Fast Enumeration of Regional Link Failures Caused by Disasters With Limited Size

János Tapolcai, *Senior Member, IEEE*, Lajos Rónyai, Balázs Vass, *Member, IEEE*, and László Gyimóthi

***Abstract*— At backbone network planning, an important task is to identify the failures to get prepared for. Technically, a list of link sets, called Shared Risk Link Groups (SRLG), is defined. The observed reliability of network services strongly depends on how carefully this list was selected and whether it contains every high-risk failure event. Regional failures often cause the breakdown of multiple elements of the network, which are physically close to each other. In this article, we show that operators should prepare a network for only a small number of possible regional failure events. In particular, we give an approach to generate the list of SRLGs that hit every possible circular disk shaped disaster of a given radius r. We show that this list has $O((n+x)\rho_r)$ SRLGs, where n is the number of nodes in the network and x is the number of link crossings, and ρ_r is the maximal number of links that could be hit by a circular disaster of radius r. We give a fast polynomial algorithm to enumerate the list of SRLGs and show that its worst-case time complexity is asymptotically optimal under some practical restrictions. Finally, through extensive simulations, we show that this list in practice has a size of $\approx 1.2n$.**

***Index Terms*— Disaster resilience, network failure modeling, shared risk link groups, SRLG enumeration, computational geometry.**

I. Introduction

BACKBONE networks are designed to protect a specific pre-defined list of failures, called Shared Risk Link Groups (SRLG) [2]. SRLG describes the relationship between links with a shared vulnerability. For example, links with shared fiber cable or conduit have a chance to fail simultaneously, or network devices with shared power-sharing, etc. The SRLGs can subtly hit the network, as each link could belong to several SRLGs. Unfortunately, SRLGs are not self-discoverable in practice [3]; thus the mapping of links to SRLGs should be defined by the network operators. Operators must very carefully decide the list of SRLGs because leaving out one likely simultaneous failure event will significantly degrade the observed reliability of the network. There is a high number of severe network outages witnessed in the last decades [4]–[7]. This present a clear evidence that selecting the proper list of SRLGs is still a challenging problem to solve [8]–[16]. To fill this gap in reliable network design, this article proposes a systematic approach for selecting the list of SRLGs. The general idea in defining SRLGs is that links close to each other have a chance for simultaneous failure. Thus we list sets of links close to each other. The main finding of this study is that surprisingly the number of such SRLGs is not too high in practice.

After the list of SRLGs is defined, the network is designed to be able to recover in case of a single SRLG failure, such that every connection operates again after a very short interruption. Current backbone networks are required to fulfill a very high level of service availability, and they can handle an arbitrary list of SRLGs. The only practical limitation is that the list of SRLGs cannot be extremely long to keep the routing algorithms, the failure localization scheme, and the failure states scalable. There is no performance guarantee when a network is hit by a disaster that hits links that are not a subset of an SRLG. Thus, the best practice is to list every *single link or node failure* as an SRLG. Here the concept is that the disaster first hits a single network element for whose protection the network is already pre-configured. Note that the list of SRLGs can be extended to protect against possible multiple failures. Several papers studied its limitation [8]–[16] that the networks have severe outages when almost every equipment in a vast physical region gets down as a result of a disaster, such as earthquakes, hurricanes, tsunamis, tornadoes, etc. For example, the 7.1-magnitude earthquake in Taiwan in Dec. 2006 caused simultaneous failures of six submarine links between Asia and North America [4], the 9.0 magnitude earthquake in Japan on March 2011 impacted about 1500 telecom switching offices due to power outages [5] and damaged undersea cables, the hurricane Katrina in 2005 caused severe losses in Southeastern US [6], Hurricane Sandy in 2012 created a power outage which silenced 46% of the network in the New York area [11], [17].

Heavy rainfalls, or in general weather-based region disruptions, can bring out correlated temporal failures of high

Manuscript received February 25, 2019; revised May 11, 2020; accepted July 2, 2020; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor G. Zussman. Date of publication August 13, 2020; date of current version December 16, 2020. The work of János Tapolcai was supported in part by the National Research, Development and Innovation Fund (OTKA) of the Development and Innovation Office of Hungary (NKFIH) under Grant 128062 and Grant 124171. The work of Lajos Rónyai was supported in part by the NKFIH/OTKA under Grant K115288. The research reported in this article was supported by the BME Artificial Intelligence TKP2020 IE grant of NKFIH Hungary (BME IE-MI-SC TKP2020) and in part by the NKFIH under Grant NKFIH-115288. An earlier version of the article appeared at IEEE INFOCOM 2017. *(Corresponding author: Balázs Vass.)*

János Tapolcai, Balázs Vass, and László Gyimóthi are with the MTA-BME Future Internet Research Group, Budapest University of Technology and Economics (BME), 1111 Budapest, Hungary, and also with the MTA-BME Information Systems Research Group, Department of Telecommunication and Media Informatics, Budapest University of Technology and Economics (BME), 1111 Budapest, Hungary (e-mail: tapolcai@tmit.bme.hu; balazs.vass@tmit.bme.hu; gyimothi@tmit.bme.hu).

Lajos Rónyai is with the Institute of Computer Science and Control, BME, 1111 Budapest, Hungary, and also with the Department of Algebra, BME, 1111 Budapest, Hungary (e-mail: lajos@info.ilab.sztaki.hu).

Digital Object Identifier 10.1109/TNET.2020.3009297

capacity wireless links (as, e.g., in Wireless Mesh Networks) in a small region. Another critical reason for disruptions on a massive scale the network operators need to prepare for is related to intentional human activities, such as the bombing or use of weapons of mass destruction attacks, electromagnetic pulse attacks. An electromagnetic pulse attack is an intense energy field that can instantly overload or disrupt numerous electrical circuits, thereby affecting networking equipment within a large geographic area [18]. Submarine cables are vulnerable to human activities such as fishing, anchors, and dredging [19]. These types of failures are called *regional failures* which are simultaneous failures of nodes/links located in specific geographic areas [8]–[16]. It is still a challenging open problem how to prepare a network to protect against such failure events, as their location and size are unknown at the planning stage. Intuitively, the number of possible regional failures can be huge. In the article, we propose a solution to this problem with a technique that can significantly reduce the number of possible SRLGs that should be added to cover the majority of the potential regional failures.

In the remainder of the article, we will call events that bring down the network in a geographic area simply as *disasters*, indifferent to their cause (natural or human-made). We define a regional failure as a failure of multiple network elements in a geographic area, which can have any location, size, and shape. In particular, we are interested in regional failures caused by single disasters. We will consider the size of disaster as its most important property. To measure the size of a regional failure, we compute the smallest circular disk (representing a disaster) that hits every failed link, and the radius of this circular disk represents its size. In this study, we are interested in enumerating the SRLGs of disasters with a given maximum size. We assume operators prepare their network to survive the failure of each SRLG, for example, by allocating SRLG-disjoint working and protection paths for each connection. Since an SRLG is a worst-case scenario, there is no need to have an SRLG, which is a subset of another SRLG. In particular, we are interested in two versions of the problem. In the first version, we list every possible failure the network can have due to a circular disk shaped disaster of a given radius r. In the second version of the problem, we assume the radius of the circular disaster is not a network-wide parameter but depends on the area. For example, the radius is larger in flat regions and smaller in hilly areas. In this case, we list every possible circular disk failures with radius at most r and leave the operator to filter out the unrealistic ones.

The main **contribution of this article** is a reduction of the number of SRLGs subject to disasters by applying computational geometric tools based on the following two assumptions: (1) The network is a geometric graph $G(V, E)$ embedded in a 2D plane, and $n = |V|$ denotes the number of nodes in the network. (2) The shape of the disaster is assumed to be a circular disk of radius r having an arbitrary center position. We show that with these assumptions:

- The number of SRLGs is small, close to $1.2n$ in a typical backbone network topology, which surprisingly does not depends on the radius r.
- We refine the bound on the number of SRLGs by introducing some practical properties of the graph: x, which is the number of link crossings of the network, ρ_r is the maximal number of links that could be hit by a circular disaster of radius r. In backbone networks, x is a small number as typically, a network node is also installed on each link crossing (similarly to road networks [20]). At the same time, ρ_r represents a density of the topology, which should not depend on the network size. Using these parameters the number of SRLGs is $O((n+x)\rho_r)$ We also give an artificial example to illustrate that these bounds are tight.
- We provide low polynomial algorithms to enumerate the SRLGs that runs in $\Theta(n \log n)$ if $\rho_r = O(1)$, $\mu = O(\log n)$, and x and x' are $O(n)$, where x' is the number of link crossings in G', where all the edges are elongated by $3\sqrt{2}r$ in both directions and μ is the square mean of numbers v_e for all $e \in E$, where v_e is the number of $w \in V \cup X$ such that $d(w, e) \leq 3r$. These assumptions hold when r is small compared to the geographical network diameter.
- Compared to prior art we handle parallel edges and collinear node triples.

Network operators can design their networks using the obtained SRLG list to protect regional and random failures. Backbone networks designed according to our new failure model should have higher reliability and leave way fewer failures unprotected at the optical layer. We believe the article contributes to closing the gap between the current SRLG based pre-planned protection and regional failures.

The article is organized as follows. In Sec. II we overview the related work and explain how our approaches can contribute to the prior art, in Sec. III we provide a mathematical definition of the problem and show some basic results. In Sec. IV we provide bounds on the number of SRLGs, which we improve in Sec. V and present our algorithm. In Sec. VI we present our numerical evaluation of real backbone networks. Finally Sec. VII concludes the article.

II. Related Work

With somewhat different motivations, similar computational geometric ideas were used in articles focusing on the most vulnerable points of physical infrastructure (communication networks or power grids [21]) to disasters. Our objective is more general as we want to enumerate all candidate failures, instead of searching for the most vulnerable according to some metric.[1] The network is embedded in the Euclidean plane and the disasters are modeled either as a disk around its epicenter (circular) [8], [22], [23], line segments [8], ellipse [24]

[1]In fact, if the metric is monotone (i.e., for any link set $E_1 \subseteq E_2$, the failure of E_2 is *worse* than the failure of E_1), the worst SRLG that can be hit by a disaster with radius r will be part of the set of exclusion-wise maximal SRLGs that can be hit by a disaster with radius r. In other words, it can be found by simply searching for the worst SRLG in the list of maximal SRLGs.

or polygons (rectangle, square, or equilateral triangle) [24]. Technically, these articles also list the candidate failures and evaluate the vulnerability metric of the residual network in case of each candidate failure, all included in [25], which is a recent tutorial on SRLG enumeration. Note that, our approach computes $O((n+x)\rho_r)$ SRLGs (in practice ρ_r is constant and $x \ll n$), while the best known general worst-case bound was $O(n^4)$ [22], which would be $O((n+x)^2)$ using our estimations with x. Besides, our approach can be used to compute the list of candidate failures for circular disasters with a varying radius.

The following vulnerability metrics are investigated: (1) the point with the maximum number of affected links [8], [22], which is ρ_r. (2) the point with the maximum average two-terminal reliability between every node-pair [8], [22]–[24]. (3) the point with the maximum average all-terminal reliability [14], [23], which allows the identification of network areas that can disconnect any component in the network. (4) the point with the maximum average value of the maximum flow between a given pair of nodes [8]. (5) the point with maximal average shortest path length between every pair of nodes [14], [24], (6) survivability as a measure of weighted spectrum based on the eigenvalues of the normalized Laplacian of a graph [14], (7) network criticality which is determined from the trace of the inverse of the Laplacian matrix and can be related to the node and link betweenness [14].

A particular case of our problem is investigated in [26] where the goal is to list all the spatially-close fiber segments. In their model, the links are not only straight line segments but can be a series of line segments connecting a set of corner points. In our model, treating the corner points as degree 2 nodes, the grouping of the spatially-close fiber segments can be directly computed.

The idea of defining SRLGs for disasters was also proposed in [12]; however, the SRLGs (called disaster zones) were decided manually. For example, in the 24-node US topology, they determine 15 distinct SRLGs for earthquakes and 19 distinct SRLGs for tornadoes by matching a seismic hazard map and a tornado activity map with US topology considering that the damage of earthquakes and hurricanes (clustered in a region) may span up to 96 and 160 km, respectively.[2] Besides, the 10 most-populated US cities and Washington DC as possible mass destruction targets are added. It is in total 45 SRLGs, while our approach automatically lists 20-30 failures depending on the radius (see also Fig. 9b).

Our approach can be used as a tool for any studies where the set of potentially vulnerable geographic cuts are taken as input, such as for multilayer networks [29], SRLG disjoint paths [30], etc.

Related to our work is the research in computational geometry on the smallest intersecting ball problem [31], [32], which has its origins in the classical 19th-century problem of Sylvester [33] about the smallest enclosing circle for a given set of points in the plane.

[2]Papers [27] and [28] deal with earthquakes at planning stage.

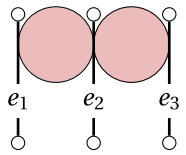

Fig. 1. In the figure above, the solid circular disks are disasters with radius r, $d(e_1, e_2) = d(e_2, e_3) = 2r$, while $d(e_1, e_3) = 4r$. The set of regional failures is $\mathcal{H}_r = \{\{e_1\}, \{e_2\}, \{e_3\}, \{e_1, e_2\}, \{e_2, e_3\}\}$. The set of maximal regional failures is $\mathcal{S}_r = \{\{e_1, e_2\}, \{e_2, e_3\}\}$.

III. Problem Definition and Basic Results

The input is a real number $r \geq 0$ and an undirected connected graph $G = (V, E)$ embedded in the $2D$ plane, where V denotes the set of nodes and E the set of edges (which are also called links). Let $n := |V|$ and $m := |E|$. We assume $n \geq 3$. The edges of G are embedded as line segments, which we call *intervals* in the geometric proofs.[3] A *disk* with centre point p *hits* an edge e if its distance to p is at most r.

Definition 1: A ***regional failure*** *F is a non-empty subset of E, for which there exists a disk with radius r hitting every edge in F.*

Note that the failure of node v is modeled as the failure of all edges incident to node v. Therefore listing the failed nodes beside listing failed edges would not give us additional information from the viewpoint of connectivity.

Definition 2: Let $\mathcal{H}_r$ be the ***set of regional failures*** *of a network for a given radius r.*

According to Def. 1, a subset of a regional failure is also a regional failure. Thus, $\mathcal{H}_r$ is a downward closed set minus the empty set.

An *SRLG* is a regional failure the network is prepared for. Recall the network can recover if an SRLG or a subset of links (and nodes) in the SRLG fail simultaneously. In other words, if a regional failure F is listed as an SRLG, then there is no need to list any subset of the links $F' \subsetneq F$ as a new SRLG. Our goal is to define a set of SRLGs which covers every possible regional failure and which is of minimal size.

Definition 3: Let $\mathcal{S}_r \subseteq 2^E$ denote the set of SRLGs, for which

$$\mathcal{S}_r = \{F \quad \textit{is a regional failure and there is no regional failure } F' \textit{ such that } F' \supsetneq F\}\ . \tag{1}$$

In other words, the set of SRLGs $\mathcal{S}_r$ is a set of failures caused by disks with radius at most r in which none of the failures is contained in another. Figure 1 illustrates Definitions 1-3. Note that $\mathcal{H}_r$ is the set of regional failures, which is the downward closed extension of $\mathcal{S}_r$ minus the empty set. In combinatorics, a Sperner system is a family of sets in which none of the sets is contained in another. A Sperner family is also sometimes called an independent system or a clutter. Note

[3]The case, when edges are considered to be embedded as polygonal chains between their endpoints consisting of at most a constant number of line segments, can be also handled in polynomial time based on our results via splitting the polygonal chains up into line segments, running our proposed algorithm (sketched in Table I) for the resulting problem instance, merging the line segments of each polygonal chain, and finally, filtering out the non-maximal sets.

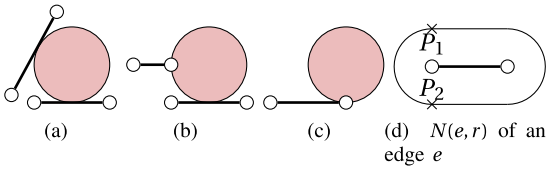

Fig. 2. Case (a),(b) and (c) of Thm. 1 and the neighbourhood $N(e,r)$ of an edge e.

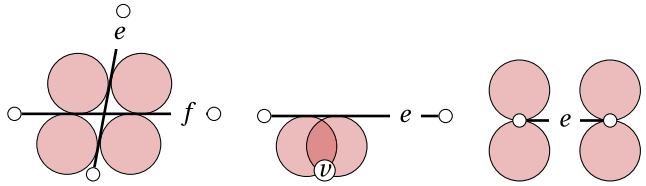

Fig. 3. The circular disasters examined in Thm. 1.

that, $\mathcal{S}_r$ is a Sperner system. Due to the minimality of SRLGs, we have the following proposition.

Proposition 1: For each ***SRLG*** $F \in \mathcal{S}_r$, $F \subseteq E$, *there is a circular disk c of radius r such that F is exactly the set of edges hit by c.*

Let r be a tiny positive number. In this case, the list of possible regional failures consists of every *single link or node failure* and link crossings. In other words, our model is a generalization of the 'best practice.' The corresponding Sperner system can be the set of single node failures, i.e., $|\mathcal{S}_r| = n + x$, where is x is the number of edge crossings. Informally speaking, protecting node failures is sufficient to protect link failures as well.

We aim to determine the set $\mathcal{S}_r$. At first glance, it is not clear that the cardinality of $\mathcal{S}_r$ is 'small.' We will prove polynomial upper bounds on $|\mathcal{S}_r|$, and we will show that $|\mathcal{S}_r|$ is $\sim n$ in practice.

To estimate the size of the SRLG list, let ρ_r denote the maximum number of edges a disk with radius r can hit in the plane, i.e., for every failure F caused by a disk with radius r, $|F| \leq \rho_r$. We observe that if $\rho_r = O(\log n)$ then there is a polynomial blowup when we switch from $\mathcal{S}_r$ to $\mathcal{H}_r$, as $|\mathcal{H}_r| \leq |\mathcal{S}_r| 2^{\rho_r}$. We often treat $\mathcal{S}_r$ as a compact representation for $\mathcal{H}_r$. It is also immediate that from $\mathcal{H}_r$ we can obtain $\mathcal{S}_r$ by $O(|\mathcal{H}_r|^2)$ comparisons of subsets of E.

We say a disk c hits a set of edges E_c if it hits all the edges in E_c. Note that several disks can hit the same set of edges.

First, we give a slight variant of Lemma 9 from [8]. Our assumptions allow somewhat more general topologies with more than 2 collinear points. The segments $e \in H$ are assumed to be nondegenerate.

Theorem 1: Let r be a positive real, and H be a nonempty set of intervals (i.e., edges) from $\mathbb{R}^2$ which is hit by a circular disk of radius r. Then there is a disk c of radius r which hits the intervals of H such that at least one of the following holds (see Fig. 2 for illustrations).

(a) There are two non-parallel intervals in H such that c intersects both of them in a single point. These two points are different.

(b) There are two intervals in H such that c intersects both of them in a single point. These two points are different, and one of them is an endpoint of its interval.

(c) Disk c touches the line of an interval $e \in H$ at an endpoint of e.

Proof: For a line segment e on the plane and a nonnegative real number r the *r-neighborhood*[4] $N(e,r)$ of e is defined as the set of all points P on the plane which have distance at most r to (some point of) e. It is immediate that $N(e,r)$ is a closed convex subset (see Fig. 2d) of the plane.

Consider the boundary B of the intersection

$$\cap_{e \in H} N(e,r). \tag{2}$$

The points of B are obviously in the union of the boundaries of the neighborhoods $N(e,r)$, where $e \in H$. The union is composed of a finite number of line segments and half circles. The circular arcs belong to circles of radius r centered at endpoints of line segments $e \in H$. We distinguish two cases.

(1) B has a point R which is on a halfcircle arc of the boundary on $N(e,r)$ for some $e \in H$. Let c_R be the disk of radius r centered at R. If R is an endpoint (P_1 or P_2 in Fig. 2d) of the halfcircle, then (c) is satisfied for c_R. We can thus assume that R is an inner point of the halfcircle connecting P_1 and P_2, and $P_i \notin B$. From the fact that B is closed, we obtain that there exists a point R' on the circular arc RP_2 which is in B, but no point of the open $R'P_2$ arc is in B. Then there must be an $f \in H$ such that $N(f,r)$ passes through R' but does not contain a larger arc $R'R''$ from $R'P_2$. Then R' is on the boundary of $N(f,r)$. We argue that (b) holds for $c_{R'}$ and the intervals e, f. This is immediate if the tangent lines to $N(e,r)$ and $N(f,r)$ at R' are different. If they are the same line ℓ then e and f must be in different halfplanes defined by ℓ, hence $e \cap f = \emptyset$ and hence (b) holds for $c_{R'}$. This reasoning settles the case (1). Note that we can also assume now that $|H| > 1$.

(2) No point of B is on a circular arc form the boundary of $N(e,r)$, with $e \in H$. Then B is a (possibly degenerate) polygon composed of some line segments. Let R be a vertex of polygon B, and $e \in H$ be a segment such that R is an interior point of one of the line segments on the border of $N(e,r)$. Let ℓ be the line of this latter segment. The fact that R is a vertex of B implies that there must be another segment $f \in H$ such that one of the line segments on the boundary of $N(f,r)$ passes through R and the line ℓ' of this segment is different from ℓ. Indeed, otherwise, for every $g \in H$ there would be an open interval form ℓ containing R in $N(g,r)$, which contradicts the extremality of R. As e is parallel to ℓ and f is parallel to ℓ', we infer that (a) holds for c_R. ■

IV. Bounds on the Number of SRLGs

Lemma 1: Let H' be a set of intervals from $\mathbb{R}^2$, $|H'| \leq 2$, and r be a positive real number. Then every circular disk described in Thm. 1 for $H = H'$ can be determined in $O(1)$ time.

Proof: Easy elementary geometric discussion of cases (a), (b) and (c) of Thm. 1. See Fig. 3 for illustration.

[4] called hippodrome in [22].

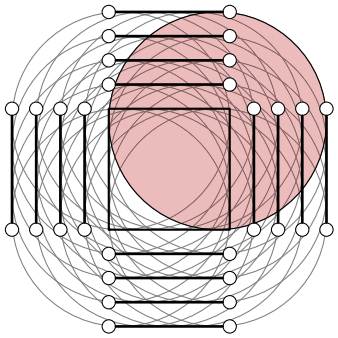

Fig. 4. An example topology ($k = 4$) where the number of maximal SRLGs hit by circular disk shaped disasters is $\Omega(m^2)$ or $\Omega((n+x)\rho_r)$.

Note that there can be at most 4 circles that intersect two line segments, as shown in Fig. 3(a), and at most two circles intersecting a line segment and a single point, as shown in Fig. 3(b), and four circles can touch a line at endpoints, as shown in Fig. 3(c). ■

From Thm. 1 and the argument of Lemma 1 we obtain the following upper bound on the number of SRLGs.

Corollary 1: $|\mathcal{S}_r| \leq 4\binom{m}{2} + 4m + 2mn$.

Note that, the graphs of Claim 1 demonstrate that the above bound is asymptotically tight.

A. Worst Case Graph

Claim 1: The graph sketched in Fig. 4 has at least $\frac{n^2}{64}$ *maximal regional failures of a radius* k.

Proof: Here we construct a set of n segments whose graph is planar (there are no edge intersections), and for a suitable radius r it has at least $\frac{n^2}{64}$, in particular a quadratic number of, incomparable failure events.[5]

Let k be a positive integer. We consider a collection of $4k$ axis parallel line segments in $\mathbb{R}^2$. We start out with the four edges of the square of edge size k whose bottom left corner is at the origin $O = (0,0)$. We consider the bottom edge connecting O to $(k,0)$, and put its copies translated i units downwards, for $i = 1,\ldots,k$ into our set of segments. For example for $i = 2$ we obtain the segment from $(0,-2)$ to $(k,-2)$. This way we obtained k segments. Similarly we translate the upper edge (from $(0,k)$ to (k,k)) of the square by i units upwards for $i = 1,\ldots,k$. These are k additional horizontal segments. We do the same in the vertical direction: we consider k translates to the left of the left edge of our starting square, and k translates to the right of the right edge of the square. We have $4k$ nonintersecting line segments of length k. The configuration for $k = 4$ is shown in Fig. 4. Consider now a disk $c = c(i,j)$ of radius k centered at the point (i,j), where i,j are integers, $0 \leq i,j \leq k$. We readily see that c intersects exactly i of the right vertical segments and $k-i$ of the left vertical segments. Similarly, c intersects exactly j of the upper horizontal edges and $k-j$ of the lower horizontal edges. We infer that no two disks of the form $c(i,j)$ can hit the same set of edges. This implies that there are at least $(k+1)^2$ maximal failure events with radius k. The number of vertices is $n = 8k$. The number of such maximal failures is at least $\frac{n^2}{64}$. ■

[5]No attempt has been made to optimize the constant. In fact, a more elaborate variant of the preceding construction gives $\frac{n^2}{16}$ maximal failures.

B. Circular Disk Failures With Radius at Most r

In this subsection, we take a more general model and assume that the radius of the failure is not a network-wide parameter but depends on the area. Our goal is to enumerate every circular disk failure for any radius at most r.

Definition 4: Let a disk c be ***smaller*** *than disk c', if c has a smaller radius than c', or if they have equal radius and the centre point of c is lexicographically smaller than the centre point of c'.*

Definition 5: Let $F \subseteq E$ be a finite nonempty set of edges (not necessarily a failure). We denote the smallest disk among the disks hitting F by c_F and we say c_F is the ***smallest hitting disk*** *of F.*

It is not difficult to see that c_F always exists. The key idea of our approach that we can limit our focus only on the smallest hitting disks c_F, for $F \in \mathcal{H}_r$, and ignore the rest of the disasters. The consequence of the next theorem is that the number of smallest hitting disks c_F, $F \in \mathcal{H}_r$ is not too large.

Theorem 2: Let H be a nonempty set of intervals from $\mathbb{R}^2$ with smallest covering disk c_H. Then there exists a subset $H' \subset H$ with $|H'| \leq 3$ such that $c_H = c_{H'}$.

Thm. 2 would be trivial if the smallest hitting disks were defined on sets of nodes because a triplet of non-collinear nodes defines a circle. In the proof in Appendix A we show that this property holds for edges (considered as line segments) too. Compared to the algorithm of Thm. 1 here we not only shift the disks but also shrink them.

Corollary 2: $\left|\bigcup_{0<r<\infty} \mathcal{S}_r\right| \leq \binom{m}{3} + \binom{m}{2} + m = \frac{m^3}{6} + \frac{5m}{6}$.

Theorem 3: Let H be a set of intervals from $\mathbb{R}^2$, $|H| \leq 3$. Then c_H can be determined in $O(1)$ time.

The proof is relegated to Appendix B.

Remark. Thm. 3 outlines an efficient algorithm for c_H in an exact symbolic computational setting. A good numerical algorithm for approximating the radius r of c_H and the center P of c_H is also possible: for a positive real number r' we can efficiently test if

$$N(e_1,r') \cap N(e_2,r') \cap N(e_3,r') \neq \emptyset.$$

Indeed $N(e_i,r')$ is a union of two half disks and a rectangle, and the intersection of such objects is easily computable. Using such tests for emptiness, r can be approximated by binary search as the smallest r' providing nonempty intersection.

Since the smallest hitting disk of a triplet of edges can be calculated in $O(1)$ time, we could solve the problem by processing $O(m^3)$ triplets of edges. However, we will achieve better upper bounds on the running time and of $|\mathcal{S}_r|$ with the help of some further observations.

V. Improved Bounds and Algorithm to Enumerate the Set of SRLGs

Next, we define five practical parameters of the input to better estimate the number of SRLGs and computing time.

- ρ_r is the *link density* of the network which is measured as the maximal number of links that could be hit by a circular disk shaped disaster of radius r.
- x is the number of link crossings of the network G.
- x' is the number of link crossings in G', where all the edges are extended by $3\sqrt{2}r$ in both directions.[6]
- μ is the square mean of numbers v_e for all $e \in E$, where v_e is the number of $w \in V \cup X$ such that $d(w, e) \leq 3r$.

In backbone networks, x is a small number as typically a network node is also installed on each link crossings [20], while the link density ρ_r practically should not depend on the network size. We also know that ρ_r is at least the maximal nodal degree in the graph. For simplicity, we assume that edges intersect in at most one point.

Definition 6: Let X be the set of points p which are not in V and there exist at least 2 non-parallel edges crossing each other in p. Let $x = |X|$.

Although on arbitrary graphs x can be even $\Theta(n^4)$, in backbone network topologies typically $x \ll n$. This is because a switch is usually installed if two cables are crossing each other.[7] It gives us the intuition that G is "almost" planar, and thus it has few edges.

Claim 2: The number of edges in G is $\Omega(n)$ and $O(n+x)$.

Proof: Since G is connected, $m = \Omega(n)$ is immediate.

Let $G'(V \cup X, E')$ be the planar graph obtained from dividing the edges of G at the crossings. Since every crossing increases the number of edges by at least two, $|E'| \geq m + 2x$. On the other hand, $|E'| \leq 3(n + x) - 6$ since G' is planar. Thus $m \leq |E'| - 2x \leq 3n + x - 6$. ∎

A. Lower Bound on Computing the Maximal Failures

Now we present a straightforward lower bound on the time needed to determine $\mathcal{S}_r$. As it will turn out (in Cor. 5), in specific circumstances, this lower bound is asymptotically tight.

Proposition 2 Lemma 4 of [35]: Any algorithm that correctly computes segment intersections among any set of m segments will take on the order of $m \log m + k$ time in the worst case, for any value of the input size m and any feasible value of the output size k, where k denotes the number of pairwise intersections of line segments.

Corollary 3: The complexity of computing $\mathcal{S}_r$ is $\Omega(n \log n)$.

Proof: By combining Prop. 2 and Claim 2, we get that reporting that there are no intersecting line segments takes $\Omega(n \log n)$. In other words, this means that computing $\mathcal{S}_r$ in the special case of $r = 0$ needs $\Omega(n \log n)$ time. ∎

[6] A smaller extension (reducing the number of new intersections) would be enough using an algorithm more complicated than in the proof of Thm. 6.

[7] Recent experimental studies give empirical evidence that real-world road networks typically have $\Theta(\sqrt{n})$ edge crossings [34].

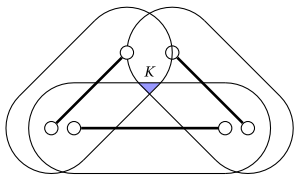

Fig. 5. Illustration to Thm. 4.

B. Upper Bounds and Algorithm for Computing the Maximal Failures

The set of link intersections X can be computed in near-linear time, for example, with the help of algorithm Bentley-Ottmann [36]:

Proposition 3 Theorem 2.4 of [36]: All intersection points of E, together with the segments giving the intersection, can be reported in $O((m+I)\log m)$ time and $O(m)$ space, where I is the number of intersection points.[8]

Claim 3: X can be reported in $O((n + x)\log n)$ time and $O(n + x)$ space.

Proof: To easily distinguish nodes and edge intersections geometrically, edges are shortened in both directions with a tiny fraction of their length. The statement follows by using Proposition 3 and Claim 2 by noting also that $O(\log(n+x)) = O(\log n)$. ∎

The next theorem states, it is enough to process the edge triplets in the neighborhood with radius $3r$ of every point in $V \cup X$.

Theorem 4: For every failure $H \in \mathcal{H}_r$ there exists a disk c of radius at most r hitting H with centre point at distance at most $2r$ from $V \cup X$.

The proof of the theorem is relegated to Appendix C.

Theorem 5: Let r be a positive real number, $F \in \mathcal{S}_r$ be a set of line segments which can be hit by a disk of radius r. Then there exists a segment $e \in F$ and a disk c described in Thm. 1 (disk c has radius r, hits F, intersects e in a single point Q, and (a), or (b), or (c) holds with $H = F$), such that the centre point of c is at distance at most $2r$ from either an endpoint of e or a point of crossing (of e and an other segment $f \in F$).

Proof: We proceed along the lines of the proof of Thm. 1. If we are in case (1) of the proof of Thm. 1, then (b) or (c) holds for the statement of the theorem, as Q can be an endpoint of a segment $e \in F$.

We may turn our attention to the case (2) from Thm. 1. Then $K = \cap_{e \in F} N(e, r)$ is a closed bounded convex set on the plane whose boundary is a polygon composed of line segments. If K has no interior points in the plane, then r is an optimal hitting radius for F. Then $c = c_F$ will be a suitable disk. The proof of Thm. 4 can be extended to show that the requirements of Theorem 1 will be valid for c_F in the place of c. This follows from a simple but tedious analysis of the Cases 1-4 of Theorem 3, which we omit here.

We may, therefore, assume that K has an interior point (see also Fig. 5). Then K is a proper convex k-gon for some $k \geq 3$,

[8] Note that it is easy to modify the algorithm used in proof of Proposition 3 in [36] to determine $\mathcal{S}_r$ for $r = 0$. The case of $\frac{r}{2}$ being larger than the geographical diameter of the network is also trivial. In this article we fill in the gap between these two values.

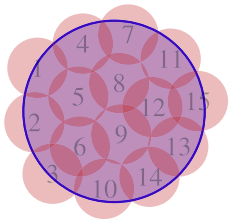

Fig. 6. A disk with radius $3r$ can be hit with 15 disks with radius r.

TABLE I
ALGORITHM FOR DETERMINING $\mathcal{S}_r$ AND COMPLEXITY OF ITS TASKS

#	Task	Complexity
1	Determine X	$O((n+x)\log n)$
2	For $w \in V \cup X$ determine E_w	$O((n+x')\log n + (n+x)\rho_r \log \rho_r)$
3	For $e \in E$ determine V_e	$O((n+x')\log n + (n+x)\rho_r \log \rho_r)$
4	For $w \in V \cup X$ determine $\mathcal{L}_{r,w}$	$O((n+x)\rho_r^3)$
5	For $e \in E$ for $w_1, w_2 \in V_e$ compare $\mathcal{L}_{w_1}$ with $\mathcal{L}_{w_2}$	$O((n+x)\mu\rho_r^5)$
6	Merge resulting lists in $\mathcal{S}_r$	$O(n+x)$

hence there exists a vertex R of K with angle $\alpha \geq \frac{\pi}{3}$. The circle of radius r centered at R will meet the requirements of the theorem. Indeed, there will be then two segments $e, f \in F$ such that their supporting lines are tangent to c, and c is seen at angle α from their point of intersection. Q will be the point of tangency of e or f with c. See the last case in the proof of Thm. 4 for further details. ■

Next, we will give better upper bounds on the number of SRLGs. As a consequence of Theorem 5, when considering circular disasters of radius r, then in a sense, we may ignore the points on the edges $e \in E$ which are more than $3r$ away from $V \cup X$. Consider the pairs (e, v) where $e \in E$, $v \in V \cup X$, and $v \in e$. If we have an SRLG of radius r as in Theorem 5 with edge e such that the distance of c is at most $2r$ from v, then the edges of this SRLG must intersect the disk of radius $3r$ centered at v. This gives at most $15\rho_r$ possibilities for the other edge besides e in Theorem 5 (a) or (b) (see Fig. 6, where 15 circular disks of radius r cover a disk of radius $3r$). The number of pairs (e, v) can be counted by looking at the contribution of node v: it will be $\deg v$, where $\deg$ is the degree in the planarized graph. The sum of the degrees is twice the number of the edges of the latter graph, which is $O(n+x)$. Thus we have the following bound:

Corollary 4: $|\mathcal{S}_r| = O((n+x)\rho_r)$.

This bound is asymptotically tight[9] on the graphs in Claim 1 because $\rho_r = \frac{n}{2}$ for $r = k$. Next, we discuss the algorithm to generate the list of SRLGs.

Theorem 5 together with other formerly presented results inspire an improved algorithm with a running time near linear in n described in Table I. The main idea is to build up local data structures, pre-compute the lists of candidate members of $\mathcal{S}_r$, then merge these lists, all in nearly linear time. With this aim, we make the following definitions.

Definition 7: For a given r and $w \in V \cup X$, let $E_w := \{e \in E |\ d(w,e) \leq 3r\}$; and let the edges in E_w be given in sorted order with respect to the lexicographic ordering of their endpoints. For a given $e \in E$, let $V_e := \{w \in V \cup X |\ d(e,w) \leq 3r\}$.

Definition 8: Let $G'(V, E')$ be the graph resulting from elongating the edges of E by $3\sqrt{2}r$ in both directions. Let X' be the set of edge intersections in G', and let $x' = |X'|$.

Theorem 6: All the sets E_w for $w \in V \cup X$ can be determined in time $O((n+x')\log n + (n+x)\rho_r \log \rho_r)$. Similarly, all the sets V_e for $e \in E$ can be computed in the same time complexity.

The proof of Thm. 6 is relegated to Appendix D.

Lemma 2: The set of SRLGs circular disk shaped disasters of radius r can be computed in $O(((n+x)(\log n + \rho_r^3) + x' \log n))$

Proof: Based on Claim 3 and Thm. 6, E_w can be determined in the proposed complexity for all $w \in V \cup X$.

Then for every node w, we compute list $\mathcal{L}_{r,w}$ containing the set of edges hit by an element of disk set $\mathcal{C}_{r,w}$ defined as follows: for $e, f \in E_w$ we compute disks c of radius r (if exist) according to Thm. 2: either case a) applies if e and f are not parallel, and c intersects them in two different points, or case b) when c intersects e and f in two different points, one being an endpoint of e, or case c) when c touches e at an endpoint; moreover we require that formerly computed disks c have centres not farther than $2r$ from w. These disks are collected in $\mathcal{C}_{r,w}$. This takes $O((n+x)\rho_r^3)$ time, since there are $O(\rho_r^2)$ disks c to determine and store in $\mathcal{C}_{r,w}$, and for each $c \in \mathcal{C}_{r,w}$ the set of edges hit by c can be determined in $O(\rho_r)$ time based on E_w. It follows readily from Thm. 5 that for every $F \in \mathcal{S}_r$ there exists a $w \in V \cup X$ such that F is a subset of an element of list $\mathcal{L}_{r,w}$. ■

Please note that lists $\mathcal{L}_{r,w}$ together may contain duplicates and non-maximal sets as well, those will be eliminated later at a subsequent phase.

Finally, based on Cor. 2 we give an upper bound on the total number of circular disk failures with radius at most r.

Proposition 4: $\left| \bigcup_{0<r'<r} \mathcal{S}_{r'} \right| = O((n+x)\rho_r^2)$.

Proof: We can use Theorems 2 and 4 and the fact that a disk of radius $3r$ hits $O(\rho_r)$ segments. From Theorem 2, we see that it suffices to construct disks of the form c_H, for sets of segments H of size at most 3. Then by Theorem 4 it is enough to calculate for every $v \in V \cup X$ the smallest hitting disk of every set H containing an edge going through v and containing 1 or 2 edges from the $3r$ neighborhood of v. For a fixed v we have $O(\deg v \cdot \rho_r^2)$ SRLGs, and the claim follows. ■

As mentioned after Lemma 2, the final task for determining $\mathcal{S}_r$ is to merge lists $\mathcal{L}_{r,w}$ by eliminating duplicates and non-maximal elements. To do this in subquadratic time in n, one must avoid comparing all pairs of lists $\mathcal{L}_{r,w_1}$, $\mathcal{L}_{r,w_2}$.

Definition 9: Let μ be the mean square of numbers $|V_e|$ for all $e \in E$, i.e. $\mu := \frac{\sum_{e \in E} |V_e|^2}{m}$.

Theorem 7: The maximal circular disk failures with radius exactly r can be computed in time $O((n+x)(\log n + \mu\rho_r^5) + x' \log n)$ and this is tight in n.

[9]No attempt have been made to optimize the constant.

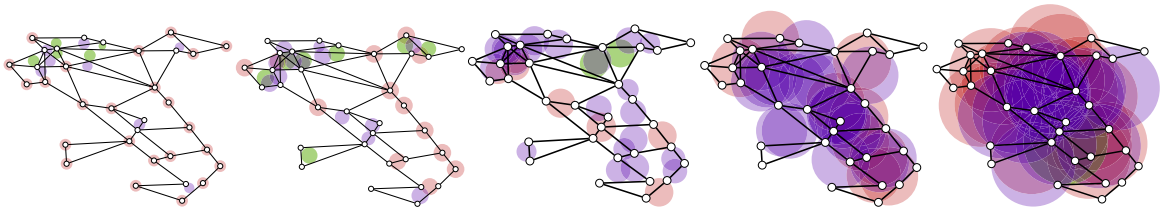

(a) $r = 20$ (b) $r = 30$ (c) $r = 50$ (d) $r = 100$ (e) $r = 150$

r	ρ_r	y_r	$\lvert\mathcal{S}_r\rvert$
50	7	243	37
70	7	301	37
90	7	338	38
110	7	408	38
130	7	481	38
150	8	546	35
170	8	590	34
190	8	635	37
210	8	705	38
230	12	758	35

(f) Table of results

Fig. 7. The set of SRLGs in the 33-node Italian network for various radius sizes. The network has $n = 33$ nodes, $m = 56$ links, and $x = 4$ edge crossings.

Proof: According to Lemma 2, all sets of failures $\mathcal{L}_{r,w}$ can be determined in time $O(((n+x)(\log n+\rho_r^3)+x'\log n))$.

We observe that it is enough to compare lists $\mathcal{L}_{r,w_1}$ and $\mathcal{L}_{r,w_2}$ for possible containment or duplicates only if $E_{w_1} \cap E_{w_2} \neq \emptyset$, or in other words there exists an $e \in E$ for which $\{w_1, w_2\} \subseteq V_e$. We deduce that it is enough to compare for all $e \in E$ and $w_1, w_2 \in V_e$ list pairs $\mathcal{L}_{r,w_1}$, $\mathcal{L}_{r,w_2}$. This means comparing at most

$$\sum_{e\in E} \frac{|V_e|(|V_e|-1)}{2} < m\frac{\sum_{e\in E}|V_e|^2}{m} 0 = m\mu \overset{\text{Claim 2}}{=} O((n+x)\mu)$$

pairs of lists, with each list having $O(\rho_r^2)$ elements. Taking into consideration that a comparison of two elements (SRLG candidates) can be done in $O(\rho_r)$, we obtain a complexity of $O((n+x)\mu\rho_r^5)$, confirming claim for the total complexity. The lower bound is provided by Corollary 3. ∎

Table I summarizes the steps of our proposed algorithm. Note that parameters ρ_r, x, x' and μ are theoretically upper bounded by m, $\frac{m(m-1)}{2}$, $\frac{m(m-1)}{2}$ and $(n+x)^2$, respectively, meaning that our algorithm for determining $\mathcal{S}_r$ is clearly polynomial in n or m. Furthermore, based on Thm. 7 using that x is $O(n)$ in practice, and that ρ_r is more or less proportional to $\frac{2r}{diam}m$ (see Sec. VI) in the interval $(0, diam/2]$, where *diam* is the geometric diameter of the network, we get a complexity bound of $O\left(n(\log n + \mu(\frac{r}{diam})^5) + x'\log n\right)$ for determining $\mathcal{S}_r$. Also, as in practice $x = O(n)$, and for r much smaller than network diameter, $\rho_r = O(1)$, $\mu = \log(n)$, and $x' = O(n)$, we can state the following corollary:

Corollary 5: If $\rho_r = O(1)$, $\mu = O(\log n)$, and x and x' are $O(n)$, $\mathcal{S}_r$ can be calculated in $O(n \log n)$ optimal time. These assumptions hold in practice when r is much smaller than the geographical network diameter.

Proof: Combining Thm. 7 and Cor. 3 yields the proof. ∎

VI. Numerical Results

In this section, we present numerical results that demonstrate the use of the proposed algorithms on some real backbone networks. The algorithm was implemented in C++ using the Geometric Tools Engine, a library for computing in the fields of mathematics, graphics, and image analysis (Wild Magic 5 distribution, version 5.13). The output of the algorithm is a list of SRLGs so that no SRLG contains the other. The network topologies with the obtained list of SRLGs for various radii are available online.[10]

We visualize each SRLG in the obtained list of SLRGs $\mathcal{S}$ by its smallest hitting disk. According to Thm. 3 the smallest hitting disk is computed using at most 3 nodes or edges. The different cases are shown with different colors: The red circles go through 2 or 3 nodes, and the disks hitting 1 node are represented as red disks with radius r and the center being the given node. The green disks have 3 edges on the boundary. All other disks are violet.

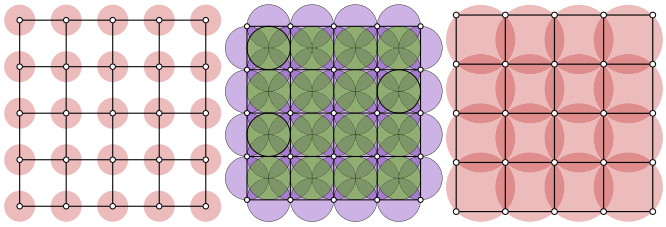

(a) $r = 0.3$, $|\mathcal{S}_r| = 25$ (b) $r = 0.5$, $|\mathcal{S}_r| = 56$ (c) $r = \sqrt{2}/2$, $|\mathcal{S}_r| = 16$

Fig. 8. The set of SRLGs of a 5×5 grid network.

Fig. 7 shows the Italian optical backbone network with circular disk shaped disasters of three different radii $r = 20, 30, 50, 100, 150$km. For the smallest radius, the SRLGs are the nodes and the edge crossing points. In our bounds, the number of SRLGs was $O((n+x)\rho_r)$, and ρ_r increases with the radius. Surprisingly, as the radius increases, the number of SRLGs does not increase but stays close to $n + x$. It is because the SRLGs, which are a subset of another SRLGs are filtered out. Note that our bounds are asymptotically tight for the artificial networks in Fig. 4. In other words, it seems the number of SRLGs does not depend on the radius in practical scenarios.

To understand this phenomenon let us consider a perfect 2D grid network of $k \times k$ nodes, where the length of each edge is 1. Until the radius is less than $\frac{1}{2}$ only node failures must be considered, as shown on Fig. 8a. The total number of such failures is $|\mathcal{S}_r| = k \times k$. As the radius increases reaching $\frac{1}{2} \leq r < \frac{\sqrt{2}}{2}$ we have the SRLGs of every link with the neighbouring links, and every facet (each square) of links, that is $|\mathcal{S}_r| = 3(k-1) \times k - (k-1)$ in total (see Fig. 8b). As the radius further increases to $r = \frac{\sqrt{2}}{2}$ the SRLGs will be every facets with the neighbouring links, which is $|\mathcal{S}_r| = (k-1) \times (k-1)$. Finally, when $r = \frac{\sqrt{2}}{2}k$ we have only one SRLG hitting every link of the network, i.e. $|\mathcal{S}_r| = 1$. This example illustrates that $|\mathcal{S}_r|$ does not increase or decrease monotonously, and may have local maxima and minima.

To analyze the relationship between the radius and the number of SRLGs, we analyzed 6 real-world backbone networks: 3 European and 3 US topologies. We have plotted the length of the SRLG list $|\mathcal{S}_r|$ compared to the radius of the circular disaster. Fig. 9 shows our results where ρ_r the maximal

[10] https://github.com/jtapolcai/regional-srlg

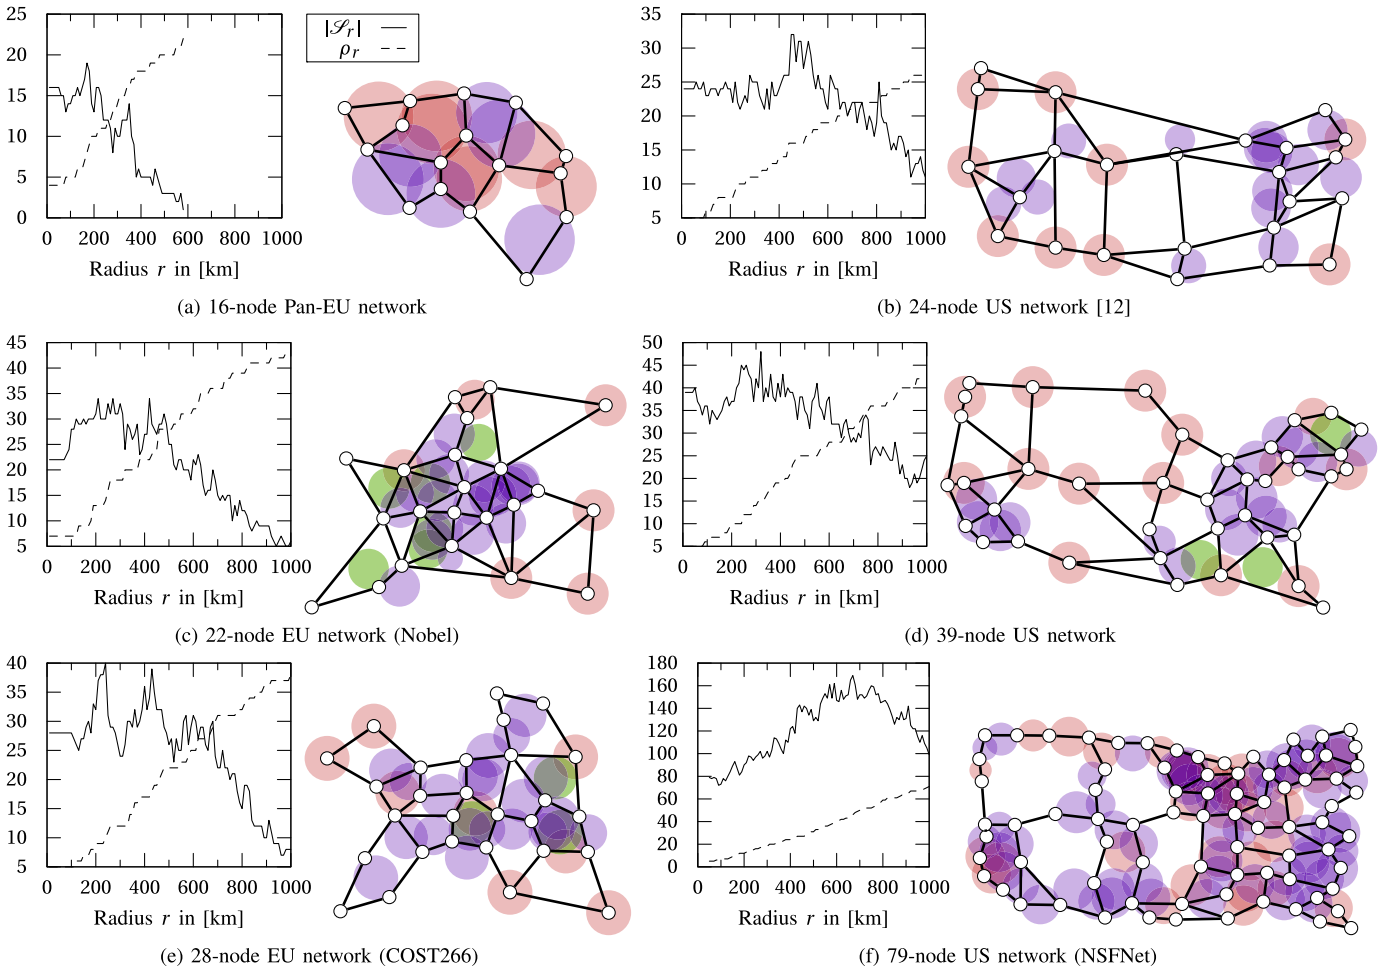

Fig. 9. The number of SRLGs $|\mathcal{S}_r|$ vs. the radius r of the disaster. The number of edges in the largest SRLG, ρ_r, is also plotted. The graph topologies with the SRLGs of radius 200km is plotted next to the charts.

TABLE II
RESULTS ON SOME BACKBONE TOPOLOGIES FROM [37]

Network	n	m	x	ρ_r	y_r	$\lvert\mathcal{S}_r\rvert$	Runtime [s]
Pan-EU	16	22	0	5	44	14	0.131
German	17	26	0	7	69	15	0.28
EU	22	45	0	13	176	34	9.569
US [12]	24	42	0	8	124	24	1.233
US	26	42	0	10	122	25	5.668
EU (Nobel)	28	41	0	9	94	39	3.983
Italian	33	56	4	13	199	31	14.17
EU (COST266)	37	57	0	7	134	41	0.537
US	39	61	0	7	152	33	0.83
US (NFSNet)	79	108	0	9	217	92	7.102

number of links in the SRLG is also listed. ρ_r increases with the radius; however, the number of SRLGs slightly increases until 100-200 km radius, and after that, it has a flat period with local maxima and minima, and finally, it decreases as the radius becomes extremely large. Surprisingly the number of SRLGs was never more than $2.3n$ for any radius, and often it is less than the number of links. We also plotted the SRLGs for disasters of radius 200km. Note that, in the European networks, the nodes are closer to each other compared to the US. The SRLGs are mostly node failures and in the densely connected areas small sets of links. The list of SRLGs obtained with our approach for the 24-node US network hits the disaster zones for earthquakes, tornadoes, and weapons of mass destruction attacks defined in [12].

Table II shows a comparison among the networks, where the radius r is the length of the shortest edge in every network. The columns are: network name, the number of nodes and links and link crossings, the two link density metrics ρ_r, y_r the total number of edge pairs whose distance is at most $2r$, the number of SRLGs, and running time. The runtime corresponds to the slower algorithm, which enumerates every circular disk failure with radius at most r. It was measured on a commodity laptop with Core i5 CPU at 1.8 GHz with 4 GB of RAM.

VII. CONCLUSION

In this article, we view networks as geometric graphs and regional disasters (natural on human-made) as circular disks of a given radius and propose a fast and systematic approach to enumerate the list of possible link failures caused by the disasters. Although the number of possible disasters is infinite, we show that under reasonable and realistic assumptions, the list of failures to be considered is short, it is close to linear in the network size. Note that, we do not assume the failed

region has a disk shape, but overestimate the size of the failed region. Intuitively, the radius defines the minimum distance between the working and backup resources. We present two fast polynomial-time algorithms, the first lists every link set that can be hit by a circular disk shaped disaster of a fixed radius. Its central idea is to move a disk of radius r to every candidate location. The second algorithm lists every disk failure with a radius at most r. This allows a more sophisticated regional failure model where different radii of disasters are used in flat or hilly areas. It also helps in understanding the number of SRLGs compared to the network size. The algorithm moves and shrinks the disks at every candidate location. Through numerical evaluation of several specific networks, we show that the algorithms are fast enough for network design problems, and the obtained list of SRLGs is surprisingly small $\approx 1.2n$.

APPENDIX

A. Proof of Thm. 2

We need the following simple lemma.

Lemma 3: Let C_1, C_2, C_3 be convex subsets of the plane $\mathbb{R}^2$ such that $t = C_1 \cap C_2 \cap C_3$ is a line segment with more than 1 point. Then there exist two indices i, j such that $C_i \cap C_j$ is collinear.

Proof: Let R, S be two different points of t. If the statement is false, then the pairwise intersections $C_i \cap C_j$ all contain points not on the line of t. Without loss of generality we may assume that $C_1 \cap C_2$ and $C_1 \cap C_3$ contain points P_2 and P_3 from the same open halfplane defined by the line of t. If $P_3 = P_2$ then we obtain $P_2 \in t$ which is a contradiction. We infer that P_3, P_2, R, S are four different points. Radon's lemma (Theorem 1.3.1 in [38]) can be applied to them. The Radon point X will be on one hand on the open halfplane defined by the line of t and containing the P_i. On the other hand

$$X \in C_1 \cap C_2 \cap C_3.$$

This gives a contradiction. ∎

Proof: [Proof of Thm. 2] Let r be the radius of c_H. We have then

$$\cap_{e \in H} N(e, r) \neq \emptyset, \tag{3}$$

but $\cap_{e \in H} N(e, r') = \emptyset$ for any $r' < r$.

The statement of the theorem is immediate if H has at most 2 sets. Suppose now that $|H| \geq 3$. If for every 3-element subset H' of H there exists a radius $r_{H'} < r$ such that $\cap_{e \in H'} N(e, r_{H'}) \neq \emptyset$, then with $r^* = \max r_{H'}$ we have $\cap_{e \in H} N(e, r^*) \neq \emptyset$ by the planar Helly's theorem (Theorem 1.3.2 in [38]) applied to the convex sets $N(e, r^*)$, hence H can be hit by a disk of radius r^*, which is impossible. We obtain that there exists a 3 element subset H' of H such that the radius of $c_{H'}$ is r.

Note also, that the intersection on the left of (3) is necessarily a (possibly degenerate) nonempty closed bounded line segment s. This follows from the fact that the intersection is a nonempty closed bounded convex subset *without an interior point*. Indeed an interior point would allow a hitting radius for H, which is less than r. Note also that the lexicographically smallest (end)point P of s is the center of c_H.

We observe next that for H' above the hitting radius r is also minimal, hence the intersection $\cap_{e \in H'} N(e, r) = s'$ is also a line segment which contains s. If the smallest point of s' is P then we are done, as P will be the center of $c_{H'}$. We may therefore suppose that s' contains a point Q smaller than P.

Suppose that $H' = \{e_1, e_2, e_3\}$. We verify that there exist i, j, $1 \leq i < j \leq 3$, such that the intersection $N(e_i, r) \cap N(e_j, r)$ is a subset of the line of s'. Indeed, this follows from Lemma 3 applied to the neighborhoods $N(e_i, r)$ and $t = s'$.

We conclude by noting that there exists an edge $f \in H$ such that $N(f, r)$ does not have a point on the line of s' which is smaller than P (otherwise s itself had such a point). These imply that the lexicographically smallest point of $N(e_i, r) \cap N(e_j, r) \cap N(f, r)$ is P and the proof is complete. ∎

B. Proof of Thm. 3

Let H be a set of intervals from $\mathbb{R}^2$, $|H| \leq 3$. We show that c_H can be determined in $O(1)$ time.

Again, Thm. 3 would be trivial in case of nodes instead of edges. *Proof:* [Proof of Thm. 3] We assume that the intervals $e \in H$ are given as input by their endpoints whose coordinates are real numbers. We allow degenerate segments consisting of just one point, too. In the complexity count, we use the unit cost arithmetic model of computation, where the cost of a basic operation with real numbers is 1. The basic operations allowed here are $+, -, *, /$, and taking the nonnegative square root of a nonnegative real number.

Case 1. *The intervals of H have a point in common.* Then c_H is the lexicographically smallest point of $\cap_{e \in H} e$. This point is easy to compute in $O(1)$ time.

From here on we have $|H| > 1$, and $\cap_{e \in H} e = \emptyset$. Henceforth let r stand for the radius of c_H.

Fact. *As before, the intersection $s = \cap_{e \in H} N(e, r)$ is a collinear set. It is either a single point P, or a nondegenerate line segment PQ with P, Q different points of the plane.* Indeed, otherwise s would have an interior point. That point would be closer than r to every $e \in H$.

Case 2. $|H| = 2$.

Subcase 2a. $|H| = 2$, $H = \{e_1, e_2\}$, $s = PQ$. Then s is on the boundary of both $N(e_i, r)$, e_1, e_2 are parallel segments, $2r$ is exactly the distance of their respective lines, and P, Q is easily computed. The center of c_H will be the lex-smaller of P and Q (see also Fig. 10a).

Subcase 2b. $|H| = 2$, $H = \{e_1, e_2\}$, *s is a single point* P. This is possible only if $N(e_1, r)$ and $N(e_2.r)$ touch each other. An endpoint P_1 of e_1 is at distance $2r$ from e_2, or an endpoint of e_2 is at distance $2r$ from e_1. In the former subcase a candidate disk for c_H will be the disk whose diameter is the segment connecting P_1 to the unique closest pont X of P_1 to e_2. Similarly we have to consider the smallest circles passing through an endpoint of e_2 and intersecting e_1 in a single point (see also Fig. 10b). These circles can be calculated efficiently.

Note that we do not necessarily know in advance which subcase we are dealing with, and hence may need to calculate as many as 5 circles and select the smallest among them.

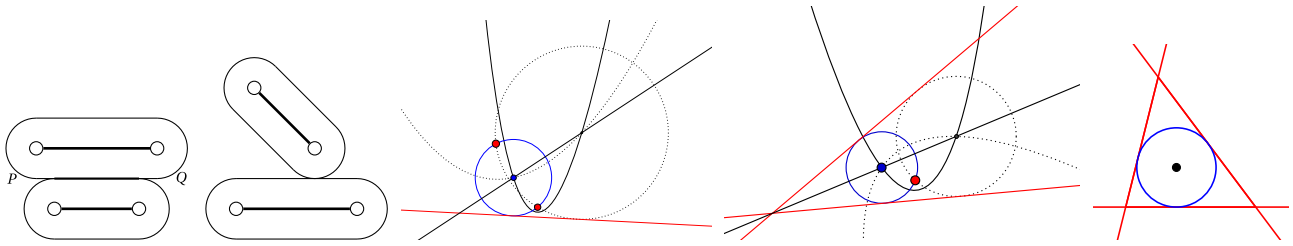

(a) Case 2a (b) Case 2b (c) Smallest hitting disk of 2 points and 1 line (d) Smallest hitting disk of 1 point and 2 lines (e) Smallest hitting disk of 3 lines

Fig. 10. Illustration of the cases of Thm. 3.

Case 3. *$|H| = 3$, $H = \{e_1, e_2, e_3\}$, and there exist two different indices i, j such that $N(e_i, r) \cap N(e_j, r)$ is a collinear pointset.*

Subcase 3a. *$|H| = 3$, $H = \{e, e_2, e_3\}$, and there exist different indices i, j such that $N(e_i, r) \cap N(e_j, r) = s'$ is nondegenerate line segment.* Then e_i and e_j are necessarily parallel segments. We calculate s' and r as in Subcase 2a. The center of our candidate for c_H will be the smallest point R on s' which is at distance at most r to e_k (if there is any such point), where k is the index from $\{1, 2, 3\}$ different form i and j. Note that all the segments may be parallel. Then we may have to work with more than one pair e_i, e_j, and take the smallest of the resulting circles for c_H.

Subcase 3b. *$|H| = 3$, $H = \{e_1, e_2, e_3\}$, and there exist different indices i, j such that $P = N(e_i, r) \cap N(e_j, r)$ is a single point P.* Then necessarily $N(e_1, r)$ and $N(e_2.r)$ touch each other and we may proceed as in Subcase 2b. We have to test if the resulting candidate for c_H does indeed itersect the third segment.

Case 4. *$|H| = 3$, $H = \{e_1, e_2, e_3\}$, and for every pair of different indices i, j the intersection $N(e_i, r) \cap N(e_j, r)$ is not a collinear set.*

Here by Lemma 3 we know, that the intersection

$$N(e_1, r) \cap N(e_2, r) \cap N(e_3, r)$$

is a single point P. Moreover P is on the boundary of every $N(e_i, r)$ as otherwise we would be in Case 3. These imply that for every i segment e_i has exactly one point $P_i \in e_i$ at distance r form P (similar to Fig. 5 but K is a single point).

First we dispense with the subcase when there exist an $i \neq j$ such that $P_i \in e_j$. To account for these possibilities, for each pair $i \neq j$ we determine the smallest disk hitting $e_i \cap e_j$ (if it is not empty) and e_k for $k \neq i, j$. This can be done as in Case 2, and put the resulting disks onto the list of candidates for c_H.

From that point on, we may assume that the points P_i are not included on e_j whenever i is different from j. In particular, they are not collinear.

Subcase 4a. *P_i is an endpoint of e_i for every i.* Then put on the list of candidates for c_H the unique circle passing through these 3 points. This circle is easily computed. We do that for every possible selection of the endpoints $P_i \in e_i$ which give three different points (there may be as many as 8 possibilities).

Subcase 4b. *Two of the P_i are endpoints of the e_i the third of them (say P_k) is an inner point of e_k.* Then the circle of radius r with center P touches the line of e_k. The candidate for P with these data is uniquely determined as follows: we compute (a quadratic equation of) the parabola which is the locus of the points equidistant to P_i and e_k (whose focus is P_i and directrix is the line of e_k). By our assumptions, this is a nondegenerate parabola. We intersect the parabola with the perpendicular bisector of the segment P_iP_j, where $j \neq i, k$. This gives at most two possible points for P. We take the smallest of the resulting circles as a candiadte for c_H. We do this for every pair i, j, $i \neq j$. This gives at most 6 candidates for c_H (see also Fig. 10c).

Subcase 4c. *One of the P_i is an endpoint of e_i, the other two are inner points of, say e_j and e_k.* In this case the lines ℓ_j of e_j and ℓ_k of e_k can not be parallel, because then we would be in Subcase 3a. Point P will be at the intersection of the nondegenerate parabola with focus P_i and directrix ℓ_j, and the angular bisectors of the intersecting lines ℓ_j and ℓ_k. This gives at most four possibilities for P. These points and the corresponding candidates for c_H are easily calculated. We have to do this computation for every $i \in \{1, 2, 3\}$, giving at most 12 candidates for c_H (see also Fig. 10d).

Subcase 4d. *The points P_i are all inner points of their segments e_i.* Then, as in the previous subcase, we infer that no two of the lines ℓ_i determined by e_i can be parallel. In this setting, as a noted special case of Apollonius' problem (see pages 346-355 in [39]) four circles are touching ℓ_1, ℓ_2, ℓ_3 and they can be calculated easily by taking intersections of angular bisectors for the line pairs ℓ_i, ℓ_j. We note however that only the inscribed circle for the triangle determined by the by ℓ_1, ℓ_2, ℓ_3 can be a candidate for c_H. the other three disks cannot be optimal[11] (see also Fig. 10e).

The subcases we have given cover all the possibilities. When we process the input in the order of the subcases discussed above, we do not necessarily know in advance whether the subcase just considered is a valid one for that particular instance. One only computes the candidates for c_H described at the subcase, if the input data permits it.[12] There is always just a constant number of circles to form. The minimal one can

[11]The shortest arc of the other circles containing the tree points of intersection with the ℓ_i is less than half of the complete circle, hence the points can be hit by a smaller disk.

[12]For example, if some of the lines ℓ_i are parallel, then we do not compute circles at Subcase 4d, as we do not have a proper triangle to work with. Finally, 4d does not apply to the actual input.

be selected at the end from the list of candidates calculated in the course of the computation. ■

C. Proof of Thm. 4

By Thm. 2 it suffices to show that the statement holds for a subset $H \subset \mathcal{H}_r$ and $c = c_H$ whenever $|H| \leq 3$. We may proceed along the Cases in the proof of Thm. 3.

Proof: [Proof of Thm. 4] In Case 1, the center of c will be a point from X. In Case 2, the circle c passes through a point from V (an endpoint of an interval from H). In Subcase 3a, the center I of c will be on the segment s'. If I is an endpoint of s', then c passes through an endpoint of e_i or e_j. If I is an interior point of s', then I must lie on the boundary of $N(e_k, r')$ where r' is the radius of c, hence c passes through exactly one point of e_k. Now, if c intersects e_k in one of its endpoints, then we are done again. If not, then e_k cannot be parallel to e_i and e_j, because otherwise, we would have a hitting circle with center smaller than I. We have that c touches all of the three intervals in three different interior points. Moreover, exactly two of the intervals are parallel. We defer the analysis of this situation until the end, where the closely analogous 4d is considered.

In Subcase 3b c passes through an endpoint of e_i or e_j.

We now turn to Case 4. If the points $P_i \in e_i$, $(i = 1, 2, 3)$ are not all different, then the resulting candidates for c_H all pass through either an endpoint of an e_i or a point of an intersection $e_i \cap e_j$,[13] and we are done in these cases.

We can henceforth assume that the points P_1, P_2, P_3 are all different. Now Subcases 4a, 4b and 4c are done, as c must pass through an endpoint of an e_i. We turn our attention to Subcase 4d. Then the lines $\bar{e}_1, \bar{e}_2, \bar{e}_3$ of the respective intervals e_1, e_2, e_3 determine a nondegenerate triangle ABC, and c will necessarily be the inscribed circle of this tringle (see Fig. 11), as otherwise c would not be minimal among the circles passing through P_1, P_2, P_3. We assume w.l.o.g. that the largest angle of the triangle ABC is $\widehat{BAC}$. Then the radian measure of the angle $\widehat{BAC}$ is in interval $[\frac{\pi}{3}, \pi)$. This way $\sin(\widehat{IAC}) \geq \frac{1}{2}$ since AI is the angular bisector of $\widehat{BAC}$. Moreover,

$$|AI| = \frac{|LI|}{\sin \widehat{IAC}} \leq \frac{r}{\frac{1}{2}} = 2r.$$

If $A \in V \cup X$, then we are done. Else there exists a node $T \in V$ in $[AK] \cup [AL]$. Clearly, $d(TI) \leq 2r$, and we are done.

Finally, we consider the remaining case from 3a. This is similar to the preceding one, but as two of the three lines are parallel, ABC has a point at infinity. At one of the vertices, say A, the triangle will have an angle at least $\pi/4$, and we exhibit a suitable point $T \in X \cup V$ with $d(T, I) \leq 2r$ as in the preceding paragraph. ■

D. Proof of Thm. 6

We need the following two simple lemmas.

Lemma 4: Let $A = (x, y)$ be a point in the plane of distance at most 3 from the origin. Then, any line going through A intersects either the x- or the y-axis not farther than $3\sqrt{2}$ from the origin.

[13] If e_i and e_j are from the same line, then only the endpoints of $e_i \cap e_j$ are to be considered in X.

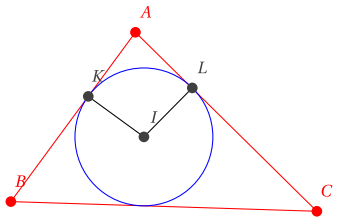

Fig. 11. Illustration for the proof of Thm. 4.

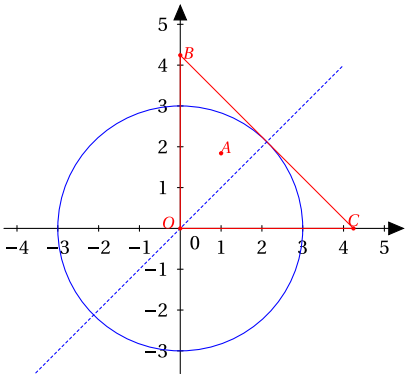

Fig. 12. Illustration for proof of Lemma 4.

Proof: Without loss of generality, we can assume A is in the first quadrant of the plane. Let $B = (0, 3\sqrt{2})$ and $C = (3\sqrt{2}, 0)$, respectively. Then line of BC is tangent to the circle centered at the origin O and having radius 3 (at the point $\left(\frac{3}{\sqrt{2}}, \frac{3}{\sqrt{2}}\right)$, see Fig. 12). Now any line ℓ passing through A must intersect a side of the triangle OBC, hence it intersects at least two sides (Pasch's axiom), therefore ℓ intersects either OB or OC. ■

Definition 10: Let ρ' be the maximum number of edges of E' intersecting a disk with radius $3r$.

Lemma 5: ρ' is $O(\rho_r)$.

Proof: For any point p the number of edges of G' hit by the disk with radius r and center point p is less or equal to the number of edges of G hit by the disk with radius $(1 + 3\sqrt{2})r$ and center point p, which is $O(\rho_r)$ since a disk with radius $(1 + 3\sqrt{2})r$ clearly can be covered by a constant number of disks with radius r. ■

Proof: First, let us concentrate on determining sets E_w for $w \in V \cup X$. We shall use a fast algorithm for line segment intersections (e.g., the Bentley-Ottmann described in [36]) for slightly modified versions of G'.

The most important observation is that, based on Lemma 4, if an edge $e \in E$ is also part of E_w for a $w = (x, y) \in V \cup X$, then the corresponding edge e_{3r} in E' (extended in length by $3\sqrt{2}r$ in both directions) intersects either $I_w^{|} := [(x - 3\sqrt{2}r, y), (x + 3\sqrt{2}r, y)]$ or $I_w^{-} := [(x, y - 3\sqrt{2}r), (x, y + 3\sqrt{2}r)]$. Here we use also the simple fact that the diameter of a square (the length of the longest segment within the square) of side length $3r$ is $3\sqrt{2}r$.

Let $G'^{|}$ be the graph resulting by adding intervals $I_w^{|}$ to G' for every $w \in V \cup X$ as edges of the graph. Let $E_w'^{|}$ denote

the set of edges (of E') intersecting $I_w^{|}$. $G_w'^{-}$ and $E_w'^{-}$ can be defined similarly. It is easy to see that $E_w'^{|} \cup E_w'^{-}$ contains all the edges, which in the original graph G are not farther from w than $3r$, however, it may contain some outliers. Thus in order to get E_w, one can check the distance of the original (i.e., not extended) edges from w, which correspond to edges in $E_w'^{|} \cup E_w'^{-}$ from w.

It is easy to see that $G'^{|}$ has still $O(n+x)$ edges, and $O(n+x')$ edge intersections (with probability 1, if the 'vertical' direction is chosen randomly), thus the intersections of $G'^{|}$ can be determined in $O((n+x')\log n)$ time alongside with the sets $E_w'^{|}$ for $w \in V \cup X$. The same reasoning applies to the sets $E_w'^{-}$.

For any given $w \in V \cup X$, $E_w'^{|} \cup E_w'^{-}$ contains $\leq 2\rho'$ edges, this way based on Lemma 5, E_w can be determined in $O(\rho_r \log \rho_r)$ time in such a way that the edges are given in E_w in sorted order with respect to the lexicographic ordering of their endpoints. This means a total complexity of $O((n+x)\rho_r \log \rho_r)$ for this second phase.

The inverse mapping, i.e., sets of nodes V_e for $e \in E$, can be done in the course (or after) the preceding algorithm. Let V_e be initialized as empty set for all edges e, then, when an E_w is confirmed, w is added to sets V_e for all $e \in E_w$. Clearly, this also can be done in the proposed complexity. ■

ACKNOWLEDGMENT

The authors would like to express our sincere gratitude to Erika Bérczi-Kovács and Attila Kőrösi for their work in the early stage of this research work.

REFERENCES

[1] J. Tapolcai, L. Ronyai, B. Vass, and L. Gyimothi, "List of shared risk link groups representing regional failures with limited size," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Atlanta, GA, USA, May 2017, pp. 1–9.

[2] L. F. Mollenauer, J. P. Gordon, P. V. Mamyshev, I. Kaminow, and T. Koch, *Optical Fiber Telecommunications IIIA*. San Diego, CA, USA: Academic, 1997.

[3] J. Strand, A. L. Chiu, and R. Tkach, "Issues for routing in the optical layer," *IEEE Commun. Mag.*, vol. 39, no. 2, pp. 81–87, Feb. 2001.

[4] D. M. Masi, E. E. Smith, and M. J. Fischer, "Understanding and mitigating catastrophic disruption and attack," *Sigma J.*, pp. 16–22, Sep. 2010.

[5] Y. Nemoto and K. Hamaguchi, "Resilient ICT research based on lessons learned from the great east Japan earthquake," *IEEE Commun. Mag.*, vol. 52, no. 3, pp. 38–43, Mar. 2014.

[6] A. Kwasinski, W. W. Weaver, P. L. Chapman, and P. T. Krein, "Telecommunications power plant damage assessment for hurricane Katrina–site survey and follow-up results," *IEEE Syst. J.*, vol. 3, no. 3, pp. 277–287, Sep. 2009.

[7] Y. Ran, "Considerations and suggestions on improvement of communication network disaster countermeasures after the Wenchuan earthquake," *IEEE Commun. Mag.*, vol. 49, no. 1, pp. 44–47, Jan. 2011.

[8] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of the fiber infrastructure to disasters," *IEEE/ACM Trans. Netw.*, vol. 19, no. 6, pp. 1610–1623, Dec. 2011.

[9] O. Gerstel, M. Jinno, A. Lord, and S. J. Yoo, "Elastic optical networking: A new dawn for the optical layer?" *IEEE Commun. Mag.*, vol. 50, no. 2, pp. s12–s20, Feb. 2012.

[10] M. F. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of disaster-resilient optical datacenter networks," *J. Lightw. Technol.*, vol. 30, no. 16, pp. 2563–2573, Aug. 15, 2012.

[11] J. Heidemann, L. Quan, and Y. Pradkin, *A Preliminary Analysis of Network Outages During Hurricane Sandy*. Los Angeles, CA, USA: Univ. Southern California, Information Sciences Institute, 2012.

[12] F. Dikbiyik, M. Tornatore, and B. Mukherjee, "Minimizing the risk from disaster failures in optical backbone networks," *J. Lightw. Technol.*, vol. 32, no. 18, pp. 3175–3183, Sep. 15, 2014.

[13] I. B. B. Harter, D. A. Schupke, M. Hoffmann, and G. Carle, "Network virtualization for disaster resilience of cloud services," *IEEE Commun. Mag.*, vol. 52, no. 12, pp. 88–95, Dec. 2014.

[14] X. Long, D. Tipper, and T. Gomes, "Measuring the survivability of networks to geographic correlated failures," *Opt. Switching Netw.*, vol. 14, pp. 117–133, Aug. 2014.

[15] B. Mukherjee, M. F. Habib, and F. Dikbiyik, "Network adaptability from disaster disruptions and cascading failures," *IEEE Commun. Mag.*, vol. 52, no. 5, pp. 230–238, May 2014.

[16] R. S. Couto, S. Secci, M. M. Campista, and L. M. K. Costa, "Network design requirements for disaster resilience in IaaS clouds," *IEEE Commun. Mag.*, vol. 52, no. 10, pp. 52–58, Oct. 2014.

[17] G. O'Reilly, A. Jrad, R. Nagarajan, T. Brown, and S. Conrad, "Critical infrastructure analysis of telecom for natural disasters," in *Proc. Netw. 12th Int. Telecommun. Netw. Strategy Planning Symp.*, Nov. 2006, pp. 1–6.

[18] J. S. Foster, Jr., *et al.*, "Report of the commission to assess the threat to the united states from electromagnetic pulse (EMP) attack: Critical national infrastructures," DTIC Document, Fort Belvoir, VA, USA, Tech. Rep., 2008.

[19] R. Wilhelm and C. Buckridge. (2008). *Mediterranean Fibre Cable Cut–a Ripe NCC Analysis*. [Online]. Available: https://www.ripe.net/data-tools/projects/archive/mediterranean-fibre-cable-cut

[20] D. Eppstein, M. T. Goodrich, and D. Strash, "Linear-time algorithms for geometric graphs with sublinearly many edge crossings," *SIAM J. Comput.*, vol. 39, no. 8, pp. 3814–3829, Jan. 2010.

[21] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman, "Power grid vulnerability to geographically correlated failures—Analysis and control implications," in *Proc. IEEE INFOCOM*, Apr./May 2014, pp. 2634–2642.

[22] P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, "The resilience of WDM networks to probabilistic geographical failures," *IEEE/ACM Trans. Netw.*, vol. 21, no. 5, pp. 1525–1538, Oct. 2013.

[23] M. T. Gardner and C. Beard, "Evaluating geographic vulnerabilities in networks," in *Proc. IEEE Int. Workshop Tech. Committee Commun. Qual. Rel. (CQR)*, May 2011, pp. 1–6.

[24] S. Trajanovski, F. A. Kuipers, and P. Van Mieghem, "Finding critical regions in a network," in *Proc. IEEE Conf. Comput. Commun. Workshops (INFOCOM WKSHPS)*, Apr. 2013, pp. 223–228.

[25] B. Vass, J. Tapolcai, D. Hay, J. Oostenbrink, and F. A. Kuipers, "How to model and enumerate geographically correlated failure events in communication networks," in *Guide to Disaster-Resilient Communication Networks*. Cham, Switzerland: Springer, Jul. 2020.

[26] F. Iqbal, S. Trajanovski, and F. Kuipers, "Detection of spatially-close fiber segments in optical networks," in *Proc. 12th Int. Conf. Design Reliable Commun. Netw. (DRCN)*, Mar. 2016, pp. 95–102.

[27] P. N. Tran and H. Saito, "Enhancing physical network robustness against earthquake disasters with additional links," *J. Lightw. Technol.*, vol. 34, no. 22, pp. 5226–5238, Nov. 15, 2016.

[28] P. N. Tran and H. Saito, "Geographical route design of physical networks using earthquake risk information," *IEEE Commun. Mag.*, vol. 54, no. 7, pp. 131–137, Jul. 2016.

[29] M. T. Gardner, R. May, C. Beard, and D. Medhi, "Finding geographic vulnerabilities in multilayer networks using reduced network state enumeration," in *Proc. 11th Int. Conf. Design Reliable Commun. Netw. (DRCN)*, Mar. 2015, pp. 49–56.

[30] J.-C. Bermond, D. Coudert, G. D'Angelo, and F. Z. Moataz, "SRLG-diverse routing with the star property," in *Proc. 9th Int. Conf. Design Reliable Commun. Netw. (DRCN)*, Mar. 2013, pp. 163–170.

[31] N. M. Nam, T. A. Nguyen, and J. Salinas, "Applications of convex analysis to the smallest intersecting ball problem," *J. Convex Anal.*, vol. 19, no. 2, pp. 497–518, 2012.

[32] S. Jadhav, A. Mukhopadhyay, and B. Bhattacharya, "An optimal algorithm for the intersection radius of a set of convex polygons," *J. Algorithms*, vol. 20, no. 2, pp. 244–267, Mar. 1996.

[33] J. J. Sylvester, "A question in the geometry of situation," *Quart. J. Pure Appl. Math.*, vol. 1, no. 1, pp. 79–80, 1857.

[34] D. Eppstein and M. T. Goodrich, "Studying (non-planar) road networks through an algorithmic lens," in *Proc. 16th ACM SIGSPATIAL Int. Conf. Adv. Geographic Inf. Syst. (GIS)*, 2008, p. 16.

[35] B. Chazelle and H. Edelsbrunner, "An optimal algorithm for intersecting line segments in the plane," *J. ACM*, vol. 39, no. 1, pp. 1–54, Jan. 1992.

[36] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars, *Computational Geometry: Algorithms and Applications*. Berlin, Germany: Springer, 2008.
[37] S. Orlowski, R. Wessäly, M. Pióro, and A. Tomaszewski, "SNDlib 1.0—Survivable network design library," *Netw., Int. J.*, vol. 55, no. 3, pp. 276–286, 2010.
[38] J. Matoušek, *Lectures Discrete Geometry*, vol. 108. New York, NY, USA: Springer, 2002.
[39] R. Hartshorne, *Geometry: Euclid Beyond*. Berlin, Germany: Springer, 2013.

János Tapolcai (Senior Member, IEEE) received the M.Sc. degree in technical informatics and the Ph.D. degree in computer science from the Budapest University of Technology and Economics (BME), Budapest, in 2000 and 2005, respectively, and the D.Sc. degree in engineering science from the Hungarian Academy of Sciences (MTA) in 2013. He is currently a Full Professor with the High-Speed Networks Laboratory, Department of Telecommunications and Media Informatics, BME. He has authored more than 150 scientific publications. He was a recipient of several best paper awards, including ICC'06, DRCN'11, HPSR'15, and NaNa'16. He is a Winner of the MTA Lendület Program and the Google Faculty Award in 2012 and the Microsoft Azure Research Award in 2018. He is a TPC Member of leading conferences, e.g., the IEEE INFOCOM 2012 and the General Chair of the ACM SIGCOMM 2018.

Lajos Rónyai received the Ph.D. degree from the Eötvös Loránd University, Budapest, in 1987. He is currently a Research Professor with the Informatics Laboratory, Computer and Automation Institute, Budapest, Hungary. He also leads a research group there which focuses on theoretical computer science and discrete mathematics. He is also a Full Professor with the Mathematics Institute, Budapest University of Technology and Economics. His research interests include efficient algorithms, complexity of computation, algebra, and discrete mathematics. He is a member of the Hungarian Academy of Sciences and a recipient of the Count Széchenyi Prize.

Balázs Vass (Member, IEEE) received the M.Sc. degree in applied mathematics from Eötvös Loránd Science University (ELTE), Budapest, in 2016, and the Ph.D. degree in informatics from the Budapest University of Technology and Economics (BME) in 2016. His research interests include networking, survivability, combinatorial optimization, and graph theory. He received the Best Paper Award of NaNa'16 and the Best-in-Session Presentation award on INFOCOM'18. He was an invited speaker of COST RECODIS Training School on Design of Disasterresilient Communication Networks in 2019.

László Gyimóthi received the M.Sc. degree in electrical engineering from the Budapest University of Technology and Economics (BME) in 2015. He is currently pursuing the Ph.D. degree. His research interests include efficient failure localization in optical networks, combinatorial optimization, and efficient algorithm design. He founded an IT start-up in 2017 that provides a commercially available automatic scorekeeping solution for steel-tip darts.