This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE Globecom 2011 proceedings.

Network Reliability Under Random Circular Cuts

Sebastian Neumayer and Eytan Modiano
Massachusetts Institute of Technology (MIT)
Email: {bastian,modiano}@mit.edu

***Abstract*—Optical fiber networks consist of fibers that are laid out along physical terrestrial paths. As such, they are vulnerable to geographical physical failures, such as earthquakes and Electromagnetic Pulse (EMP) attacks. Moreover, such disasters can lead to multiple, geographically correlated, failures on the fiber network. Thus, the geographical layout of the fiber infrastructure has a critical impact on the robustness of the network in the face of such geographical physical failures.**

In this paper, we develop tools to analyze network connectivity after a 'random' geographic disaster. The random location of the disaster allows us to model situations where the physical failures are not targeted attacks. In particular, we consider disasters that take the form of a 'randomly' located disk in a plane. Using results from geometric probability, we are able to approximate some network performance metrics to such a disaster in polynomial time. We present some numerical results that make clear geographically correlated failures are fundamentally different from independent failures and then discuss network design in the context of random disk-cuts.

I. Introduction

Optical fiber networks consist of fibers that are laid out along physical terrestrial paths. As such, they are vulnerable to geographical physical failures, such as earthquakes, bombs, and Electromagnetic Pulse (EMP) attacks [16], [29], [32]. Moreover, such disasters can lead to multiple, geographically correlated, failures on the fiber network. Thus, the geographical layout of the fiber infrastructure has a critical impact on the robustness of the network in the face of such geographical physical failures. In this paper, we attempt to evaluate the impact of geographically correlated failures on network connectivity with respect to a random disk. We first develop the necessary tools to evaluate network performance metrics under a geographic failure model. This, in turn, allows us to study network designs that can lessen the effects of regional disasters.

Previous works have considered the problem of finding the worst-case location for a geographic failure (represented as a disk or line segment) in a geographic network with respect to certain network connectivity measures [1], [2], [27]. This models the scenario where the network is attacked with the intention to reduce its capacity or connectivity. On the other hand, in this paper we consider the impact of a randomly located disaster on network connectivity. The random location of the disaster can model failure resulting from a natural disaster such as a hurricane or collateral (non-targeted) damage in an EMP attack. In [25] we studied this problem with respect to a randomly located line-failure; in this work we focus on a randomly located disk-failure. Similarly, in [33] the survivability of undersea cables with respect to a randomly located disk is studied, however only a two node topology was considered.

Wide-spread failures have been studied extensively in the context of logical topology [8], [9], [14], [19]. Most of these papers model the topology of the Internet as a random graph [5] and use percolation theory to study the effects of random link failures. Based on various measurements (e.g., [13]), it has been recently shown that the topology of the Internet is influenced by geographical factors such as population density [4], [17], [34]. Yet, these works do not consider the effect of failures that are geographically correlated. Network survivability has also been studied extensively (e.g., [6], [15], [20], [35]), but most of the previous work in the study of physical topology and fiber networks (e.g., [23], [24]) focus on a small number of fiber failures. In contrast, here we focus on events that induce a large number of failures in a certain geographical region (e.g., [3], [7], [16], [29]). Only recently has the impact of geographically correlated failures started to receive attention (e.g. [31], [27], [1]).

In this paper we model a random disaster by a disk of a particular radius that is 'randomly' located on the plane. Any links (modeled as line segments) that are intersected by this disk are removed from the network. In Sections II and III we introduce geometric probability and present a method that allows us to approximate joint link failure probabilities under a random disk-cut. To the best of our knowledge [25], [33], and this paper are the first to apply geometric probability techniques to network survivability. In Section IV we use these tools to show how to approximate certain metrics under this geographic failure model. In Section V we present some numerical results that make clear that geographically correlated failures are fundamentally different from independent failures and then in Section VI we present some network design problems in the context of random-disk cuts. We conclude and discuss future research directions in Section VII.

Our ultimate goal is to understand the effect of regional failures on the connectivity and reliability of a network, and to expose the design tradeoffs related to network survivability under a regional disaster. These tradeoffs may imply that there is a need to protect electronic components in certain regions (e.g., shielding against EMP attacks [16], [29]) or there may be a need to redesign parts of the network in order to make it less vulnerable to regional failures.

II. Modeling Random Circular Cuts

In this section we describe how to model random disasters using geometric probability. For simplicity, we focus only on

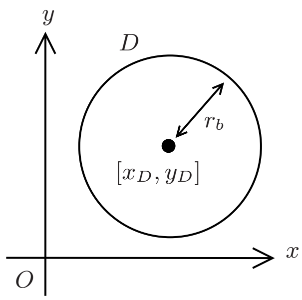

Fig. 1. Every disk in the plane of radius r_b can be parameterized by the location of its center. Denote the center of disk D as $[x_D, y_D]$.

disasters which remove links intersecting a random disk. After introducing some basic definitions from geometric probability, we review classical results which allow us to find single link failure probabilities. These results are requisite for Section III where we show how to find joint link failure probabilities to random disk-cuts.

A. Geometric Probability

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

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

Geometric probability tells us how to assign a measure to sets of disks; let this measure be denoted by μ. The rest of this section reviews results from geometric probability (see [21], [30]) that are necessary for the development of this work.

Note that every disk in the plane of radius r_b can be parameterized by the location of its center. Denote the center of disk D as $[x_D, y_D]$ (see Fig. 1). Let D_O be the disk of radius r_b centered at the origin.

We now present the definition of the measure μ.

Definition 1 (Measure of a set of disks)**.** *The measure μ of a set of disks G is defined as the integral*

$$\mu(G) = \int_G dxdy$$

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

Definition 2 (Minkowski Sum)**.** *The Minkowski sum of two sets in the plane A and B in Euclidean space, denoted by $A \oplus B$, is given by*

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

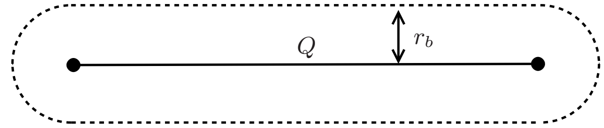

Fig. 2. The dotted shape above represents $Q \oplus D_O$. This shape is known as a hippodrome and it represents the set of all points whose shortest distance to Q is less than or equal to r_b [12]. Denote the hippodrome corresponding to a link Q and radius r_b by $H(Q, r_b)$. Note that a disk D of radius r_b intersects Q iff $[x_D, y_D] \in H(Q, r_b)$.

Intuitively, every point in the Minkowski sum $C \oplus D_O$ represents a center of a disk of radius r_b that intersects C. We will now discuss an important example. Let Q be a line segment link; consider $Q \oplus D_O$ (see Fig. 2). This shape is known as a hippodrome and it represents the set of all points whose distance to Q is less than or equal to r_b [12]. Denote the hippodrome corresponding to a link Q and radius r_b by $H(Q, r_b)$. Note that a disk D of radius r_b intersects Q iff $[x_D, y_D] \in H(Q, r_b)$.

Lemma 1 ([30])**.** *Let C be a bounded closed convex set of points in the plane, then*

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

Intuitively, every point in the Minkowski sum $C \oplus D_O$ represents a center of a unique disk of radius r_b that intersects C. Integrating over the set of centers of these disks yields the measure of $\langle C\rangle$. For example, consider a line segment link Q of length d. Now the measure of the set of disks of radius r_b that intersect Q is $\mu(\langle Q\rangle) = area(Q \oplus D_O) = area(H(Q, r_b)) = 2dr_b + \pi {r_b}^2$

B. Single Link Failures

Let $\langle Q\rangle$ and $\langle C\rangle$ be sets of disks of radius r_b in the plane such that $\langle Q\rangle \subset \langle C\rangle$. Given μ, the probability that a 'random' disk is in the set $\langle Q\rangle$ given it is in the set $\langle C\rangle$ is given by the ratio $\frac{\mu(\langle Q\rangle)}{\mu(\langle C\rangle)}$ [30]. Note that C contains the centers of all possible disk failures and is required for normalization purposes.

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

$$\begin{aligned} \Pr(Q \text{ cut} | C \text{ cut}) &= \frac{\mu(\langle Q\rangle)}{\mu(\langle C\rangle)} = \frac{area(Q \oplus D_O)}{area(C \oplus D_O)} \\ &= \frac{2dr_b + \pi {r_b}^2}{ab + 2(a+b)r_b + \pi {r_b}^2} \end{aligned} \tag{1}$$

III. Geographically Correlated Link Failures

In this section we present an algorithm that calculates the measure of disks of radius r_b intersecting only a particular set of links. This result will allow us to calculate the probability that a random disk-cut intersects a certain set of links in

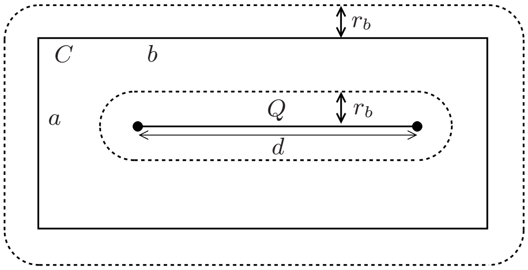

Fig. 3. The rectangle C and line segment link Q are shown in solid. The center of every disaster that intersects Q is given by the dotted hippodrome $H(Q, r_b)$. The center of every disaster that intersects C is given by $C \oplus D_O$ which is shown as the larger dotted shape.

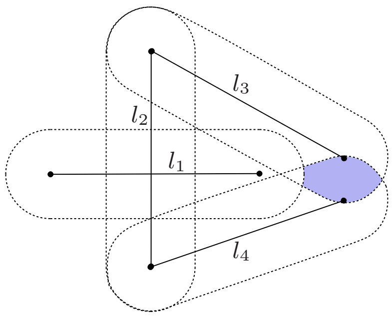

Fig. 4. The measure of disks that intersect C, l_3, and l_4 but neither l_1 or l_2 is given by the area of the shaded region above (for ease of presentation we do not picture C). This measure can be written as $\mu(\langle C\rangle \cap \langle l_3\rangle \cap \langle l_4\rangle - \langle l_1\rangle \cup \langle l_2\rangle)$ as well as $area((C \oplus D_O) \cap H(l_3, r_b) \cap H(l_4, r_b) - H(l_1, r_b) \cup H(l_2, r_b))$.

a network (e.g. links whose removal would disconnect the network). We will then use this to efficiently calculate network performance measures with respect to random disk-cuts.

Let L be the set of all line segment links in the network and C be a convex polygon that contains L. Consider some set of links $K \subset L$. We wish to find the measure of all disks of radius r_b that intersect C and every link in K but intersect no links in $L-K$. See Fig. 4 for an example. This measure is given by $\mu\left(\langle C\rangle \cap (\cap_{k\in K}\langle k\rangle) - \cup_{q\in(L-K)}\langle q\rangle\right)$. It is clear that a disk D belongs to this measured set iff i) $[x_D, y_D] \in C \oplus D_O$, ii) $[x_D, y_D] \in H(k, r_b)\ \forall k \in K$, and iii) $[x_D, y_D] \notin H(q, r_b)\ \forall q \in (L-K)$. So, this measure can also be written as $area\left((C \oplus D_O) \cap (\cap_{k\in K} H(k, r_b)) - \cup_{q\in(L-K)} H(q, r_b)\right)$. For ease of presentation we abuse notation and denote this measure by $area(K)$.

Definition 3 ($area(K)$)**.** *Let area(K) be given by the measure of all disks of radius r_b that intersect C and every link in K but intersect no links in $L-K$.*

A. Approximation

We note that finding $area(K)$ seems difficult because it requires finding the area of intersections and unions of hippodromes. In the following we describe a method for approximating $area(K)$ which is based on approximating hippodromes by polygons for which there are known methods to calculate intersections, unions, and area. We approximate $H(l, r_b)$ by the inscribing polygon $\hat{H}_n(l, r_b)$ such that $H(l, r_b)$ shares the line segment portion of its boundary with $\hat{H}_n(l, r_b)$ and each end of $\hat{H}_n(l, r_b)$ forms half of a regular $2n$-sided polygon (see Fig. 5). Let $\widehat{area}_n(K)$ be defined the same as $area(K)$ except that every hippodrome is replaced by its polygon approximation. Using techniques for finding the intersection, union, and area of polygons [28], we can find $\widehat{area}_n(K)$ in polynomial time.

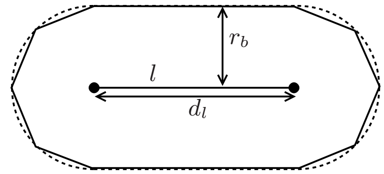

Fig. 5. We approximate $H(l, r_b)$, shown as a dashed hippodrome above, by the inscribing polygon $\hat{H}_n(l, r_b)$ such that $H(l, r_b)$ shares the line segment portion of its boundary with $\hat{H}_n(l, r_b)$ and each end of $\hat{H}_n(l, r_b)$ forms half of a regular $2n$-sided polygon. The solid polygon above is $\hat{H}_4(l, r_b)$.

We now present a lemma that shows $\widehat{area}_n(K)$ is a good approximation for $area(K)$ for large enough n. A proof is omitted due to space constraints and can be found in [26].

Lemma 2. $\lim_{n\to\infty} \widehat{area}_n(K) = area(K)\ \forall\ K \subset L$

IV. Evaluating Network Reliability Metrics

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

A. Network Model and Metrics

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

We now introduce two important network performance metrics which are evaluated after the removal of the intersected links. We will use tools from geometric probability to evaluate 'average' values of these metrics with respect to a random disk-cut. The first metric we consider is the total capacity of intersected links. Let this metric be denoted by TC. The other metric we consider is the fraction of node pairs that remain connected to each other. Let this metric be denoted by $A2TR$.

B. Evaluation of the Metrics

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

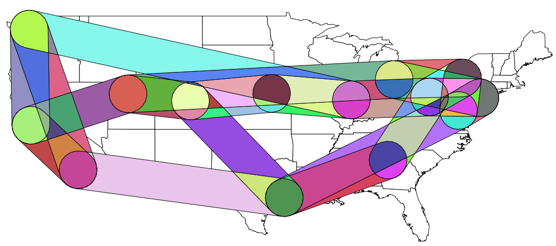

Fig. 6. This figure shows the hippodromes and related intersections with respect to NSFNET [22] and a circular cut of $r_b = 2$.

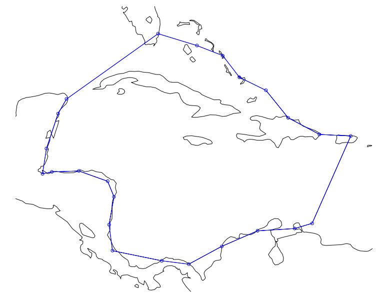

Fig. 7. ARCOS-1 network circa 2009 [10].

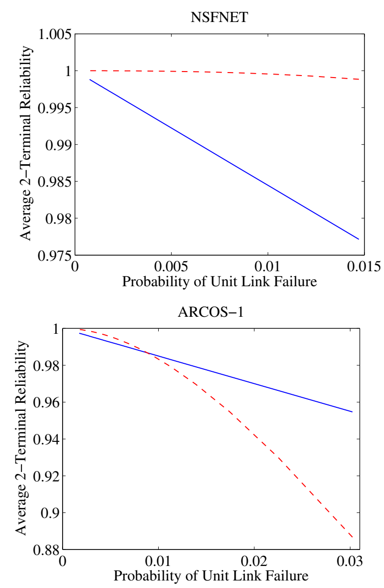

Fig. 8. The solid line shows $A2TR$ versus the probability a unit (latitude/longitude) of fiber is cut by a random disk of $r_b = 2$. The dashed line shows $A2TR$ assuming links fail *independently* such that links fail with the same probability as in the random disk-cut case.

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

$$\sum_{K \in P} \frac{area(K)}{area(C \oplus D_O)} Y(K) \tag{2}$$

Section III shows how to approximate $area(K)$ in polynomial time. $Y(K)$ for the performance metrics can also be calculated in polynomial time. In the following, we apply the theory of arrangements to show that the size of P grows polynomially with respect to N. For technical reasons this theory cannot be directly applied to this setting and requires modification.

Let ∂ denote the boundary of a set. Consider the set of curves $R = \partial C \cup \{\partial H(l, r_b) | l \in L\}$. These curves partition C, the set containing the network, into maximally connected regions called faces that are bounded by the curves in R. By enumerating these faces, we can enumerate every element in P (since every disk in a particular face intersects the same links). Arrangements, a computational geometry tool, allow us to enumerate the faces of a set of curves in $\mathbb{R}^2$ in polynomial time. However, the theory requires that every pair of curves intersect in a finite number of locations [11] which does not hold in our setting. Nonetheless, the theory can be applied with a minor perturbation to the geometry.

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

V. Numerical Results

In this section we evaluate some network metrics using the results of the previous section. We consider NSFNET as found in 1991 [22] and the ARCOS-1 ring network [10]. The NSFNET network we consider has 14 nodes and connects major universities across the U.S. (see Fig. 6). ARCOS-1 has 24 nodes and connects regions on the Dominican Republic, Florida, Mexico, Panama, and Venezuela (see Fig. 7). All distance units mentioned in this section are in longitude and latitude coordinates (one unit is approximately 60 miles) and for simplicity we assume latitude and longitude coordinates are projected directly to [x, y] pairs on the plane. We assume that all the link capacities are equal to 1. We also assume each network is contained within a rectangular set C.

A. Independent Versus Correlated Failures

Using the results of Section IV, we calculate $A2TR$ of NSFNET and ARCOS-1 to random-disk cuts of $r_b = 2$ while the size of C varies. The size of C is varied to change the probability a unit of fiber is cut. So we can plot $A2TR$ versus the probability a unit of fiber is cut. See Fig. 8 for results. Note the linear form of the result in the figure; this agrees with Equation 2 since $1/area(C \oplus D)$ is proportional to the probability a unit of fiber is cut.

Next, we calculate $A2TR$ of the networks assuming independent link failures such that links fail with the same probability as in the random disk-cut case. Thus the probability a link fails is still a function of its length, however links fail independently. Since the total number of links is small in each network, calculating $A2TR$ by enumerating all possible failures is still feasible (possible failures are exponential in number of links). Note the total expected number of removed links is the same for both the independent and geographically failure models. See Fig. 8 for results.

Notice that in NSFNET $A2TR$ under independent failures is greater than in the case of random disk-cuts. Perhaps this is because in most cases at least three links must fail independently to disconnect the network; however a disk that intersects a node is guaranteed to disconnect the network. Since most backbone networks are likely to be well connected, we expect a random disk-cut to lead to lower $A2TR$ than independent link failures in this type of mesh network setting. We also note that similar results were found for the random line-cut setting [25].

Looking at the results for the ARCOS-1 network we see the opposite tendency; $A2TR$ under independent failures is typically less than the case of random disk-cuts. Perhaps this is because a single disk that intersects ARCOS-1 usually only removes two adjacent links creating components of size 1 and $|N|-1$ (where $|N|$ is the number of nodes) whereas just two independent link failures on opposite sides of the ring create components of size $|N|/2$ and $|N|/2$ (which results in lower A2TR).

B. Multiple Failures

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

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

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

In [27] we propose an algorithm to evaluate network reliability metrics after an *intentional* disk failure. To calculate a network metric after sequential intentional failures, we simply apply the algorithm in [27] iteratively.

Fig. 9 shows the results for multiple failures, both intentional and random for NSFNET (similar results for ARCOS-1 are not shown). As expected, the plots are sub-linear since each additional failure is being placed on a smaller network. Note that random failures result in much less disruption than intentional failures.

VI. Network Design Under Random Circular Cuts

In this section we discuss some network design problems in the context of random disk-cuts. In all the proposed problems the location of every node is fixed; the problem is to find a set

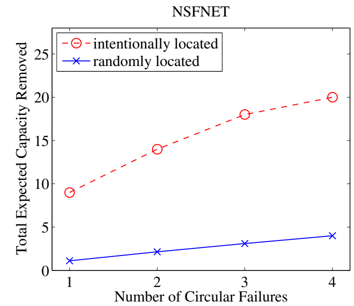

Fig. 9. The solid line shows the total expected capacity removed versus the number of *randomly* located circular disasters of $r_b = 2$. Using the algorithm in [27], the dashed line shows the total capacity removed versus the number of *intentionally* located circular disasters of $r_b = 2$.

of links most robust to some metric under some constraints. In the following, let N be a set of nodes fixed on the plane and assume all links are represented as line segments between the nodes.

(i) A reasonable goal is to design a connected network with the least expected number of links cut by a random disk of radius r_b. By Eqn. 1 and linearity of expectation, the expected number of links cut is proportional to $\pi r_b^2(|N|-1)+2r_b\sum_{l\in U}d_l$ where U is the set of links chosen and d_l is the length of link l. So, this problem reduces to minimizing the total length of links in the network while ensuring the graph is connected. This is equivalent to finding a Euclidean minimum spanning tree of N which can be done in polynomial time. Note however that the resulting network is not robust because a single link failure will disconnect it. We also note this is the same result we get for optimizing for random line-cuts [25]. If fact, it can be shown that this result will hold for any convex-shaped cut.

(ii) We next consider ring networks. A reasonable goal is to design a connected ring network with the least expected number of links cut by a random disk of radius r_b. As before, the expected number of links cut is affine in the total length of the links. So this problem reduces to finding the minimum length Hamiltonian cycle. This is equivalent to the Euclidean traveling salesman problem which is hard to compute [18]. We again note this is the same result we get for optimizing for random line-cuts [25]. If fact, it can be shown that this result will hold for any convex-shaped cut.

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

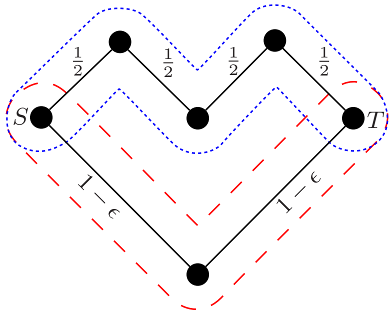

Fig. 10. The bottom path is shorter than the top path by 2ϵ. However the top path is more robust to a random circular failure as compared to the bottom path since the blue dotted shape has less area than the red dashed shape. This shows the shortest path is not necessarily the most robust to failure.

An interesting example is given in Fig. 10. The uppermost path here gives the most robust path to a random disk failure; however, the bottom path is the shortest. This shows the shortest path is not necessarily the most robust to failure.

VII. Conclusions

We focused on randomly located circular cuts on a network which can model the 'random' nature of a natural disaster or collateral damage. Using tools from geometric probability we demonstrated how to compute failure probabilities and showed how to approximate some network performance metrics in polynomial time under this failure model. We then presented some numerical results that make clear geographically correlated failures are fundamentally different from independent failures and then discussed some network design problems. Some future research directions include the consideration of convex cuts (e.g., oval cuts) and robust network design in the face of geographical failures.

Acknowledgments

This work was supported by NSF grants CNS-0830961 and CNS-1017800, and by DTRA grants HDTRA1-07-1-0004 and HDTRA-09-1-005.

References

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

[2] ——, "Network vulnerability to single, multiple, and probabilistic physical attacks," in *Proc. IEEE MILCOM'10*, Nov. 2010.

[3] 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.

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

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

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

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

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

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

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

[13] 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.

[14] 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.

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

[16] 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.

[17] 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.

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

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

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

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

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

[23] 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.

[24] 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.

[25] S. Neumayer and E. Modiano, "Network reliability with geographically correlated failures," in *Proc. IEEE INFOCOM'10*, Mar. 2009.

[26] ——, "Network reliability with random circular cuts," Massachusetts Institute of Technology, EE, Tech. Rep. 2011-02-28, Feb. 2009. [Online]. Available: http://www.mit.edu/%7Ebastian/NRWRCC-Tech-Report.pdf

[27] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of the fiber infrastructure to disasters," *IEEE/ACM Transactions on Networking*, no. 99, pp. 1–1, To Appear.

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

[29] 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

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

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

[32] 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

[33] 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.

[34] 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.

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