu.ac.il

Eytan Modiano
MIT
Cambridge, MA 02139
modiano@mit.edu

Abstract—Communication networks are vulnerable to natural disasters, such as earthquakes or floods, as well as to human attacks, such as an electromagnetic pulse (EMP) attack. Such real-world events have geographical locations, and therefore, the geographical structure of the network graph affects the impact of these events. In this paper we focus on assessing the vulnerability of (geographical) networks to such disasters. In particular, we aim to identify the location of a disaster that would have the maximum effect on network capacity. We consider a geometric graph model in which nodes and links are geographically located on a plane. Specifically, we model the physical network as a bipartite graph (in the topological and geographical sense) and consider the set of all vertical line segment cuts. For that model, we develop a polynomial time algorithm for finding a worst possible cut. Our approach has the potential to be extended to general graphs and provides a promising new direction for network design to avert geographical disasters or attacks.

Index Terms—Network survivability, geographic networks, Internet, cut capacity, network design, Electromagnetic Pulse (EMP)

I. INTRODUCTION

The U.S. military and civilian communications infrastructure is primarily based on fiber-optic networks. This infrastructure exists physically and has physical vulnerabilities. Fiber links and backbone nodes can be destroyed by anything from electromagnetic pulse (EMP) attacks [8], [12] to dragging anchors [4], [13]. Such real-world disasters have geographical locations, and therefore, the geography of the infrastructure affects the impact of these events. For example, since an EMP attack can affect electronic components in a large geographic area, such an attack over a city which is a telecommunications hub would have a disastrous impact on the U.S. telecommunications capabilities. Our approach is to gain insight into robust network design by developing the necessary theory to find the most geographically vulnerable areas of a network. This can provide important input to the construction of network design tools and can support the efforts to mitigate the effects of regional disasters.

There are several works on the topology of the Internet as a random graph [2] and on the effect of link failures in these graphs [7], [10]. 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. Some works such as [9], [14] model the Internet using geographical notions. Yet, these works do not consider the effect of failures that are geographically correlated. Finally, variants of the network inhibition problem in which a set of links has to be removed from a graph such that the effect on the graph will be maximized have been studied in [11] and reviewed in [6]. Yet, to the best of our knowledge, this problem was not studied under the assumption of geographically correlated failures. 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 (i.e., the effect on the fibers rather than on the logical links). It should be noted that fibers are subject to regional attacks, such as earthquakes, floods, and even an EMP attack; as these may affect the electronic amplifiers that are needed to operate the fiber plant.

Our long-term goal is to understand the effect of a regional failure on the bandwidth and connectivity of the Internet and to expose the design tradeoffs related to network survivability under an attack/disaster with regional implications. In this paper, we are interested in the location of geographical disasters that maximize the capacity of disconnected links. That is, we want to identify the worst-case location for a disaster or an attack.

We focus on a graph model which can serve as an abstraction of the U.S. fiber plant, the transpacific fiber links, or the transatlantic fiber links. In this model, nodes, links, and cuts are geographically located on a plane. We consider a bipartite graph (in the topological and geographical sense) which 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).¹ Since the continental U.S. has a width that is greater than its height, vertical disasters seemingly cut mostly east-west links. Therefore, and since vertical line segment cuts are somewhat simpler to analyze, we focus in this paper on such cuts.

Under this model, we study the effect of a north-south

¹Similarly, nodes can represent major cities on both sides of an ocean.




Fig. 1. A bipartite network and an example of a cut

regional disaster on the east-west capacity of the network. We assume that a vertical regional disaster (e.g., a tornado) affects the electronic components of the network within a certain region. Hence, the fibers that pass through that region are effectively cut due to such a disaster. The more communication links a vertical cut intersects, the greater the effect on the communication capacity between the east and west coasts. We start by considering a simple case of the bipartite model and providing motivating examples. We then present a polynomial time algorithm for finding the location of a cut which maximizes the expected capacity of the removed links.

The rest of the paper is organized as follows. In Section II, we introduce the network model and formulate the problem. In Section III, we consider a simple case of the bipartite model and provide numerical examples. In Section IV, we develop a polynomial-time algorithm for finding the worst-case cut. In Section V we present numerical results for worst-case cuts in the Pacific Ocean and in Section VI we review our results and discuss future research directions.

II. MODEL AND PROBLEM FORMULATION

In this section we present a geographical bipartite network model (in the topological and geographic sense) and assume that the cuts are vertical line segments.

We define the *geometric* bipartite graph as follows. 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; that is $r_i \neq r_j \forall i, j$ and $l_i \neq l_j \forall i, j$. Also, the number of nodes on the left side is denoted by L , the number of nodes on the right side is R , and the total number of nodes is N . We denote a link from l_i to r_j as (i, j) . We define p_{ij} as the probability that link (i, j) exists, and c_{ij} as the capacity of link (i, j) where $c_{ij} \in [0, \infty)$. In order 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 a disaster results in a vertical line segment cut of height h whose lowest point is at point $[x, y]$. We define this cut as $\text{cut}_h(x, y)$. Such a cut removes all links which intersect it. To be clear, in this paper 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. 1 demonstrates a particular construction of the model and an example of a cut.

There are many ways to define a loss of communication capacity in a network. For simplicity, we define a worst-case cut as follows.

Definition 1 (Worst-Case Cut): A worst-case cut, denoted by $\text{cut}_h(x^*, y^*)$, is a cut which maximizes the total expected capacity of the intersected links.

It should be noted that because this is a bipartite graph setting, a worst-case cut is also one which minimizes maximum flow between the two sides.

In order to formulate an optimization problem which finds a worst-case cut, we define the following $(0, 1)$ variables:

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

Given a cut height, link probabilities, and capacities, the solution to the optimization problem below is an endpoint of the worst-case cut.

$$\begin{aligned} & \max \sum_{(i,j)} p_{ij} c_{ij} z_{ij}(x, y) \\ & \text{such that} \\ & 0 \leq x \leq 1 \\ & -h \leq y \leq h_G \\ & z_{ij}(x, y) = \mathbf{1}_{y \leq (r_j - l_i)x + l_i} \mathbf{1}_{y + h \geq (r_j - l_i)x + l_i} \end{aligned} \quad (1)$$

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.

$$\begin{aligned} & \max \sum_{(i,j)} p_{ij} c_{ij} z_{ij} \\ & \text{such that} \\ & (r_j - l_i)x - (y - l_i) \geq u_{ij} - 1 \quad \forall i, j \quad r_j \geq l_i \\ & (y + h - l_i) - (r_j - l_i)x \geq d_{ij} - 1 \quad \forall i, j \quad r_j \geq l_i \\ & (l_i - r_j)(1 - x) - (y - r_j) \geq u_{ij} - 1 \quad \forall i, j \quad l_i > r_j \\ & (y + h - r_j) - (l_i - r_j)(1 - x) \geq d_{ij} - 1 \quad \forall i, j \quad l_i > r_j \\ & u_{ij} + d_{ij} \geq 2z_{ij} \quad \forall i, j \\ & x \geq 0 \\ & y \geq 0 \\ & u_{ij}, d_{ij}, z_{ij} \in \{0, 1\} \end{aligned}$$

In general, solving integer programs can be computationally intensive. However, the geographical (geometric) nature of the problem lends itself to relatively low complexity algorithms. Hence, in Section IV we show that the worst-case cut can be found by searching over a polynomial number of potential cuts.




Fig. 2. An example of the fully connected bipartite graph with $L = R = 4$.




Fig. 3. Number of links intersected by a worst-case cut ($\text{cut}_h(x^*, y^*)$) as a function of the cut height (h) in a bipartite graph with 15 nodes on each side.

III. A MOTIVATING EXAMPLE

In this section we consider a simple case of a fully connected bipartite graph. We let $L = R$ (number of nodes is the same on both sides), $p_{ij} = 1$, and $c_{ij} = 1 \forall i, j$. We also place nodes evenly on each side such that they are separated by distance $a = h_G/L = h_G/R$. An example is shown in Fig. 2. We first develop a lower bound on the worst-case cut by considering cuts down the center. Then, we provide numerical results for the bipartite graph optimization problem (1).

A. A Lower Bound

In this simple model we can bound the worst-case by looking at worst-case cuts down the center where $x = 0.5$. 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 loss of capacity from a the worst-case cut of height h (assuming $0 < h \leq h_G$) is lower bounded by:

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

B. Intuition from Numerical Results

We employ genetic algorithms to solve specific cases of the optimization problem (1). Since the objective function has many points of discontinuity, we rely on Matlab's genetic algorithm to find probable² optimal solutions (worst-case cuts).

²Note that genetic algorithms do not provide definitive results. Due to their probabilistic nature, they only find lower bounds for optimization problems.




Fig. 4. The maximum number of removed links as a function of the x -location of the cut for $h = 1.6$. Note that there are no 'stand-out' local maxima.




Fig. 5. The maximum number of removed links as a function of the x -location of the cut for $h = 0.1$. Note that there are 'stand-out' local maxima.

We used the genetic algorithm in order to obtain solutions for a graph with 15 nodes on each side ($L = R = 15$) and with $a = 1$ ($h_G = 15$). Fig. 3 illustrates the results obtained regarding the number of removed links for each cut height (h). The result is nearly identical to the lower bound for the center cuts in (2). This implies that a worst-case cut is indeed likely at the center of the graph.

Next, we study the effect of the horizontal cut location on the number of removed links. 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. Figures 4 and 5 illustrate the maximum number of removed links versus the horizontal (x) position of the cut on the graph with $L = R = 15$. With $h = 1.6$ the results were relatively monotonic, with the largest cut appearing at the center while the number of removed links more or less descends from there (Fig. 4). When the cut height is reduced to 0.1, significant local maxima begin to appear (Fig. 5). 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 section, we focus on developing a polynomial-time algorithm for identifying a worst-case cut.

IV. A WORST-CASE CUT

In this section we present an $O(N^6)$ algorithm for finding a worst-case cut. 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 set of all possible cuts is compact and 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 (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, 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 7) takes $O(1)$ time and it has to be evaluated for all (i, j) , finding the capacity of a candidate cut takes $O(N^2)$.³

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

```

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

³Computational geometry results can probably be used to reduce the complexity of Algorithm WCBG. Particularly, [5] (based on [3]), 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 [5] can be used within Algorithm WCBG.




Fig. 6. 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 how (i_ω, j_ω) is not unique.




Fig. 7. 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 .

intersect at least one link. This implies (i_ω, j_ω) and (i_α, j_α) exist for all worst-case cuts.

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 which has an endpoint on a node or link intersection at $[x^*, y_\omega]$ or $[x^*, y_\alpha]$.

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 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. 7. 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 which has an endpoint on a link intersection or node.

Proof:

Let $y_\omega(x) = (r_{j_\omega} - l_{i_\omega})x + l_{i_\omega}$ be the equation of (i_ω, j_ω) on $x \in [0, 1]$. Let $y_\alpha(x) = (r_{j_\alpha} - l_{i_\alpha})x + l_{i_\alpha}$ 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 $y_\alpha(x)$.

$$r_{j_\omega} - l_{i_\omega} \leq r_{j_\alpha} - l_{i_\alpha}$$

- 2) The slope of $y_\omega(x)$ is greater or equal to the slope of




Fig. 8. $\text{cut}_h(x^*, y^*)$ 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 $\text{cut}_h(x', y_\alpha(x'))$, a worst-case cut which has an endpoint on a link intersection.

$$y_\alpha(x) \cdot r_{j_\omega} - l_{i_\omega} \geq r_{j_\alpha} - l_{i_\alpha}$$

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 $y(x') = y(x^*) + (x' - x^*)(r_j - l_i) \forall i, j$. Since we know $y_\omega(x^*) \leq y_\alpha(x^*) + h$ ($\text{cut}_h(x^*, y^*)$ hits both $y_\omega(x)$ and $y_\alpha(x)$) and $(r_{j_\omega} - l_{i_\omega})(x' - x^*) \leq (r_{j_\alpha} - l_{i_\alpha})(x' - x^*)$ (case 1 above and $x' - x^* \geq 0$), we have $y_\omega(x^*) + (r_{j_\omega} - l_{i_\omega})(x' - x^*) \leq y_\alpha(x^*) + (r_{j_\alpha} - l_{i_\alpha})(x' - x^*) + h$. Thus $y_\omega(x') \leq y_\alpha(x') + h$. See Fig. 8.

This means $\text{cut}_h(x', y_\alpha(x'))$ will cut 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 cut by $\text{cut}_h(x^*, y^*)$ are also cut by $\text{cut}_h(x', y_\alpha(x'))$ (they are “trapped” between (i_ω, j_ω) and (i_α, j_α) in $x^* \leq x < x'$).

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. ■

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. 8).

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. ■

V. NUMERICAL RESULTS

In this section we present a numerical result that demonstrates the use of the WCBG algorithm. The result highlights




Fig. 9. Results obtained by the WCBG algorithm for the Pacific Ocean fiber network found in [1]. The thick segments are locations where cuts of length $h = 1$ (1 degree of latitude) can disconnect at least 3 cables.

the vulnerabilities of a specific fiber network that has a bipartite structure. Clearly, the WCBG algorithm can be used to analyze other networks that have mostly bipartite structure. The results were obtained using MATLAB.

We modeled the Pacific Ocean submarine cables as a bipartite graph. We used the Alcatel submarine network map found in [1] to identify the nodes and links in the graph. The nodes on the left side of the graph are large Asian cities which are connected to these cables. The nodes on the right side of the graph are U.S. west coast cities. We assume two cities are connected by a fiber of unit capacity if there is a submarine fiber connecting them. We also assume that the fibers follow straight lines. We model a cut as a vertical line segment of height equal to one degree of latitude (~ 60 miles). Such a cut could represent an underwater earthquake or an intentional cut by a dragging anchor. We used the WCBG algorithm to identify high capacity cuts for this model.

Fig. 9 presents the results obtained by the algorithm. The thick segments in the figure show cuts of length $h = 1$ (1 degree of latitude) which cut at least three cables. These results are intuitive; some cuts simply disconnect nodes (major cities) which have many fibers attached to them (e.g. Tokyo, San Luis Obispo, and Bandon). There are also locations in the mid-Pacific where cuts will disconnect 3 or more links. These represent locations where fibers have been laid over each other and thus make for a more vulnerable area for attack. In the future we plan to extend the algorithm such that it will deal with general graphs and to obtain numerical results for the continental U.S.

VI. CONCLUSIONS

Motivated by applications in the area of network robustness and survivability, in this paper, we focused on the problem of geographical cuts in a graph, whose nodes and links are located in Euclidean space. We provided a preliminary study

of the properties and impact of geographical line segment cuts in bipartite graphs. These graphs can represent the fiber links between the east and west coasts in the U.S. or the fiber links across an ocean. For that model, we developed a polynomial-time algorithm for finding a worst-case cut and used it in order to obtain numerical results.

Our approach provides a fundamentally new way to look at network survivability to disasters or attacks that takes into account the geographical correlation between links. Some future research directions include the design of algorithms for the case in which nodes can be in *arbitrary* locations on the plane and the disaster model includes line segment cuts in any direction. In addition, we plan to consider non-linear cuts (e.g., cuts that have various shapes such as circles or rectangles), more sophisticated metrics for measuring the impact of a cut, as well as the impact of geographical failures on survivable network designs.

ACKNOWLEDGMENT

This work was supported by the NDSEG Fellowship Program, DTRA grant number HDTRA1-07-1-0004, and by a Marie Curie International Fellowship within the 6th European Community Framework Programme.

REFERENCES

- [1] Alcatel-Lucent, Optical fibre submarine network map. [Online]. Available: <http://www1.alcatel-lucent.com/submarine/refs/index.htm>
- [2] A. L. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, October 1999.
- [3] J. Bentley and T. Ottmann, "Algorithms for reporting and counting geometric intersections," *IEEE Trans. Comput.*, vol. 28, no. 9, pp. 643–647, 1979.
- [4] J. Borland, "Analyzing the Internet collapse," *MIT Technology Review*, Feb. 2008. [Online]. Available: <http://www.technologyreview.com/Infotech/20152/?a=f>
- [5] T. M. Chan, "A simple trapezoid sweep algorithm for reporting red/blue segment intersections," in *Proc. 6th Canad. Conf. Comp. Geom. (CCCG)*, 1994.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] D. Magoni, "Tearing down the Internet," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 6, pp. 949–960, Aug. 2003.
- [11] C. A. Phillips, "The network inhibition problem," in *Proc. ACM STOC'93*, 1993.
- [12] 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>
- [13] R. Wilhelm and C. Buckridge, "Mediterranean fibre cable cut – a RIPE NCC analysis," Apr. 2008. [Online]. Available: <http://www.ripe.net/projects/reports/2008cable-cut/index.html>
- [14] 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.