Assessing the Vulnerability of the Fiber Infrastructure to Disasters

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 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, and model the disaster event as a line segment or a circular cut. We develop algorithms that find a worst-case line segment cut and a worst-case circular cut. Then, 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, geographic networks, fiber-optic, Internet, 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 [14], [24] to dragging anchors [6], [27]. 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 [28]. Hence, such an attack over a U.S. 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 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 [3] and on the effect of link failures in these graphs [8], [17] (for more details see Section II). However, most of these works are motivated by failures of routers due

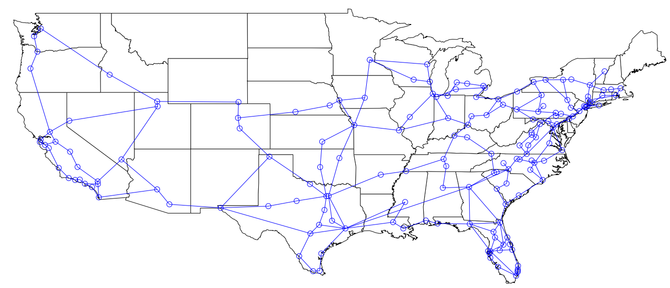

Fig. 1. The fiber backbone operated by a major U.S. network provider [16].

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 [15], [29]. Yet, these works do not consider the effect of failures that are geographically correlated. Finally, [22] 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 [28].

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 [14], [24]). 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. In [21] we 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 in [21] on such cuts.

However, the bipartite model does not consider the impact on nodes located within the continent; nor does it consider the impact of a disaster that is not simply a vertical cut. In this paper, we 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 (circles) 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 average two-terminal reliability [25], (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.

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, this paper and [21] are the first attempts to study this problem.

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

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 briefly discuss the solution for the bipartite case. In Sections V and VI we study the general model with line segment and circular cuts. In Section VII we present numerical results. We conclude and discuss future research directions in Section VIII.

II. Related Work

The issue of network survivability and resilience has been extensively studied in the past (e.g., [4], [12], [18], [31] 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., [19], [20]) focused on a *small number of fiber failures*. On the contrary, in this paper we focus on events that cause a large number of failures in a specific geographical region (e.g., [6], [14], [24], [27]). To the best of our knowledge, [13] is one of the only papers that considered geographically correlated failures. Yet, it focused on a specific routing solution.

The theoretical problem most closely related to the problem we consider is known as the *network inhibition problem* [22]. Under that problem, each edge in the network has a destruction cost, and a fixed budget is given to attack the network. A feasible attack removes a subset of the edges, whose total destruction cost is no greater than the budget. The objective is to find an attack that minimizes the value of a maximum flow in the graph after the attack. Several variants of this problems were studied in the past (see for example [23] and the review in [7]). However, as mentioned above, the removal of (geographically) neighboring links has not been considered (perhaps the closest to this concept is the problem formulated in [5]).

When the logical (i.e., IP) topology is considered, wide-spread failures have been extensively studied [8], [9], [11], [17]. Most of these works consider the topology of the Internet as a random graph [3] and use percolation theory to study the effects of random link and node failures on these graphs. The focus on the logical topology rather than on the physical topology is motivated by failures of routers due to attacks by viruses and worms. Based on various measurements (e.g., [10]), it has been recently shown that the topology of the Internet is influenced by geographical concepts [2], [15], [29]. These observations motivated the modeling of the Internet as a scale free geographical graph [26], [30]. 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 circles on the plane.

A. Bipartite Model with Vertical Line Segment Cuts

We define the *geometric* bipartite graph model 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 $l_i \neq l_j \; \forall \; i,j$. Denote the total number of nodes on the left and right side by N. We also 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 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 $\mathrm{cut}_h(x,y)$. Such a cut removes all links which 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.

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 average two terminal reliability of the network.[2]
- *MFST* - The maximum flow between a given pair of nodes s and t.
- *AMF* - The average value of maximum flow between all pairs of nodes.

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

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

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

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

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

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

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

[2]The two terminal reliability between two nodes is defined here as 1, if there is a path between them and 0, otherwise [25].

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

The above problem can be formulated as a Mixed Integer Linear Program (MILP) [21]. 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 IV). Although in [21] we focus only on 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 by (i,j). 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 $\mathrm{cut}_h([x,y],[v,w])$ (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 $\mathrm{cut}_h([x^*,y^*],[v^*,w^*])$ 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:

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

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

$$\begin{gathered} \max \textstyle\sum_{(i,j)} p_{ij} c_{ij} z_{ij}([x,y],[v,w]) \\ \text{such that} \\ [x,y] \neq [v,w] \\ \sqrt{(x-v)^2 + (y-w)^2} \leq h \\ x_i \leq x \leq x_j \text{ for some } i \text{ and } j \\ y_i \leq y \leq y_j \text{ for some } i \text{ and } j \end{gathered} \tag{1}$$

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 $\mathrm{cut}_r(x,y)$. Such a cut removes all links which intersect it (including the interior of the circle). We call the set of points for which the Euclidean distance is r away from $[x,y]$ the boundary of $\mathrm{cut}_r(x,y)$. For brevity, we sometimes denote the worst-case cut $\mathrm{cut}_r(x^*, y^*)$ as cut_r^*. We now define the following problem and demonstrate its formulation.

Geographical Network Inhibition by Circles (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:

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

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

$$\begin{aligned} &\max \textstyle\sum_{(i,j)} p_{ij} c_{ij} z_{ij}([x, y]) \\ &\qquad\quad \text{such that} \\ &x_i \le x \le x_j \text{ for some } i \text{ and } j \\ &y_i \le y \le y_j \text{ for some } i \text{ and } j \end{aligned} \tag{2}$$

Similar GNIL and GNIC problems can be formulated for performance measures *ATTR*, *MFST*, and *AMF* (for these measures we assume that $p_{ij} \in \{0, 1\}\, \forall i, j$). For example, under *MFST*, flow conversation constraints should be added to the set of constraints, the flow through links for which $z_{ij}([x, y], [v, w]) = 1$ is 0, and the flow between s and t has to be maximized. In sections V and VI 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. Worst-Case Cut – Bipartite Model

In order to facilitate the discussion regarding the general model, we now briefly discuss the $O(N^6)$ algorithm (to be introduced in [21]) for solving the BGNI Problem. In this section, a worst-case cut refers to a worst-case *vertical line segment* cut.

In [21], we show that there exists a worst-case cut which has an endpoint at a link intersection or a node. An algorithm that solves the BGNI Problem simply evaluates the capacity of every cut that has an endpoint on a link intersection or a node. The complexity of this algorithm 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 a node (cuts have two endpoints), and therefore, the total number of candidate cuts is at most $O(N^4)$; (iii) since evaluating if a link intersects a cut takes $O(1)$ time, finding the capacity of a candidate cut takes $O(N^2)$.

We now explain the methodology we use in [21] to develop the algorithm. Given a worst-case cut, we showed that there exists a translation of this cut such that it is also worst-case and has an endpoint on a link intersection or a node. For example, consider $\mathrm{cut}_h(x^*, y^*)$ and the link intersection $[x^*, y_\alpha]$ in Fig. 2. There does not exist a link that intersects $\mathrm{cut}_h(x^*, y^*)$ between $[x^*, y^*]$ and $[x^*, y_\alpha]$. Therefore, we know that $\mathrm{cut}_h(x^*, y_\alpha)$, which is a vertical translation of $\mathrm{cut}_h(x^*, y^*)$, intersects every link which $\mathrm{cut}_h(x^*, y^*)$ intersects. Thus, if $\mathrm{cut}_h(x^*, y^*)$ is a worst-case cut, then $\mathrm{cut}_h(x^*, y_\alpha)$ is a worst-case cut also. Using similar arguments, we presented lemmas which show that any worst-case cut can be translated such that the resulting cut is also worst-case and has an endpoint at a link intersection or a node. The algorithm follows as a direct result.

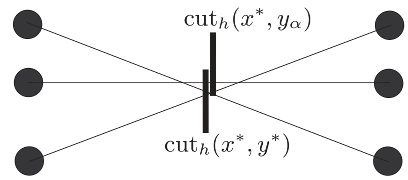

Fig. 2. A geographical bipartite graph with two cuts. $\mathrm{cut}_h(x^*, y_\alpha)$ is a vertically translated version of $\mathrm{cut}_h(x^*, y^*)$. $\mathrm{cut}_h(x^*, y_\alpha)$, which has an endpoint on a link intersection, is guaranteed to intersect every link that $\mathrm{cut}_h(x^*, y^*)$ intersects, since there exist no links at x^* from y^* to y_α.

Although this algorithm finds a worst-case cut, there may be other worst-case cuts that intersect the same total expected capacity of links. The endpoints of these cuts do not necessarily have to be on a link intersection or a node. Yet, there cannot be a cut with a higher value than the one obtained by the algorithm.

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

A. TEC Performance Measure

Before proceeding, note that the set of all possible cuts is compact and the objective function in (1) 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 an algorithm that finds a worst-case line segment cut under the *TEC* measure in the general model. This algorithm considers all cuts that (i) have an endpoint on a link intersection and contain a node not at the intersection, (ii) have an endpoint on a link intersection and another endpoint on a link, (iii) contain two distinct nodes and have an endpoint on a link, and (iv) contain a node and have both endpoints on links.

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

Theorem 1: Algorithm WLGM has a running time of $O(N^8)$ and finds a worst-case line segment cut that is a solution to the GNIL Problem.

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

Lemma 1: There exists a worst-case cut that contains a node or has an endpoint at a link intersection.

Algorithm 1 Worst-Case Line Segment Cut in the General Model (WLGM)

input: h, height of cut
worstCaseCapacityCut $\leftarrow 0$
$L \leftarrow \{\}$
for every link intersection $[x_k, y_k]$ **do**
  for every node i such that $[x_i, y_i] \neq [x_k, y_k]$ **do**
    $L = L \cup \{$cut that has an endpoint on $[x_k, y_k]$ and contains $[x_i, y_i]\}$
  for every (i, j) **do**
    $L = L \cup \{$cuts that have an endpoint on $[x_k, y_k]$ and another endpoint on $(i, j)\}$
for every (i, j) and node k **do**
  for every node l such that $k \neq l$ **do**
    $L = L \cup \{$cuts that have an endpoint on (i, j) and contain $[x_k, y_k]$ and $[x_l, y_l]\}$
  for every (m, n) **do**
    $L = L \cup \{$cuts that have an endpoint on (i, j), another endpoint on (m, n), and contain $[x_k, y_k]\}$
for every $\mathrm{cut}_h([x_k, y_k], [v_k, w_k]) \in L$ **do**
  call evaluateCapacityofCut(x_k, y_k, v_k, w_k**)**
return cut_h^*
Procedure evaluateCapacityofCut(x_k, y_k, v_k, w_k**)**
capacityCut $\leftarrow 0$
for every (i, j) **do**
  if $z_{ij}([x_k, y_k], [v_k, w_k]) = 1$ **then**
    capacityCut $\leftarrow$ capacityCut $+ c_{ij} p_{ij}$
if capacityCut $\geq$ worstCaseCapacityCut **then**
  $\mathrm{cut}_h^* \leftarrow \mathrm{cut}_h([x_k, y_k], [v_k, w_k])$
  worstCaseCapacityCut $\leftarrow$ capacityCut

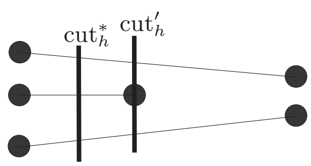

Fig. 3. cut_h' contains a node as well as intersects all links which cut_h^* does.

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. Fig. 2 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 $\mathrm{cut}_h'([x^*+a, y^*+b], [v^*+a, w^*+b])$ and $\mathrm{cut}_h''([x^*-c, y^*-d], [v^*-c, w^*-d])$ to be translated versions of cut_h^* such that (i) $sign(a) = sign(c)$ and $sign(b) = sign(d)$, (ii) there does not exist any nodes in the parallelogram defined by cut_h^* and cut_h' (which we denote "parallelogram B") except those contained in cut_h' and in the parallelogram defined by cut_h^* and cut_h'' (which we denote "parallelogram C") except those contained in cut_h'', and (iii) no link intersects (i_α, j_α) or (i_ω, j_ω) in either parallelogram except on cut_h' or cut_h''. Since a node does not exist within the interior of either parallelogram all links intersected by cut_h^* must also cut one of the other three edges of each parallelogram.

Now choose the maximum a and c such that the edge $([x^*, y^*], [x^*+a, y^*+b])$ of parallelogram B and the edge $([x^*, y^*], [x^*-c, y^*-d])$ of parallelogram C are both parallel to the link (i_α, j_α) and the parallelograms satisfy the constraints in the paragraph above. This implies both cut_h' and cut_h'' contain a node or contain a point where (i_α, j_α) or (i_ω, j_ω) intersects a link. Since (i_α, j_α) is parallel to both edges $([x^*, y^*], [x^*+a, y^*+b])$ and $([x^*, y^*], [x^*-c, y^*-d]))$ and since (i_ω, j_ω) can cut at 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. 3) or (ii) contains a point where (i_α, j_α) or (i_ω, j_ω) intersects a link. In the latter case, we can translate this worst-case cut in a similar fashion to the first case to construct a worst-case cut which satisfies the lemma. ■

We now consider two cases of worst-case cuts. The first case is a worst-case cut which has an *endpoint* at a link intersection. The second case is a worst-case cut which *contains* a node. In both cases, let the node or link intersection that is in the cut be denoted by A. Lemma 2 considers the first case where A is a link intersection.

Lemma 2: If there exists a worst-case cut that has an endpoint on 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 the boundary of 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. ■

The following two lemmas consider the second case where A is a node.

Lemma 3: If there exists a worst-case cut that contains A

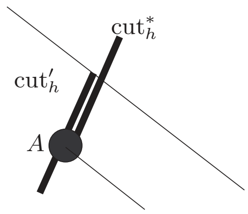

Fig. 4. Translate cut_h^* along its own direction until one of its endpoints intersects a link; we call this new cut cut_h'. Because every link remains intersected during translation, cut_h' intersects all links cut_h^* does.

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

Lemma 4: If there exists a worst-case cut that contains A and has an endpoint on a link, then there exists a worst-case cut which 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 2 implies Lemma 4. 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 its endpoints by $[x_1, y_1]$ and $[x_2, y_2]$. Denote the point at which cut_h^* intersects L by $[x_0, y_0]$. Now 'slide' 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 $0 \le a \le 1$. For $a = 0$ this is just cut_h^*. We slide along L by increasing a until a new cut, called cut_h', either has an endpoint that is h away from A (we cannot slide further) or cut_h' can no longer satisfy $\sum_{(i,j)} p_{i,j} c_{i,j} \mathrm{cut}_h' = \sum_{(i,j)} p_{i,j} c_{i,j} \mathrm{cut}_h^*$. In the first case, the cut has both endpoints on links. In the second case, cut_h' may no longer be able to slide along L and be a worst-case cut, if cut_h' has an endpoint on L which is a link intersection (considered in Lemma 2), cut_h' intersects a node which is not A, or cut_h' has an endpoint on L and the other endpoint on a link (a link 'goes out of reach' from L). The first two possibilities are demonstrated in Fig. 5. They imply cut_h' can have endpoint on a link intersection or can contain another node that is not A. Fig. 6 shows cut_h' which contains A and has both endpoints on links. This can occur when an endpoint

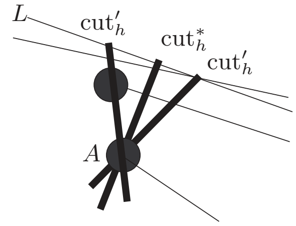

Fig. 5. Slide an endpoint of cut_h^* right along L until it intersects a link intersection. This new cut is the cut_h' on the right. We can also slide and endpoint of cut_h^* left along L until it intersects a node. This new cut is the cut_h' on the left.

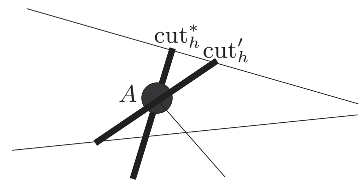

Fig. 6. Slide an endpoint of cut_h^* along the top link until it can no longer intersect the bottom link. This new cut is cut_h'.

of cut_h^* slides along L and the other endpoint intersects a link. ■

Using the lemmas above we now prove Theorem 1.

Proof of Theorem 1: The lemmas presented in this section imply we only need to consider a polynomially sized set of cuts. By Lemma 1 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 2 implies that for every link intersection, $O(N^4)$, there exists a possible worst-case cut for every link and node, $O(N^2)$. In the second case, Lemmas 3 and 4 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 (1) 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 VII, 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 any node pair.

VI. 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 worst-case cut refers to a worst-case *circular* cut.

Before proceeding, note that the set of all possible cuts is compact and 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 which finds a worst-case circular cut under the TEC measure in the general model.

Theorem 2: 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 circles and line segments and then present some lemmas to reduce the set of candidate cuts.

Lemma 5: If a line segment intersects only the boundary of a circle, then the line segment and circle intersect at exactly one point.

Proof: Proof by contradiction. Assume a line segment intersects only the boundary of a circle and this intersection contains more than one point. Since a line segment and a circle are both convex, their intersection must be convex as well. However, we assumed at least two points on the boundary of the circle 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. However, we also assumed only the boundaries intersect and since part of the chord is in the interior of the circle this leads to a contradiction. ■

Lemma 6: If there exists a worst-case cut, denoted by cut_r^*, which intersects exactly one link, then there exists a worst-case cut, denoted by cut_r', which intersects that link at exactly one point such that $[x', y']$ lies on the line which contains the link.

Algorithm 2 Worst-Case Circular Cut in the General Model (WCGM)

input: r, radius of cut
worstCaseCapacityCut $\leftarrow 0$
$L \leftarrow \{\}$
for every (i,j) **do**
  $L = L \cup \{$center $[x_k, y_k]$ of every circle that intersects (i,j) at exactly one point and is centered on the line which contains $(i,j)\}$
  for (k,l) such that $(i,j) \neq (k,l)$ **do**
    if (i,j) is parallel to (k,l) **then**
      $L = L \cup \{$center $[x_k, y_k]$ of every circle having node i or j on its boundary that intersects (k,l) at exactly one point$\}$
    else
      $L = L \cup \{$center $[x_k, y_k]$ of every circle that intersects (i,j) and (l,k) at exactly one point each such that these points are distinct$\}$
for every $(x_k, y_k) \in L$ **do**
  call evaluateCapacityofCut(x_k, y_k)
return cut_r^*

Procedure evaluateCapacityofCut(x_k, y_k)
capacityCut $\leftarrow 0$
for every (i,j) **do**
  if minimum distance from (i,j) to $[x_k, y_k]$ is $\leq r$ **then**
    capacityCut $\leftarrow$ capacityCut $+ c_{ij}p_{ij}$
if capacityCut $\geq$ worstCaseCapacityCut **then**
  $\text{cut}_r^* \leftarrow \text{cut}_r(x_k, y_k)$
  worstCaseCapacityCut $\leftarrow$ capacityCut

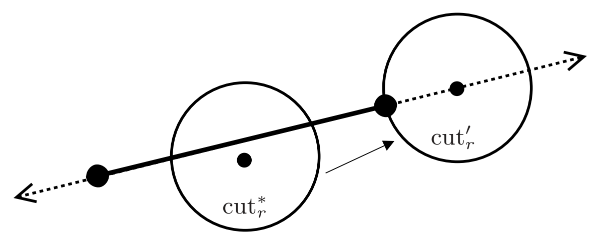

Fig. 7. An example illustrating the observation in Lemma 6. cut_r' is a translated version of cut_r^* such that $[x', y']$ lies on the line which contains the intersected link and cut_r' intersects the link at exactly one point.

Proof: Since cut_r^* is a worst-case cut and only intersects a single link, any cut which intersects the same link is also a worst-case cut. See Fig. 7. ■

Lemma 7: If there exists a worst-case cut 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 which contains one of the two links.

The proof of the lemma above is similar to the proofs of the lemmas in Section V. 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: Let a link which intersects cut_r^* have endpoints 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 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 5 this cut must intersect at least one link at exactly one point (see Fig. 9). Every link which is intersected by cut_r^* must intersect cut_r'' because as a line segment and a circle 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 the circle 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 value such that two links intersect only the boundary of this cut at distinct and non-diametrically opposing points or two links intersect only the boundary of this cut and one of these intersection points is a node. Denote this translated cut by cut_r'. Now, 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.

In second sub-case, two links intersect the circle at a single point, C. This implies C is an endpoint of at least one of these links. Now choose a link with an endpoint at 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 which intersects the line that contains (k, l) (see Fig. 8). 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 circle are continuously translated away from each other, the last non-empty intersection is an intersection of their boundaries. Also, the links that intersect cut_r'' at C remain intersected throughout the translation because $\text{cut}_r[p_x(t), p_y(t)]$ intersects C on $0 \le t \le t_C$. Thus, cut_r' is a worst-case cut and by Lemma 5 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 or ii) $[x', y']$ lies on a line which contains one of the two links ($[x', y'] = \mathbf{p}(t_C)$).

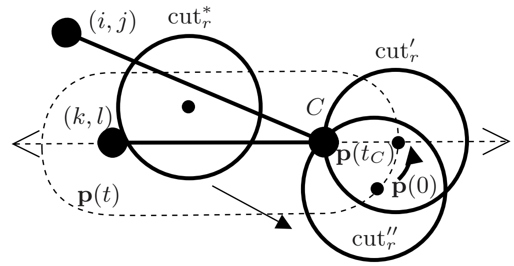

Fig. 8. This figure illustrates a case in the proof of Lemma 7. cut_r^* is first translated in the direction of (i, j) to become cut_r'' which intersects only (k, l) at exactly one point and intersects another link (in this case (i, j)) at exactly the same point. Then cut_r'' is translated along $\mathbf{p}(t)$ towards $\mathbf{p}(t_a)$ to cut_r' such that $[x', y']$ lies on the line which contains (k, l).

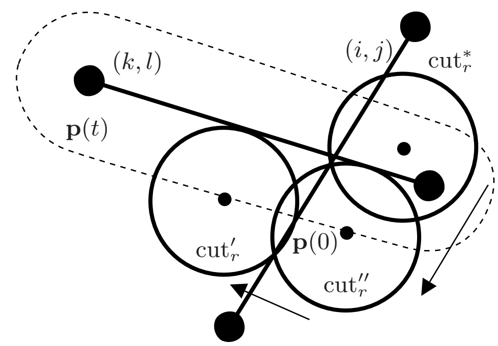

Fig. 9. This figure illustrates a case in the proof of Lemma 7. cut_r^* is first translated in the direction of (i, j) to become cut_r'' which intersects (k, l) at exactly one point. Then cut_r'' is translated along $\mathbf{p}(t)$ to cut_r' where (i, j) and (k, l) intersect only the boundary of the cut at two distinct points.

Now we consider the case where cut_r'' intersects exactly one link at exactly one point (in addition to possibly 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 defined as before (see Fig. 9). 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. By Lemma 5 we know these two links intersect this cut in at most one point each. So this case reduces to the first case for which we know the lemma holds. ■

Lemma 8: There are at most 20 circles of radius r which intersect two non-parallel line segments of positive length at exactly one point each such that these points are distinct.

Proof: If a line segment intersects a circle at exactly one point, then either the endpoint of the line segment intersects the boundary of the circle or the line segment is tangent to the boundary of the circle. This implies every circle which satisfies the lemma falls into at least one of three cases: i) the boundary of the circle intersects two endpoints, ii) the boundary of the circle intersects an endpoint of one segment and the other segment is tangent to the boundary of the circle, or iii) both segments are tangent to the boundary of the circle.

In case one, if two endpoints do not correspond to the same point, by geometry we know there are at most two circles of radius r whose boundary contains these points. If these endpoints correspond to the same point, no circle can intersect this point and satisfy the lemma because this single point belongs to two distinct line segments. In case two, given a point and a line segment we know by geometry there are at most two circles of radius r for which the line segment is tangent to and whose boundary contains the point. In case three, given two non-parallel line segments the lines containing these segments divide the plane into four quadrants. There exists at most one circle tangent to both lines at each of these quadrants. Thus, there are at most four circles tangent to both the line segments. Since for a pair of non-parallel line segments there are four pairs of endpoints, four endpoint-segments pairs, and one segment-segment pair, we know there exists at most 20 circles which may satisfy the lemma. ■

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

Using the above lemmas, we now prove Theorem 2.

Proof of Theorem 2: The lemmas presented in this section imply there exists a worst-case cut which intersects a link at exactly one point such that the center of this cut lies on the line which contains this link or there exists a worst-case cut which intersects two links at exactly one point and at least one of the following: (i) at least two of the points are distinct and are not diametrically opposite or (ii) at least two of the points are distinct and one of them is a node. Algorithm WCGM enumerates all these possible cuts. It considers each link, $O(N^2)$, and finds both cuts which 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 8 we know there are at most 20 of these cuts for every pair of links. If the links are parallel, we need only consider circles 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 (2), 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 V-B).

VII. Numerical Results

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

We used Algorithm WLGM, presented in Section V, to compute worst-case cuts under the TEC, $ATTR$, $MFST$, and AMF performance measures for a fiber plant of a major network provider [16]. In all cases, we found that the results are intuitive. We also used Algorithm WCGM, presented in Section VI, to compute worst-case circular cuts under the *MFST* performance measure for the same fiber plant. We found these circular cuts are in similar locations to their line segment counterparts. All distance units mentioned in this section are in longitude and latitude coordinates (one unit is approximately 60 miles) and for simplicity we assume latitude and longitude coordinates are projected directly to $[x, y]$ pairs on the plane. We also assume that all the link capacities are equal to 1.

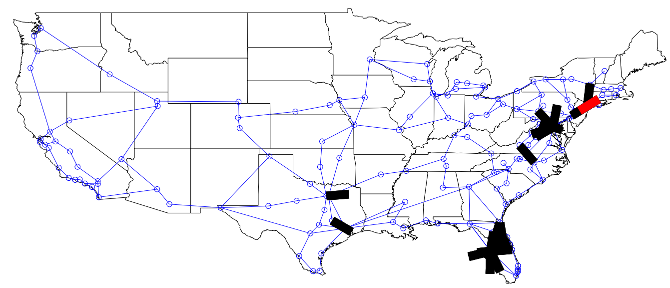

Fig. 10. Line segments cuts optimizing TEC for $h = 2$ - the red cuts maximize TEC and the black lines are nearly worst-case cuts.

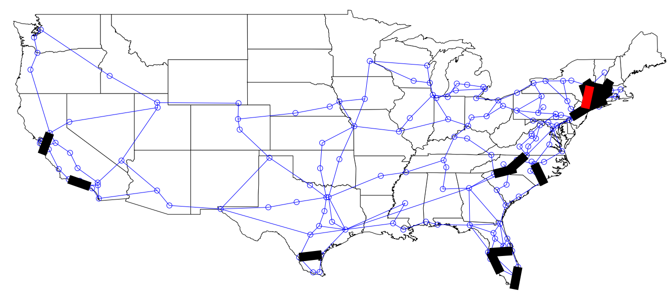

Fig. 11. Line segments cuts optimizing the $ATTR$ for $h = 2$ - the red cuts minimize $ATTR$ and the black lines are nearly worst-case cuts.

Fig. 10 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. 11 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. 12 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

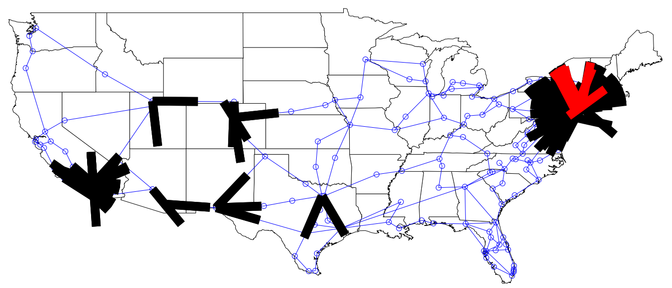

Fig. 12. Line segments cuts optimizing $MFST$ between Los Angeles and NYC for $h = 4$ - the red cuts minimize $MFST$ and the black lines are nearly worst-case cuts. Cuts which intersect the nodes representing Los Angeles or NYC are not shown.

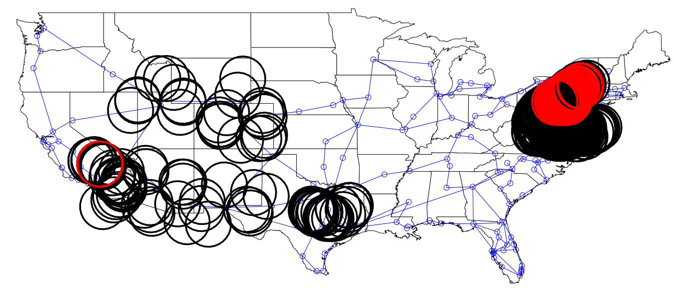

Fig. 14. The impact of circular cuts of radius 2 on the $MFST$ between Los Angeles and NYC. Red circles represent cuts that result in $MFST = 0$ and black circles result in $MFST = 1$. Cuts which intersect the nodes representing Los Angeles or NYC are not shown.

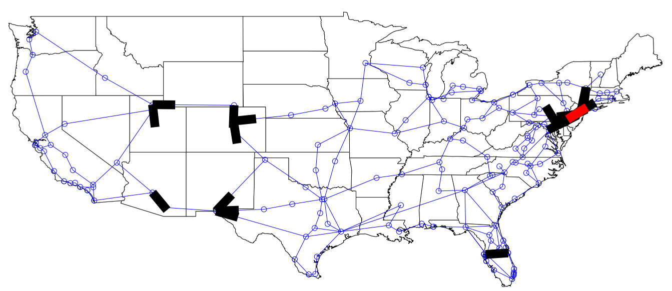

Fig. 13. Line segments cuts optimizing the AMF for $h = 2$ - the red cuts minimize AMF and the black lines are nearly worst-case cuts.

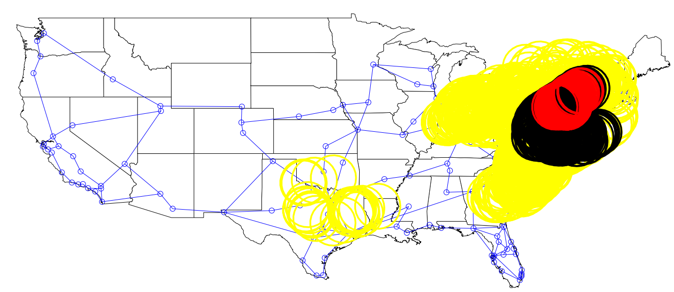

Fig. 15. The impact of circular cuts of radius 2 on the $MFST$ between Fort Worth and NYC. Red circles represent cuts that result in $MFST = 0$, black circles result in $MFST = 1$, and yellow circles result in $MFST = 2$. Cuts which intersect the nodes representing Fort Worth or NYC are not shown.

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. 13 line segment cuts of $h = 2$ which minimize the AMF performance measure. AMF is smallest 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. 14). 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. 15). Due to the complexity of the network along the east coast, the results were less straightforward than in the Los Angles-NYC case.

VIII. 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 represented by either a line segment cut or a circular cut in the physical network graph. 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 circles. For both cases, we developed polynomial-time algorithms for finding worst-case cuts. We 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.

Acknowledgments

This work was supported by a DTRA grant number HDTRA1-07-1-0004, NSF grant number CNS-0830961, the NDSEG Fellowship Program, and by a Marie Curie International Fellowship within the 6th European Community Framework Programme.

References

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

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

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

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

[5] D. Bienstock, "Some generalized max-flow min-cut problems in the plane," *Math. Oper. Res.*, vol. 16, no. 2, pp. 310–333, 1991.
[6] J. Borland, "Analyzing the Internet collapse," *MIT Technology Review*, Feb. 2008. [Online]. Available: http://www.technologyreview.com/Infotech/20152/?a=f
[7] 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.
[8] 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.
[9] 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.
[10] 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.
[11] L. K. Gallos, R. Cohen, P. Argyrakis, A. Bunde, and S. Havlin, "Stability and topology of scale-free networks under attack and defense strategies," *Phys. Rev. Lett.*, vol. 94, no. 18, 2005.
[12] O. Gerstel and R. Ramaswami, "Optical layer survivability: a services perspective," *IEEE Commun.*, vol. 38, no. 3, pp. 104–113, Mar. 2000.
[13] 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.
[14] 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.
[15] 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.
[16] Level 3 Communications, Network Map. [Online]. Available: http://www.level3.com/interacts/map.html
[17] D. Magoni, "Tearing down the Internet," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 6, pp. 949–960, Aug. 2003.
[18] 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.
[19] 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.
[20] 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.
[21] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the impact of geographically correlated network failures," *to appear in Proc. IEEE MILCOM'08*, Nov. 2008.
[22] C. A. Phillips, "The network inhibition problem," in *Proc. ACM STOC'93*, 1993.
[23] A. Pinar, Y. Fogel, and B. Lesieutre, "The inhibiting bisection problem," in *Proc. ACM SPAA'07*, 2007.
[24] W. Radasky, "High-altitude electromagnetic pulse (HEMP): A threat to our way of life," *IEEE-USA Today's Engineer*, Sept. 2007. [Online]. Available: http://www.todaysengineer.org/2007/Sep/HEMP.asp
[25] S. Rai and D. P. Agrawal, *Distributed Computing Network Reliability*. Los Alamitos, CA, USA: IEEE Computer Society Press, 1990.
[26] C. P. Warren, L. M. Sander, and I. M. Sokolov, "Geography in a scale-free network model," *Phys. Rev. E*, vol. 66, no. 5, Nov. 2002.
[27] 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
[28] 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
[29] 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.
[30] 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.
[31] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE Network*, vol. 14, no. 6, pp. 16–23, Nov.-Dec. 2000.