

MIT Open Access Articles

Assessing the Vulnerability of the Fiber Infrastructure to Disasters

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Neumayer, Sebastian, et al. "Assessing the Vulnerability of the Fiber Infrastructure to Disasters." IEEE/ACM Transactions on Networking, vol. 19, no. 6, Dec. 2011, pp. 1610–23.

As Published: <http://dx.doi.org/10.1109/TNET.2011.2128879>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <http://hdl.handle.net/1721.1/115240>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Assessing the Vulnerability of the Fiber Infrastructure to Disasters

Sebastian Neumayer *Student Member, IEEE*, Gil Zussman *Senior Member, IEEE*, Reuven Cohen,
and Eytan Modiano *Senior Member, IEEE*

Abstract—Communication 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. In this paper, we focus on assessing the vulnerability of (geographical) networks to such disasters. In particular, we aim to identify the most vulnerable parts of the network. That is, the locations of disasters that would have the maximum disruptive effect on the network in terms of capacity and connectivity. We consider graph models in which nodes and links are geographically located on a plane. First, we consider a simplistic bipartite graph model and present a polynomial time algorithm for finding a worst-case vertical line segment cut. We then generalize the network model to graphs with nodes at arbitrary locations. We model the disaster event as a line segment or a disk and develop polynomial time algorithms that find a worst-case line segment cut and a worst-case circular cut. Finally, we obtain numerical results for a specific backbone network, thereby demonstrating the applicability of our algorithms to real-world networks. Our novel approach provides a promising new direction for network design to avert geographical disasters or attacks.

Index Terms—Network survivability, geographically correlated failures, fiber-optic, Electromagnetic Pulse (EMP).

I. INTRODUCTION

The global communications infrastructure is primarily based on fiber-optic networks, and as such has physical vulnerabilities. Fiber links and backbone nodes can be destroyed by anything from Electromagnetic Pulse (EMP) attacks [15], [16], [32] to dragging anchors [2], [9]. Such real-world disasters happen in specific geographic locations, and therefore, the geographical layout of the network affects their impact. For example, an EMP is an intense energy field that can instantly overload or disrupt numerous electrical circuits at a large distance, thereby affecting electronic components in a large geographic area [37]. Hence, such an attack over a U.S. city which is a telecommunications hub would have a disastrous impact on the

Sebastian Neumayer and Eytan Modiano are with MIT, Cambridge, MA (email: {bastian, modiano}@mit.edu). Gil Zussman is with the Department of Electrical Engineering, Columbia University, New York, NY (e-mail: gil@ee.columbia.edu). Reuven Cohen is with the Department of Mathematics, Bar-Ilan University, Ramat-Gan, Israel (e-mail: reuven@math.biu.ac.il). This work was supported by DTRA grant numbers HDTRA1-07-1-0004, HDTRA-09-1-005, and HDTRA1-09-1-0057, NSF grant numbers CNS-1017800, CNS-0830961, CNS-1018379, CNS-1054856, and CNS-0626781, the NDSEG Fellowship Program, and CIAN NSF ERC under grant EEC-0812072. Preliminary and partial versions of this paper appeared in Proc. IEEE INFOCOM’09, Apr. 2009 [29] and Proc. IEEE MILCOM’08, Nov. 2008 [28].




Fig. 1. The fiber backbone operated by a major U.S. network provider [23].

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 development 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 [5] and on the effect of link failures in these graphs [13], [24] (for more details see Section II). 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 [22], [39]. Yet, these works do not consider the effect of failures that are geographically correlated. Finally, [30] studied 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. Yet, to the best of our knowledge, the network inhibition 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 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 [37].

Our long-term goal is to understand the effect of a regional failure on the bandwidth, connectivity, and reliability of the Internet, 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 [16], [32]). In this paper, we are interested in the location of geographical disasters that have the maximum effect on the 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.

The global fiber plant has a complicated structure. For example, Fig. 1 presents the fiber backbone operated by a major network provider in the U.S. (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 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 cut 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 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. 1.¹ In particular, we illustrate the locations of cuts that optimize the different performance measures described above.

The main contributions of this paper 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 paper is organized as follows. We briefly discuss related work in Section II. In Section III, we introduce the network models and formulate the geographical network inhibition problems. In Section IV, 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 V, we develop a polynomial-time algorithm for finding the worst-case cuts in the bipartite model. In Sections VI and VII we study the general model with line segment and circular cuts. In Section VIII we present numerical results. We conclude and discuss future research directions in Section IX.

II. RELATED WORK

The issue of network survivability and resilience has been extensively studied in the past (e.g., [7], [19], [25], [41] 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., [26], [27]) focused on a *small number of fiber failures* or on the concept of *Shared Risk Link Group* (SRLG) [21]. On the contrary, in this paper we focus on events that cause a large number of failures in a specific geographical region (e.g., [2], [9], [16], [32]). To the best of our knowledge, *geographically correlated failures* have been considered only in a few papers and under very specific assumptions [3], [20], [38].

The theoretical problem most closely related to the problem we consider is known as the *network inhibition problem* [30]. 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., [10], [12], [31]). However, as mentioned above, the removal of (geographically) neighboring links has been rarely considered [8], [34]. One of the first and perhaps the closest to this concept is the problem studied in [35].

When the logical (i.e., IP) topology is considered, widespread failures have been extensively studied [13], [14], [18], [24]. Most of these works consider the topology of the Internet as a random graph [5] and use percolation theory to study

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




Fig. 2. A bipartite network and an example of a cut.

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., [17]), it has been recently shown that the topology of the Internet is influenced by geographical concepts [4], [22], [39]. These observations motivated the modeling of the Internet as a scale free geographical graph [36], [40]. 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. 1).

III. 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.

A. 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 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. 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.
- *ATTR* - The fraction of pairs of nodes that remain connected (this is similar to the average two-terminal reliability of the network²).
- *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 $\text{cut}_h(x^*, y^*)$, 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:

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

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

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

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

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

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

²The two-terminal reliability between two nodes is the probability they remain connected after random independent link failures [33].

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

cut.

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

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 V). 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.

B. 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 performance measure TEC.*

We define the following (0,1) variable:

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

⁴Notice that the assumption that links are represented by line segments is an approximation of the real deployments (e.g., [23]) in which links may not be linear.

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

$$\begin{aligned}
& \max \sum_{(i,j)} p_{ij} c_{ij} z_{ij}([x, y], [v, w]) \\
& \quad \text{such that} \\
& \quad [x, y] \neq [v, w] \\
& \quad \sqrt{(x - v)^2 + (y - w)^2} \leq h \\
& \quad x_i \leq x \leq x_j \text{ for some } i \text{ and } j \\
& \quad y_i \leq y \leq y_j \text{ for some } i \text{ and } j
\end{aligned} \tag{2}$$

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:

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

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

$$\begin{aligned}
& \max \sum_{(i,j)} p_{ij} c_{ij} z_{ij}([x, y]) \\
& \quad \text{such that} \\
& \quad x_i \leq x \leq x_j \text{ for some } i \text{ and } j \\
& \quad y_i \leq y \leq y_j \text{ for some } i \text{ and } j
\end{aligned} \tag{3}$$

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 conservation 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 VI and VII 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.

IV. 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. 3. We first obtain a lower bound for the BGNI problem by considering cuts down the center. Then, we provide numerical results for the BGNI problem.




Fig. 3. An example of a complete bipartite graph with $N = 8$.




Fig. 4. Number of links intersected (TEC) 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 ($N = 30$).

A. 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:

$$\frac{N}{2} + \sum_{i=1}^{\lfloor \frac{2h}{a} \rfloor} \left(\frac{N}{2} - 1 - \lfloor \frac{i-1}{2} \rfloor \right). \quad (4)$$

B. Intuition from Numerical Results

We now describe numerical solutions obtained for the BGNI problem (1).⁵ We obtained solutions for a network with 15 nodes on each side ($N = 30$) and with $a = 1$ ($h_G = 14$). Fig. 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 (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 5 and 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 $h = 1.6$ the results were relatively monotonic, with the worst-case cut appearing

⁵These solutions were initially obtained using MATLAB's genetic algorithms and later on verified using the algorithm described in Section V.




Fig. 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.




Fig. 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 $x \approx .7$.

at the center while the number of removed links more or less descends from there (Fig. 5). When the cut height is reduced to 0.1, significant local maxima begin to appear (Fig. 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.

V. 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 (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)




Fig. 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.

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 $O(N^2)$.⁶

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

```

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

```

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. 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.

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.

⁶Computational geometry results can probably be used to reduce the complexity of Algorithm WCBG. Particularly, [11] (based on [6]), 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 [11] can be used within Algorithm WCBG.




Fig. 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 .




Fig. 9. $\text{cut}_h(x^*, y^*)$ 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 that has an endpoint on 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. 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.

Proof: see Appendix A.

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. 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.

VI. 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 V. In this section, a worst-case cut refers to a worst-case *line segment* cut.

A. TEC Performance Measure

Before proceeding, note that the objective function in (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) (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.

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

```

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:      $L = L \cup \{\text{cut that has an endpoint at } [x_k, y_k] \text{ and contains } [x_i, y_i]\}$ 
7:   for every  $(i, j)$  do
8:      $L = L \cup \{\text{cuts that have an endpoint at } [x_k, y_k] \text{ and another endpoint on } (i, j)\}$ 
9:   for every  $(i, j)$  and node  $k$  do
10:    for every node  $l$  such that  $k \neq l$  do
11:       $L = L \cup \{\text{cuts that have an endpoint on } (i, j) \text{ and contain } [x_k, y_k] \text{ and } [x_l, y_l]\}$ 
12:    for every  $(m, n)$  do
13:       $L = L \cup \{\text{cuts that have an endpoint on } (i, j), \text{ another endpoint on } (m, n), \text{ and contain } [x_k, y_k]\}$ 
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  $\text{cut}_h^*$ 

```

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

```

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

Theorem 2: Algorithm WLGGM 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.




Fig. 10. cut_h' contains a node as well as intersects all links which cut_h^* does.

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.

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. 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 $\text{cut}_h'([x^* + a, y^* + b], [v^* + a, w^* + b])$ and $\text{cut}_h''([x^* - c, y^* - d], [v^* - c, w^* - d])$ to be translated versions of cut_h^* such that (i) $\text{sign}(a) = \text{sign}(c)$ and $\text{sign}(b) = \text{sign}(d)$, (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 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. 10) or (ii) contains a point where (i_α, j_α) or (i_ω, j_ω) intersects a




Fig. 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.

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. ■

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. 11). ■

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.

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. 12). Thus, cut'_h is a worst-case cut that contains A and has an endpoint on a link (this endpoint is $[x_\alpha, y_\alpha]$). ■

Lemma 6: If there exists a worst-case cut that contains A




Fig. 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.




Fig. 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 has an endpoint on a link, then there exists a worst-case cut that contains A , has an endpoint on a link, 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. 13. Fig. 14, which demonstrates the third possibility, shows cut'_h that contains A and has both endpoints on links. ■

Using the lemmas above we now prove Theorem 2.

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




Fig. 14. Translate an endpoint of cut_h^* along L until it can no longer intersect the bottom link. This new cut is cut'_h .

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.

B. ATTR, MFST, and AMF Performance Measures

As mentioned in Section III-B, the formulation of the GNIL Problem, presented in (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 VIII, 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., $O(N^3)$ [1]).
- *AMF* - Run a max-flow algorithm for every node pair.

VII. 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 (3) takes on a finite number of bounded values. This leads to the following observation.

Observation 3: There always exists an optimal solution to (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:    $L = L \cup \{\text{cuts that intersect } (i, j) \text{ at exactly one point and are centered on the line which contains } (i, j)\}$ 
6:   for  $(k, l)$  such that  $(i, j) \neq (k, l)$  do
7:     if  $(i, j)$  is parallel to  $(k, l)$  then
8:        $L = L \cup \{\text{cuts that contain node } i \text{ or } j \text{ on its boundary and intersect } (k, l) \text{ at exactly one point}\}$ 
9:     else
10:       $L = L \cup \{\text{cuts that intersect } (i, j) \text{ and } (l, k) \text{ at exactly one point each such that these points are distinct}\}$ 
11:   for every  $\text{cut}_r(x_k, y_k) \in L$  do
12:     call evaluateCapacityofCut( $x_k, y_k$ )
13:   return  $\text{cut}_r^*$ 
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:      $\text{cut}_r^* \leftarrow \text{cut}_r(x_k, y_k)$ 
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, 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. ■

Lemma 8: If there exists a worst-case cut, denoted by cut_r^* , which intersects exactly one link, then there exists a worst-




Fig. 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').

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. 15. ■

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 VI. 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 nonnegative 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




Fig. 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) .

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 the point on $\mathbf{p}(t)$ closest to C that intersects the line containing (k, l) (see Fig. 16). Additionally, we require that $\mathbf{p}(t)$ is exactly r units away from C for $0 \leq t \leq t_C$. 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 $0 \leq t \leq t_C$. 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. 17). Consider $\text{cut}_r[p_x(t), p_y(t)]$ where t is




Fig. 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 $\mathbf{p}(t)$ to cut_r' where (i, j) and (k, l) each intersect cut_r' at exactly one point.

the minimum nonnegative value such that two 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. ■

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 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

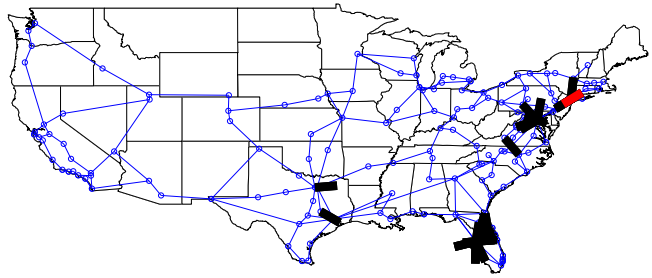


Fig. 18. Line segments cuts optimizing TEC for $h = 2$ - the red cuts maximize TEC and the black segments are nearly worst-case cuts.

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 III-B, the formulation of the GNIC Problem, presented in (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 VI-B).

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 transformation to the network node locations and then running WCGM.

VIII. NUMERICAL RESULTS

In this section we present numerical results that demonstrate the use of the algorithms presented in sections VI and VII. 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 WCGM, presented in Section VI, to compute worst-case cuts under the TEC , $ATTR$, $MFST$, and AMF performance measures for a fiber plant of a major network provider [23]. In all cases, we found that the results are intuitive. We also used Algorithm WCGM, presented in Section VII, 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




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




Fig. 21. Line segments cuts optimizing the AMF for $h = 2$ - the red cuts minimize AMF and the black segments are nearly worst-case cuts.




Fig. 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.




Fig. 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.

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. 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, such as areas in Florida, New York, and around Dallas. Fig. 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. 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. 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 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. 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. 23). Due to the complexity of the network along the east coast, the results were less straightforward than in the Los Angeles-NYC case.

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

IX. CONCLUSIONS

Motivated by applications in the area of network robustness and survivability, in this paper, we focused on the problem of geographical network inhibition. Namely, we studied the properties and impact of geographical disasters that can be




Fig. 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.




Fig. 24. The maximum value of *TEC* as a function of the cut radius for a circular cut in the fiber plant illustrated in Fig. 1.

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 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. Moreover, we plan to study the impact of geographical failures on the design of survivable components, networks, and systems.

APPENDIX A PROOF OF LEMMA 2

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 $y_\omega(x^*) \leq y_\alpha(x^*) + h$ ($\text{cut}_h(x^*, y^*)$ intersects both $y_\omega(x)$ and $y_\alpha(x)$) and $(r_\omega - l_\omega)(x' - x^*) \leq (r_\alpha - l_\alpha)(x' - x^*)$ (case 1 above and $x' - x^* \geq 0$), 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. 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. ■

REFERENCES

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows*. Prentice Hall, 1993.
- [2] 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.
- [3] 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.
- [4] S. Bar, M. Gonen, and A. Wool, “A geographic directed preferential Internet topology model,” *Computer Networks*, vol. 51, no. 14, pp. 4174–4188, 2007.
- [5] A. L. Barabasi and R. Albert, “Emergence of scaling in random networks,” *Science*, vol. 286, no. 5439, pp. 509–512, October 1999.
- [6] J. Bentley and T. Ottmann, “Algorithms for reporting and counting geometric intersections,” *IEEE Trans. Comput.*, vol. 28, no. 9, pp. 643–647, 1979.
- [7] R. Bhandari, *Survivable networks: algorithms for diverse routing*. Kluwer, 1999.
- [8] D. Bienstock, “Some generalized max-flow min-cut problems in the plane,” *Math. Oper. Res.*, vol. 16, no. 2, pp. 310–333, 1991.
- [9] J. Borland, “Analyzing the Internet collapse,” *MIT Technology Review*, Feb. 2008. [Online]. Available: <http://www.technologyreview.com/Infotech/20152/?a=f>
- [10] C. Burch, R. Carr, S. Krumke, M. Marathe, C. Phillips, and E. Sundberg, “A decomposition-based pseudoapproximation algorithm for network flow interdiction,” in *Network Interdiction and Stochastic Integer Programming*, D. L. Woodruff, Ed. Kluwer, 2003, ch. 3, pp. 51–68.
- [11] T. M. Chan, “A simple trapezoid sweep algorithm for reporting red/blue segment intersections,” in *Proc. 6th Canad. Conf. Comp. Geom. (CCCG)*, 1994, pp. 263–268.
- [12] 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.

- [13] 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.
- [14] 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.
- [15] W. R. Forstchen, *One Second After*. Tom Doherty Associates, LLC, 2009.
- [16] J. S. Foster, E. Gjelle, W. R. Graham, R. J. Hermann, H. M. Kluepfel, R. L. Lawson, G. K. Soper, L. L. Wood, and J. B. Woodard, "Report of the commission to assess the threat to the United States from electromagnetic pulse (EMP) attack, critical national infrastructures," Apr. 2008.
- [17] 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.
- [18] 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, p. 188701, 2005.
- [19] O. Gerstel and R. Ramaswami, "Optical layer survivability: a services perspective," *IEEE Commun.*, vol. 38, no. 3, pp. 104–113, Mar. 2000.
- [20] 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, pp. 1097–1105.
- [21] 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>
- [22] 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.
- [23] Level 3 Communications, Network Map, Sept. 2008. [Online]. Available: <http://www.level3.com/Resource-Library/Maps/Level-3-Network-Map.aspx>
- [24] D. Magoni, "Tearing down the Internet," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 6, pp. 949–960, Aug. 2003.
- [25] J. Manchester, D. Saha, and S. K. Tripathi (eds.), "Protection, restoration, and disaster recovery," *IEEE Network, Special issue*, vol. 18, no. 2, Mar.–Apr. 2004.
- [26] 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.
- [27] 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.
- [28] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the impact of geographically correlated network failures," in *Proc. IEEE MILCOM'08*, Nov. 2008.
- [29] —, "Assessing the vulnerability of the fiber infrastructure to disasters," in *Proc. IEEE INFOCOM'09*, Apr. 2009.
- [30] C. A. Phillips, "The network inhibition problem," in *Proc. ACM STOC'93*, 1993.
- [31] A. Pinar, Y. Fogel, and B. Lesieutre, "The inhibiting bisection problem," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-62142, 2008.
- [32] 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>
- [33] S. Rai and D. P. Agrawal, *Distributed Computing Network Reliability*. Los Alamitos, CA, USA: IEEE Computer Society Press, 1990.
- [34] A. Sen, S. Murthy, and S. Banerjee, "Region-based connectivity: a new paradigm for design of fault-tolerant networks," in *Proc. HPSR'09*, June 2009, pp. 1–7.
- [35] A. Sen, B. Shen, L. Zhou, and B. Hao, "Fault-tolerance in sensor networks: a new evaluation metric," in *Proc. IEEE INFOCOM'06*, Apr. 2006.
- [36] C. P. Warren, L. M. Sander, and I. M. Sokolov, "Geography in a scale-free network model," *Phys. Rev. E*, vol. 66, no. 5, p. 056105, Nov. 2002.
- [37] 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>
- [38] 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, pp. 1154–1159.

- [39] S. Yook, H. Jeong, and A. Barabasi, "Modeling the Internet's large-scale topology," *Proc. Natl. Acad. Sci.*, vol. 99, no. 21, pp. 13382–13386, 2002.
- [40] 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.
- [41] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE Network*, vol. 14, no. 6, pp. 16–23, Nov.-Dec. 2000.



Sebastian Neumayer (M10) received the B.S. degree (with Highest Honors) in electrical engineering from the University of Illinois at Urbana-Champaign in 2004, the M.S. degree in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, in 2007, and is currently pursuing the Ph.D. degree in electrical engineering at MIT.

Mr. Neumayer is a recipient of the NDSEG Fellowship.



Gil Zussman (S02-M05-SM07) received the B.Sc. degree in Industrial Engineering and Management and the B.A. degree in Economics from the Technion Israel Institute of Technology in 1995 (both summa cum laude). He received the M.Sc. degree (summa cum laude) in Operations Research from Tel-Aviv University in 1999 and the Ph.D. degree in Electrical Engineering from the Technion in 2004. Between 1995 and 1998, he served as an engineer in the Israel Defense Forces and between 2004 and 2007 he was a Postdoctoral Associate at MIT. He is currently an

Assistant Professor at the Department of Electrical Engineering in Columbia University. His research interests are in the area of wireless networks. Gil received the Marie Curie Outgoing International Fellowship, the Fulbright Fellowship, the IFIP Networking 2002 Best Student Paper Award, and the OPNETWORK 2002 and ACM SIGMETRICS/IFIP Performance 2006 Best Paper Awards. He received the DTRA Young Investigator Award and the NSF CAREER Award, and was a member of a group that won the 1st place in the 2009 Vodafone Foundation Wireless Innovation Project competition.



Reuven Cohen received his B.Sc. in Physics and Computer Science in 1997 and his Ph.D. in Physics from Bar-Ilan University, Ramat-Gan, Israel in 2004. He was a postdoctoral fellow in the Dept. of Mathematics and Computer Science at the Weizmann Institute, Rehovot, Israel and in the ECE department at Boston University and in the Dept. of Physics at MIT. Since October 2007 he has been at the Department of Mathematics at Bar-Ilan University in Israel, where he is an Assistant Professor. His research interests are random graphs, distributed algorithms and network

stability.



Eytan Modiano (S90-M93-SM00) received his B.S. degree in Electrical Engineering and Computer Science from the University of Connecticut at Storrs in 1986 and his M.S. and PhD degrees, both in Electrical Engineering, from the University of Maryland, College Park, MD, in 1989 and 1992 respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Post Doctoral Fellow during 1992–1993. Between 1993 and 1999 he was with MIT Lincoln Laboratory where he was a project leader for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. Since 1999 he has been on the faculty at MIT; where he is a Professor in the Department of Aeronautics and Astronautics and the Laboratory for Information and Decision Systems (LIDS). His research is on communication networks and protocols with emphasis on satellite, wireless, and optical networks. He is the co-recipient of the Sigmetrics 2006 Best paper award for the paper "Maximizing Throughput in Wireless Networks via Gossiping," and the Wiopt 2005 best student paper award for the paper "Minimum Energy Transmission Scheduling Subject to Deadline Constraints." He is currently an Associate Editor for IEEE/ACM Transactions on Networking, and had served as Associate Editor for IEEE Transactions on Information Theory. He was the Technical Program co-chair for IEEE Wiopt 2006, IEEE Infocom 2007, and ACM MobiHoc 2007.