

List of Shared Risk Link Groups Representing Regional Failures with Limited Size

János Tapolcai*, Lajos Rónyai†, Balázs Vass*‡, László Gyimóthi*

*MTA-BME Future Internet Research Group, Budapest University of Technology and Economics, tapolcai@tmit.bme.hu

‡ Computer and Automation Research Institute Hungarian Academy of Sciences, BME Department of Algebra

†Department of Operations Research, Eötvös University, Budapest, Hungary

Abstract—Shared Risk Link Group (SRLG) is a failure the network is prepared for, which contains a set of links subject to a common risk of single failure. During planning a backbone network, the list of SRLGs must be defined very carefully, because leaving out one likely failure event will significantly degrade the observed reliability of the network. Regional failures are manifested at multiple locations of the network, which are physically close to each other. In this paper we show that operators should prepare a network for only a small number of possible regional failure events. In particular, we give a fast systematic approach to generate the list of SRLGs that cover every possible circular disk failure of a given radius r . We show that this list has $O((n+x)\sigma_r)$ SRLGs, where n is the number of nodes in the network, x is the number of link crossings, and σ_r is the maximal number of links that could be hit by a disk failure of radius r . Finally through extensive simulations we show that this list in practice has size of $\approx 1.2n$.

I. INTRODUCTION

Backbone networks are designed to protect a certain pre-defined list of failures, called Shared Risk Link Groups (SRLG)¹. SRLG describes the relationship between links with a shared vulnerability. For example links with shared fibre cable or conduit have a chance to fail simultaneously, or network devices with shared power sharing, etc. The SRLGs can subtly cover the network, as each link could belong to several SRLGs. Unfortunately, SRLGs are not self-discoverable in practice [2], thus the mapping of links to SRLGs should be defined by the network operators. Operators must very carefully define the list of SRLGs, because leaving out one likely simultaneous failure event will significantly degrade the observed reliability of the network. The great number of serious network outages witnessed in the last decades [3]–[6] present clear evidence that selecting the proper list of SRLGs is still a challenging problem to solve [7]–[15]. To fill this gap in reliable network design, this paper proposes a systematic approach 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 that surprisingly the number of such SRLGs is not too high in practice.

After the list of SRLGs are 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 failure that involves links which 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 failure first hits a single network element for whose protection the network is already pre-configured. After the failure new SRLGs can be added to protect a possible multiple failure. The limitation of this approach was well studied in [7]–[15]. It turned out that the network can have serious outages when almost every equipment in a large 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 [3], the 9.0 magnitude earthquake in Japan on March 2011 impacted about 1500 telecom switching offices due to power outages [4] and damages of undersea cables, the hurricane Katrina in 2005 caused severe losses in Southeastern US [5], hurricane Sandy in 2012 caused a power outage which silenced 46% of the network in the New York area [10], [16]. Heavy rain falls, or in general weather-based region disruptions, can bring out correlated temporal failures of high capacity wireless links (as e.g., in Wireless Mesh Networks) in a small region. Another important reason for disruptions on a massive scale the network operators need to be prepared is related to intentional human activities, such as bombing or use of weapons of mass destruction attacks, electromagnetic pulse attacks. Electromagnetic pulse attack is an intense energy field that can instantly overload or disrupt numerous electrical circuits, thereby affecting networking equipments within a large geographic area [17]. Submarine cables are vulnerable to human activities such as fishing, anchors and dredging [18]. These types of failures are called *regional failures* which are simultaneous failures of nodes/links located in specific geographic areas [7]–[15]. It is still a challenging open problem how to prepare a network to protect against such failure events, as their location and size is not known at planning stage. Intuitively, the number of possible regional failures can be very large. In the paper we propose a solution to this problem with a technique that can significantly reduce the number of

¹First introduced in [1].

possible failures that should be added as SRLGs to cover all regional failures.

A regional failure is defined as a failure of multiple network elements in a geographic area, which can have any location, size and shape. We will consider the size of the regional failure as the most important property. To measure the size of a regional failure we compute the smallest circular disk that covers every failed link, and the radius of this circular disk represents its size. In this study we are interested in enumerating the SRLGs of regional failures with a given maximum size. We assume each SRLG represents a worst case scenario the network must be prepared for. For example when each connection is assigned with an SRLG-disjoint protection path. In our scenario there is no need to have an SRLG which is a subset of an other SRLG. In particular we are interested in two versions of the problem. In the first version we list every possible failures the network can have due to a circular disk failure of a given radius r . In the second version of the problem we assume the radius of the failure is not a network wide parameter, but depends on the area. For example the radius is larger on flat regions, and smaller in the hilly area. 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 paper** is a reduction of the number of SRLGs subject to regional failures 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 regional failure is assumed to be a circular disk of radius r and 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, σ_r is the maximal number of links that could be hit by a disk failure of radius r , and y_r which is the total number of link pairs whose distance is at most $2r$. In backbone networks x is a small number as typically a network node is also installed on each link crossing (similarly to road networks [19]), while σ_r represents a density of the topology, which should not depend on the network size. Using these parameters the number of SRLGs is $\Theta((n+x)\sigma_r)$ or $\Theta(m+y)$. We also give an artificial example to illustrate that these bounds are tight.
- We provide faster algorithms to enumerate the SRLGs that runs in $O((n+x)^2\sigma_r^3)$ time.
- Compared to prior art we handle parallel edges and colinear node triples.

Using the obtained SRLG list, network operators can design their networks to be protected against regional and random failures. Backbone networks designed according to our new

failure model should have higher reliability, and leave way less failures to be recovered with the convergence of higher layer intra-domain routing protocols (IS-IS, OSPF) within the next few seconds, minutes or hours after the failure. We believe the paper contributes to closing the gap between the conventional SRLG based pre-planned protection and regional failures.

The paper 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 on real backbone networks. Finally Sec. VII concludes the paper.

II. RELATED WORK

With somewhat different motivation similar computational geometric ideas were used in papers focusing on the most vulnerable points of a physical infrastructure (communication networks or power grids [20]) to regional failures or attacks. Our objective is more general as we want to enumerate all candidate failures, instead of searching for the most vulnerable according to some metric. The network is embedded in the Euclidean plane and the failures are modelled either as a disk around its epicenter (circular) [7], [21], [22], line segments [7], ellipse [23] or polygons (rectangle, square, or equilateral triangle) [23]. Technically these papers also list the candidate failures and evaluate the vulnerability metric of the residual network in case of each candidate failure. Note that, our approach computes $O((n+x)\sigma_r)$ SRLGs (in practice σ_r is constant and $x \ll n$), while the best known general worst case bound was $O(n^4)$ [21], 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 failures with varying radius.

The following vulnerability metrics are investigated: (1) the point with the maximum number of affected links [7], [21], which is σ_r . (2) the point with the maximum average two terminal reliability between every node-pair [7], [21]–[23]. (3) the point with the maximum average all-terminal reliability [13], [22] 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 given pair of nodes [7]. (5) the point with maximal average shortest path length between every pair of nodes [13], [23], (6) survivability as a measure of weighted spectrum based on the eigenvalues of the normalized Laplacian of a graph [13], (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 [13].

A special case of our problem is investigated in [24] where the goal is to list all the spatially-close fiber segments. They model links not only a straight line segment but series of line segments where the geographic location of the corner points are known. In our model the corner points can be treated as

degree 2 nodes, and with our approach the grouping of the spatially-close fiber segments can be directly computed.

The idea of defining SRLGs for disasters was also proposed in [11]; however the SRLGs (called disaster zones) were defined 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 tornadoes (clustered in a region) may span up to 96 and 160 km, respectively. 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. 8b).

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 [25], SRLG disjoint paths [26], etc.

Related to our work is the research in computational geometry on the smallest intersecting ball problem [27], [28], which has its origins in the classical 19th century problem of Sylvester [29] about the smallest enclosing circle for a given set of points in the plane.

III. PROBLEM DEFINITION AND BASIC RESULTS

The input is a real number $r > 0$ and an undirected 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, line segments or intervals in geometric proofs). 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. A *disk* with centre point p covers 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 covering every edge in F .

Note that the failure of node v is modelled 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 the 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 SRLG fails 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




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

$$\mathcal{S}_r = \{F \mid \text{is a regional failure and there is no regional failure } F' \text{ such that } F' \supsetneq F\} . \quad (1)$$

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. 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. Clearly, \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 covered by c .

Let r be a very small 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 generalisation of the 'best practice'. The corresponding Sperner system can be the set of single node failures, i.e. $|\mathcal{S}_r| = n + x$, where x is the number of edge crossings. Informally speaking protecting node failures is sufficient to protect link failures as well.

Our aim is 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 σ_r denote the maximum number of edges a disk with radius r can cover in the plane, i.e. for every failure F caused by a disk with radius r , $|F| \leq \sigma_r$. We observe that if $\sigma_r = O(\log n)$ then there is a polynomial blowup when we switch from \mathcal{S}_r to \mathcal{H}_r , as $|\mathcal{H}_r| \leq |\mathcal{S}_r| 2^{\sigma_r}$. 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 covers a set of edges E_c , if it covers all the edges in E_c . Note that several disks can result the failure of the same set of edges.

First we give a slight variant of Lemma 9 from [7]. 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 covered by a circular disk of radius r . Then there is a disk c of radius r




Fig. 2. The disk failures examined in Thm. 1

which covers the intervals of H such that at least one of the following holds (see Fig. 1 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 -neighborhood² $N(e, r)$ 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. 1d) of the plane.

Consider the boundary B of the intersection

$$\bigcap_{e \in H} N(e, r). \quad (2)$$

The points of B are obviously in the union of the boundaries of the neighbourhoods $N(e, r)$, 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. 1d) 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 covers case (1). Note that we can also assume now that $|H| > 1$.

(2) No point of B is on a circular arc from 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

²called hippodrome in [21].




Fig. 3. An example topology ($k=4$) where the number of SRLGs is $\Omega(m^2)$ or $\Omega((n+x)\sigma_r)$ for circular disk failures.

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 an other 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 to 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 , $|H'| \leq 2$. Then every circular disk described in Thm. 1 can be determined in $O(1)$ time.*

Proof: Easy elementary geometric discussion of cases (a), (b) and (c) of Thm. 1. See Fig. 2 for illustration. Note that there can be at most 4 circles that intersect two line segments as shown on Fig. 2(a), and at most two circles intersecting a line segment and a single point as shown Fig. 2(b), and four circles can touch a line at an endpoints as shown Fig. 2(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 above bound is asymptotically tight.

A. Worst Case Graph

Claim 1. *The graph sketched in Fig. 3 has at least $\frac{n^2}{64}$ 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.³

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

³No attempt has been made to optimize the constant. In fact, a more elaborate variant of the preceding construction gives $\frac{n^2}{16}$ failure events.

edge connecting O to $(k,0)$, and put its copies translated i units downwards, for $i = 1, \dots, k$ 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 $i = 1, \dots, k$. 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 on Fig. 3. 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 cover the same set of edges. This implies that there are at least $(k + 1)^2$ failure events with radius k . The number of vertices is $n = 8k$. The number of such failures is at least $\frac{n^2}{64}$. ■

B. Circular Disk Failures with Radius at Most r

In this subsection we take a more general model and assume the radius of the failure is not a network wide parameter, but depends on the area. Our goal is to enumerate every circular disk failures 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 covering F by c_F and we say c_F is the **smallest covering 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 covering disks c_F , $\forall F \in \mathcal{H}_r$, and simply ignore the rest of the disk failures. The consequence of the next theorem is that the number of smallest covering disks c_F , $\forall F \in \mathcal{H}_r$ 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 smallest covering disks were defined on sets of nodes because a triplet of non collinear nodes defines a circle. In the proof in Appendix VIII-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 VIII-B.

Remark. Thm. 3 outlines an efficient algorithm for c_H in an exact symbolic computational setting. A good numerical algorithm for approximating r and P 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 covering disk of a triplet of edges can be calculated in $O(1)$ time, we could solve the problem by processing $O(m^3)$ triplet of edges. However, we will achieve better upper bounds on 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 three practical parameters of the input to better estimate the number of SRLGs.

- y_r is the total number of edge pairs whose distance is at most $2r$.
- x is the number of link crossings of the network G .
- σ_r is the *link density* of the network which is measured as the maximal number of links that could be hit by a disk failure of radius r .

In backbone networks x is a small number as typically a network node is also installed on each link crossings [19], while the link density σ_r practically should not depend on the network size. We also know that σ_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 $x = |X|$.

Despite the fact that on arbitrary graphs x can be even $\Theta(n^4)$, in backbone network topologies typically $x \ll n$ because a switch is usually installed if two cables are crossing each other⁴. This gives us the intuition that G is “almost” planar, and thus it has few edges.

Claim 2. The number of edges in G is $O(n+x)$. More precisely for $n \geq 3$ we have $m \leq 3n + x - 6$.

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

⁴Recent experimental studies give empirical evidence that real-world road networks typically have $\Theta(\sqrt{n})$ edge crossings [30].




Fig. 4. Illustration to Thm. 4

Note that X can be computed in $O(n + x \log^k n)$ expected time, where k is a constant [19]. It is a randomized algorithm which is called geometric graph planarization. For the sake of simplicity we assume the planarized graph is simple. Since the average running time is linear for practical network topologies, we omit the complexity of computing the set of X from our algorithms in enumerating the SRLGs.

Now we can develop an upper bound on the number of SRLGs. Let y_r denote the number of edge pairs whose distance is at most $2r$.

Corollary 3. $|\mathcal{S}_r| \leq 12(n + y_r)$.

Proof: For case (a) of Thm. 1 we need to compute at most $4y_r$ circles, for case (b) we have $4y_r$, and finally for case (c) we have $4m \leq 4(3n + x - 6) < 12n + 4y_r$, because $x \leq y_r$. ■

This bound is asymptotically tight as shown by the graphs in Claim 1 because $y_r = O(n^2)$.

The next theorem states, it is enough to process the edge triplets in the neighbourhood 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 covering H with centre point at distance at most $2r$ from $V \cup X$.*

The proof of the theorem can be found in Appendix VIII-C.

Theorem 5. *Let r be a positive real number, $F \in \mathcal{S}_r$ be a set of line segments which can be covered 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 , covers 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 another 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 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 covering radius for F . Then $c = c_F$ will be a suitable disk. In fact 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. 4). Then K is a proper convex k -gon for some $k \geq 3$, 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 failures 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 $15\sigma_r$ possibilities for the other edge besides e in Theorem 5 (a) or (b) (see Fig. 5, 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. $|\mathcal{S}_r| = O((n + x)\sigma_r)$. ■




Fig. 5. A disk with radius $3r$ can be covered with 15 disks with radius r

This bound is asymptotically tight⁵ on the graphs in Claim 1 because $\sigma_r = \frac{n}{2}$ for $r = k$. Next we discuss how the list of SRLGs can be generated.

Lemma 2. *The set of SRLGs for circular disk failures of radius r can be computed in $O((n + x)^2 \sigma_r^3)$ time.*

Proof: Briefly, for every pair of edges whose distance is at most $2r$, we generate the disks c of radius r described in Lemma 1. Next, for each c we compute the set of edges F covered by disk c . Then we save F as an SRLG in the list \mathcal{S}_r , only if none of the SRLGs in \mathcal{S}_r is a superset of F .

For implementing the algorithms we use the following data structure. For each pair (e, v) where $e \in E$, $v \in V \cup X$, and $v \in e$ we will store a lists of the edges which are within distance $3r$ from node v . To do so, first we compute the distance between every edge and every point in $V \cup X$. Overall, the above data

⁵No attempt have been made to optimize the constant.




Fig. 6. 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.

structure can be built up in $O(m(n+x)) = O((n+x)^2)$ steps, and the lists have at most $O(\sigma_r(n+x))$ elements altogether.

For every pair (e, v) and an edge f from its list we form the disks c of radius r (if there is any) covering this pair of edges as in Theorem 5 (in particular, the center of c must be at most $2r$ from v). For a given pair (e, v) we have $O(\sigma_r(n+x))$ disks.

For each of these disks we compute the set of edges F . Finally, we need to check if F is not a subset of any listed SRLG, before we add F as a new SRLG to the list, and erase every SRLG in the list that is a subset of F . This is done by comparing F with every SRLGs in the lists, which requires comparing $O((n+x)^2\sigma_r^2)$ pairs of segment sets (F, F') .

Finally, to compare two SRLGs to test if they are subset of each other requires $O(\sigma_r)$ steps. To do so, we need to assign an order to the edges of the network and store the edges of an SRLG as an ordered list. ■

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

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

Proof: We can use Theorems 2 and 4 and the fact that a disk of radius $3r$ covers $O(\sigma_r)$ 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 covering disk of every set E_v containing an edge going through v and containing 1 or 2 edges from the $3r$ neighbourhood of v . For a fixed v we have $O(\sigma_r^2)$ SRLGs, and the claim follows. ■

Theorem 6. The circular disk failures with radius at most r can be computed in $O((n+x)^2\sigma_r^5)$ time.

Proof: According to Proposition 2 there are $O((n+x)^2\sigma_r^4)$ circular disk failures to examine, where each SRLG has at most σ_r edges. A pair of candidate SRLGs (F, F') can be checked in time $O(\sigma_r)$ for possible containment. ■

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




Fig. 7. The set of SRLGs of a 5×5 grid network.

other. The network topologies with the obtained list of SRLGs for various radii are available online⁶.

Each SRLG in the obtained list of SRLGs \mathcal{S} is visualized by its smallest covering disk. According to Thm. 3 the smallest covering 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 covering 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.

Fig. 6 shows the Italian optical backbone network with circular disk failures 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 were $O((n+x)\sigma_r)$, and σ_r increases with the radius. Surprisingly, as the radius increases the number of SRLGs does not increase, but stays close to $n+x$. This is because the SRLGs which are subset of an other SRLGs are filtered out. Note that our bounds are asymptotically tight for the artificial networks on Fig. 3. In other words, it seems in practical scenarios the number of SRLGs does not depend on the radius.

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. 7a. 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. 7b). As the radius further increases to $r = \frac{\sqrt{2}}{2}$ the SRLGs will be every facets with the

⁶<https://github.com/jtapolcai/regional-srlg>

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 covering 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 understand 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 disk failure. Fig. 8 shows our results where σ_r the maximal number of links in the SRLG is also listed. Clearly, σ_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 failures of radius 200km. Note that, in the European network the nodes are closer to each other compared to 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 covers the disaster zones for earthquakes, tornadoes, and weapons of mass destruction attacks defined in [11].

Table I 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 metric σ_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.

TABLE I
RESULTS ON SOME BACKBONE TOPOLOGIES FROM [31]

Network	n	m	x	σ_r	y_r	$ \mathcal{S}_r $	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 [11]	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

VII. CONCLUSIONS

In this paper we view networks as geometric graphs and propose a fast and systematic approach to enumerate the list of possible link failures caused by regional disruptions. Our approach assumes that the regional failure has a shape of disk of a given radius. Although the number of possible regional failures is infinite, we show that under reasonable and realistic assumptions the list of failures to be considered is short, it is basically linear in the network size. We present two fast

polynomial time algorithms, the first lists every disk failure of a fixed radius. Its main idea is to move a disk of radius r to every candidate location. The second algorithm lists every disk failure with radius at most r . This allows a more sophisticated regional failure model where different radii of failure are used at 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 $\approx 1.2n$.

REFERENCES

- [1] L. Mollenauer, J. Gordon, P. Mamyshev, I. Kaminow, and T. Koch, "Optical fiber telecommunications iiiia," 1997.
- [2] J. Strand, A. L. Chiu, and R. Tkach, "Issues for routing in the optical layer," *IEEE Communications Magazine*, vol. 39, no. 2, pp. 81–87, 2001.
- [3] D. M. Masi, E. E. Smith, and M. J. Fischer, "Understanding and mitigating catastrophic disruption and attack," *Sigma Journal*, pp. 16–22, 2010.
- [4] Y. Nemoto and K. Hamaguchi, "Resilient ict research based on lessons learned from the great east japan earthquake," *IEEE Communications Magazine*, vol. 52, no. 3, pp. 38–43, 2014.
- [5] 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 Systems Journal*, vol. 3, no. 3, pp. 277–287, 2009.
- [6] Y. Ran, "Considerations and suggestions on improvement of communication network disaster countermeasures after the wenchuan earthquake," *IEEE Communications Magazine*, vol. 49, no. 1, pp. 44–47, 2011.
- [7] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of the fiber infrastructure to disasters," *IEEE/ACM Transactions on Networking (TON)*, vol. 19, no. 6, pp. 1610–1623, 2011.
- [8] O. Gerstel, M. Jinno, A. Lord, and S. B. Yoo, "Elastic optical networking: A new dawn for the optical layer?," *Communications Magazine, IEEE*, vol. 50, no. 2, pp. s12–s20, 2012.
- [9] M. F. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of disaster-resilient optical datacenter networks," *Journal of Lightwave Technology*, vol. 30, no. 16, pp. 2563–2573, 2012.
- [10] J. Heidemann, L. Quan, and Y. Pradkin, *A preliminary analysis of network outages during hurricane sandy*. University of Southern California, Information Sciences Institute, 2012.
- [11] F. Dikbiyik, M. Tornatore, and B. Mukherjee, "Minimizing the risk from disaster failures in optical backbone networks," *Journal of Lightwave Technology*, vol. 32, no. 18, pp. 3175–3183, 2014.
- [12] I. B. Harter, D. Schupke, M. Hoffmann, G. Carle *et al.*, "Network virtualization for disaster resilience of cloud services," *Communications Magazine, IEEE*, vol. 52, no. 12, pp. 88–95, 2014.
- [13] X. Long, D. Tipper, and T. Gomes, "Measuring the survivability of networks to geographic correlated failures," *Optical Switching and Networking*, vol. 14, pp. 117–133, 2014.
- [14] B. Mukherjee, M. Habib, and F. Dikbiyik, "Network adaptability from disaster disruptions and cascading failures," *Communications Magazine*, vol. 52, no. 5, pp. 230–238, 2014.
- [15] R. Souza Couto, S. Secci, M. Mitre Campista, K. Costa, and L. Maciel, "Network design requirements for disaster resilience in iaas clouds," *Communications Magazine, IEEE*, vol. 52, no. 10, pp. 52–58, 2014.
- [16] G. O'Reilly, A. Jrad, R. Nagarajan, T. Brown, and S. Conrad, "Critical infrastructure analysis of telecom for natural disasters," in *Int. Telecommunications Network Strategy and Planning Symposium (Networks)*. IEEE, 2006, pp. 1–6.
- [17] J. S. Foster Jr, E. Gjeldre, W. R. Graham, R. J. Hermann, H. M. Kluepfel, R. L. Lawson, G. K. Soper, L. L. Wood, and J. B. Woodard, "Report of the commission to assess the threat to the united states from electromagnetic pulse (EMP) attack: Critical national infrastructures," DTIC Document, Tech. Rep., 2008.
- [18] R. Wilhelm and C. Buckridge, "Mediterranean fibre cable cut—a ripe ncc analysis," 2008. [Online]. Available: <https://www.ripe.net/data-tools/projects/archive/mediterranean-fibre-cable-cut>

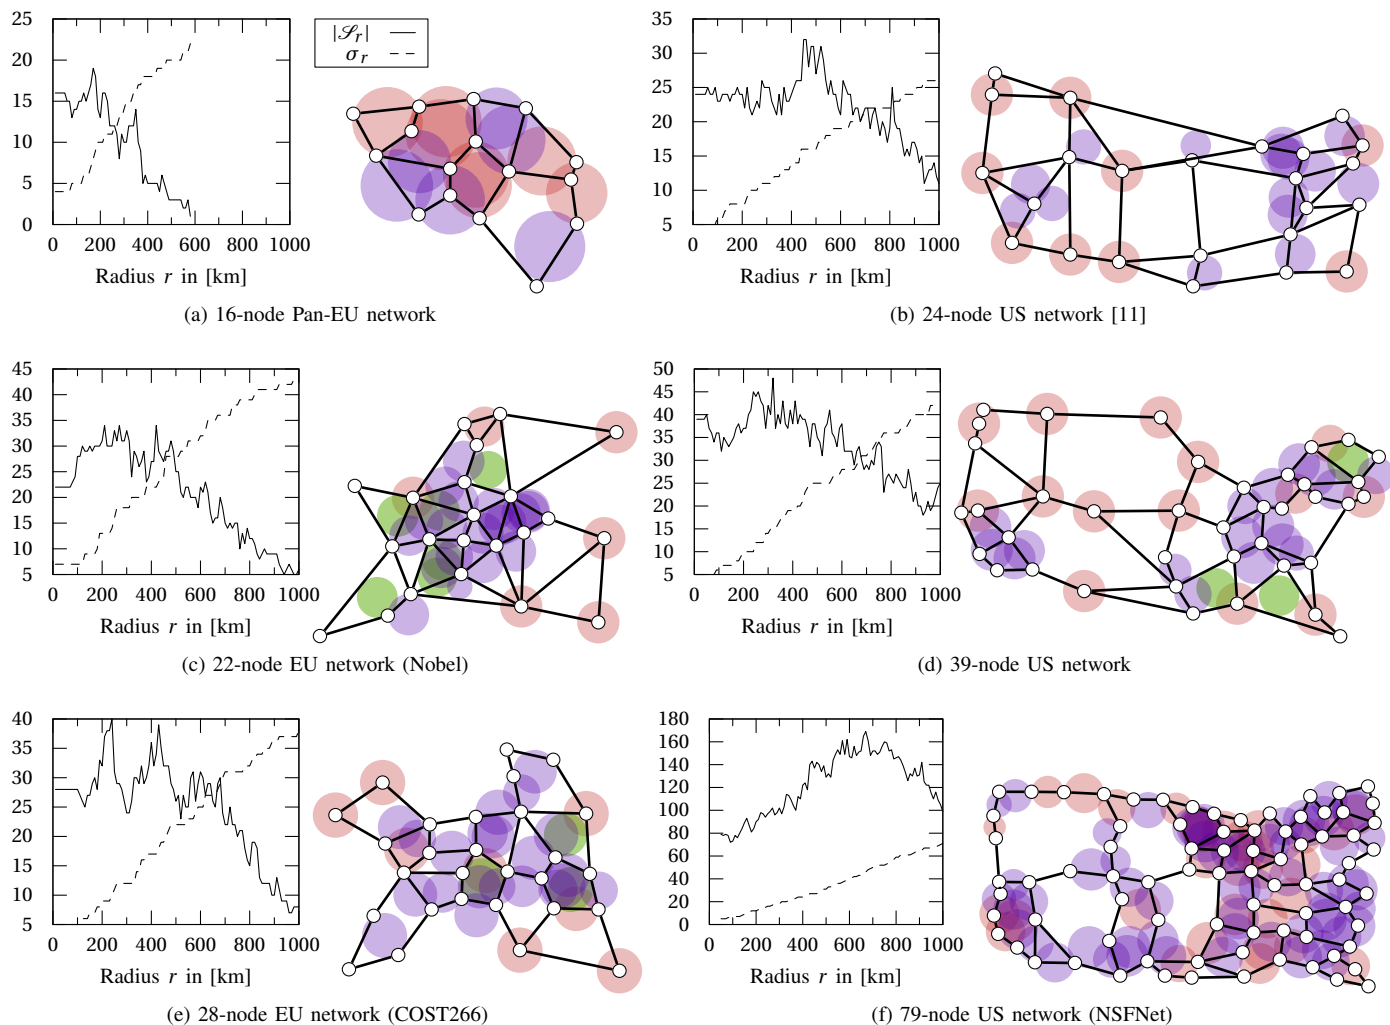


Fig. 8. The number of SRLGs $|\mathcal{S}_r|$ vs. the radius r of regional failure. The number of edges in the largest SRLG, σ_r , is also plotted. The graph topologies with the SRLGs of radius 200km is plotted next to the charts.

- [19] D. Eppstein, M. T. Goodrich, and D. Strash, "Linear-time algorithms for geometric graphs with sublinearly many edge crossings," *SIAM Journal on Computing*, vol. 39, no. 8, pp. 3814–3829, 2010.
- [20] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman, "Power grid vulnerability to geographically correlated failures: Analysis and control implications," in *IEEE INFOCOM*, 2014, pp. 2634–2642.
- [21] 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 Transactions on Networking (TON)*, vol. 21, no. 5, pp. 1525–1538, 2013.
- [22] M. T. Gardner and C. Beard, "Evaluating geographic vulnerabilities in networks," in *IEEE Int. Workshop Technical Committee on Communications Quality and Reliability (CQR)*, 2011, pp. 1–6.
- [23] S. Trajanovski, F. Kuipers, P. Van Mieghem *et al.*, "Finding critical regions in a network," in *IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2013, pp. 223–228.
- [24] F. Iqbal, S. Trajanovski, and F. Kuipers, "Detection of spatially-close fiber segments in optical networks," in *Design of Reliable Communication Networks (DRCN)*. IEEE, 2016, pp. 95–102.
- [25] M. T. Gardner, R. May, C. Beard, and D. Medhi, "Finding geographic vulnerabilities in multilayer networks using reduced network state enumeration," in *Design of Reliable Communication Networks (DRCN), 2015 11th International Conference on the*, March 2015, pp. 49–56.
- [26] J.-C. Bermond, D. Coudert, G. D'Angelo, and F. Z. Moataz, "Srlg-diverse routing with the star property," in *Design of Reliable Communication Networks (DRCN)*. IEEE, 2013, pp. 163–170.
- [27] N. M. Nam, T. A. Nguyen, and J. Salinas, "Applications of convex analysis to the smallest intersecting ball problem," *Journal of Convex Analysis*, vol. 19, no. 2, pp. 497–518, 2012.
- [28] S. Jadhav, A. Mukhopadhyay, and B. Bhattacharya, "An optimal algorithm for the intersection radius of a set of convex polygons," *Journal of Algorithms*, vol. 20, no. 2, pp. 244–267, 1996.
- [29] J. J. Sylvester, "A question in the geometry of situation," *Quarterly Journal of Pure and Applied Mathematics*, vol. 1, 1857.
- [30] D. Eppstein and M. T. Goodrich, "Studying (non-planar) road networks through an algorithmic lens," in *ACM SIGSPATIAL international conference on Advances in geographic information systems*. ACM, 2008, p. 16.
- [31] S. Orłowski, R. Wessäly, M. Pióro, and A. Tomaszewski, "Sndlib 1.0: survivable network design library," *Networks*, vol. 55, no. 3, pp. 276–286, 2010.
- [32] J. Matoušek, *Lectures on discrete geometry*. Springer New York, 2002, vol. 108.
- [33] R. Hartshorne, *Geometry: Euclid and beyond*. Springer Science & Business Media, 2013.

VIII. 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 [32]) 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. \blacksquare

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

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

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 $r^* = \max_{H'} r_{H'}$ we have $\cap_{e \in H} N(e, r^*) \neq \emptyset$ by the planar Helly's theorem (Theorem 1.3.2 in [32]) applied to the convex sets $N(e, r^*)$, hence H can be covered 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 covering 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 covering 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 $N(e_i, r)$. We have $t = s'$.

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. \blacksquare




Fig. 9. Illustration of the cases of Thm. 3

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 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 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. 9a).

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 point 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. 9b). 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.

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 from i and j . Note that it may occur that all the segments are 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 intersect 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 from P (similar to Fig. 4 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 covering $e_i \cap e_j$ (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 . Also 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 $P_i P_j$, where $j \neq i, k$. This gives at most two possible points for P . We take the smallest of the resulting circles as a candidate 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. 10).

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




Fig. 10. Smallest covering disk of 2 points and 1 line

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. 11).




Fig. 11. Smallest covering disk of 1 point and 2 lines

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 [33]) there are four circles 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 can not be optimal⁷ (see also Fig. 12).

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 just computes the candidates for c_H described at the subcase, if the input data permits it⁸. There is always just

⁷The shortest arc of the other circles containing the three points of intersection with the ℓ_i is less than half of the complete circle, hence the points can be covered by a smaller disk.

⁸For 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. Then 4d does not apply to the actual input.




Fig. 12. Smallest covering disk of 3 lines

a constant number of circles to form. The minimal one can 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 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 then 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 the third interval e_k , 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 . 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 $e_i \cap e_j$ ⁹, 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 triangle (see Fig. 13), as the other circles touching the three lines all have larger radii. We assume w.l.o.g. that the largest angle of the triangle is \widehat{BAC} . Then the radian measure of \widehat{BAC} is in interval $[\frac{\pi}{3}, \pi)$. This way $\sin(\widehat{IAC}) \geq \frac{1}{2}$ since AI is the angle bisector of \widehat{BAC} . Moreover,

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

⁹If 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 .

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 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. ■




Fig. 13. Illustration for the proof of Thm. 4