Enumerating Circular Disk Failures Covering a Single Node

Balázs Vass*†, Erika R. Bérczi-Kovács*, János Tapolcai†
*Department of Operations Research, Eötvös University, Budapest, Hungary, vassbalazs@gmail.com,koverika@cs.elte.hu
†MTA-BME Future Internet Research Group, Budapest University of Technology and Economics, tapolcai@tmit.bme.hu

***Abstract*—Current backbone networks are designed to protect a certain pre-defined list of failures, called Shared Risk Link Groups (SRLG). The list of SRLGs must be defined carefully, because leaving out one likely failure event significantly degrades the observed reliability of the network. In practice, the list of SRLGs is typically composed of every single link or node failure. It has been observed that some type of failure events manifested at multiple locations of the network, which are physically close to each other. Such failure events are called regional failures, and are often caused by a natural disaster. A common belief is that the number of possible regional failures can be large, thus simply listing them as SRLGs is not a viable solution. In this study we show the opposite, and provide an efficient algorithm enumerating all the regional failures having at a node failed as SRLGs. According to some practical assumptions this list is surprisingly short with $O(|V|)$ SRLGs in total, and can be computed in $O(|V|^2)$ time.**

I. Introduction

Current backbone networks are built to protect a certain list of failures. Each of these failures (or termed failure states) are called *Shared Risk Link Groups* (SRLG), which is a set of links that is expected to fail simultaneously. The network is designed to be able to automatically reconfigure in case of a single SRLG failure, such that every connection further operates after a very short interruption. In practice it means the connections are reconfigured to by-pass the failed set of nodes and links. Thus the network can recover if an SRLG or a subset of links and nodes in SRLG fails simultaneously; however, there is no performance guarantee when a network is hit by a failure that involves links not a subset of an SRLG. Nevertheless, the list of SRLGs must be defined very carefully, because not getting prepared for one likely simultaneous failure event the observed reliability of the network significantly degrades.

Operators have numerous best practices how to define the list of SRLGs. One extreme 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. Often there is a known risk of a simultaneous multiple failure that can be added as an SRLG, for example if two links between different pair of nodes traverse the same bridge, etc. On the other hand, we have witnessed serious network outages [1]–[9] because of a failure event that takes down almost every equipment in a physical region as a result of a disaster, such as weapons of mass destruction attacks, earthquakes, hurricanes, tsunamis, tornadoes, etc. For example the 7.1-magnitude earthquake of Taiwan in Dec. 2006 caused six submarine links between Asia and North America to fail simultaneously [10], the 9.0 magnitude earthquake in Japan Earthquake on March 2011 impacted about 1500 telecom switching offices due to power outages [5] and also damaged undersea cables, hurricane Katrina in 2005 caused severe losses in Southeastern US [11], hurricane Sandy in 2012 caused a power outage that silenced 46% of the network in the New York area [4]. These type of failures are called *regional failures* which is a simultaneous failure of nodes/links located in a specific geographic area [1]–[9]. The number of possible regional failures can be very large, thus simply listing them as SRLGs is not a viable solution. It is still a challenging open problem how to prepare a network to protect such failure events, as their location and size is not known at planning stage. In this paper we propose a solution to this problem in case of circular disk failures covering a single node with a technique that can significantly reduce the number of possible failure states that should be listed as SRLGs to cover all of this type failures.

In practice, regional failures can have any location, size and shape. It is a common best practice to fix the size or shape of regional failures, for example as cycles with given size (also called disk) [12]. In our study we assume the regional failure has a shape of cycle but do not fix its size. Instead we classify the regional failures according to the network elements they cover. For example if a failure hits a node, the node is no longer going to send traffic, which has a network wide effect. In this paper node failures are treated separately form link failures. The first class of failures is the class of disk failures that hit links only. Clearly, a disk failure that does not cover any node cannot be too large. The second class is the class of disk failures that affect nodes too besides links. Reconfiguring a network after a failure from this second class is more complicated than reconfiguring after pure link failures It is because the failure can cause a degradation in the control plane communication. Besides, after a node failure the traffic matrix often changes due to the fact that the failed node will no longer participate in communication. In particular we are interested in regional failures that cause the failure of at most one node. We assume these regional failures are large enough, thus regional failures covering multiple nodes are left for further study.

In our previous work we focused on the first class of failures covering links only. We have shown that the list of such failures is short in practice, i.e. $O(|V|)$ SRLGs in total, and can be computed very fast, in $O(|V|\log|V|)$ time [13]. The main contribution of this paper to generalise the results for second class of regional failures covering one node.

In this study the number of SRLGs to be listed is significantly reduced applying computational geometric tools based on the following assumptions:

1) The network is a geometric graph $G = (V, E)$ embedded in a 2D plane.
2) The shape of the regional failure is assumed to be a circle with arbitrary radius and center position.
3) We focus on **regional link single node failures**, failures that affect exactly one node.

We will show that the number of SRLGs corresponding to exclusion-wise maximal link single node failures is small, $O(|V|)$, in typical backbone network topology, where $|V|$ and $|E|$ denote the number of nodes and links in the network, respectively. We propose a systematic approach based on computational geometric tools that can generate the list of SRLGs in $O(|V|^2)$ steps on typical networks.

The paper is organised as follows. In Section II and III we present the core mathematical model with several observations. In Section IV we show the main result, the $O(|V|^2)$ algorithm for generating the list of SRLGs covering every regional link failure with a shape of disk covering at most one network node. Finally, in Section V we reflect on related works and in Section VI we draw the conclusions.

II. Model and Assumptions

We model the network as a connected undirected geometric graph $G = (V, E)$ with $n = |V|$ nodes and $m = |E|$ edges, we assume $n \geq 3$. The nodes of the graph are embedded as points in the Euclidean plane, and the edges are embedded as line segments. The position of node v is denoted by (v_x, v_y). A *disk* $c(x, y, r)$ is a circle with a centre point (x, y) and radius r. The failure caused by a disk is modelled as every interior node and edge with interior part is erased from the graph. One can observe that a node v is erased or becomes isolated iff every edge incident to v is erased, thus listing erased nodes beside listing erased edges has no additional information from the viewpoint of connectivity. This way for our purposes it is enough to determine the set $\mathscr{M}_1$ of the maximal sets of erased edges. In order to determine $\mathscr{M}_1$ we need to introduce some notations.

It will turn out that every element of $\mathscr{M}_1$ can be covered by a disk which has 1 node interior and 2 nodes on its boundary (Claim 2). Let denote $C^w_{1,\{u,v\}}$ the set of disks covering only node w, and having node pair $\{u, v\}$ on the boundary; let denote $C_{1,\{u,v\}}$ the set of disks exactly 1 node and having node pair $\{u, v\}$ on the boundary; let denote C^w_1 the set of disks covering only node w; and finally, let denote C_1 the set of disks having exactly one node in interior.

For every disk c let E_c denote the set of edges and nodes erased by c.

Let $\mathscr{F}_1$ and $\mathscr{F}_{1,\{u,v\}}$ be the set of failures caused by elements of C_1 and $C_{1,\{u,v\}}$, respectively. Formally, $\mathscr{F}_1 = \{E_c | c \in C_1\}$ and $\mathscr{F}_{1,\{u,v\}} = \{E_c | c \in C_{1,\{u,v\}}\}$. We call the elements of $\mathscr{F}_1$ *regional link single node failures*, or simply *link single node failures*.

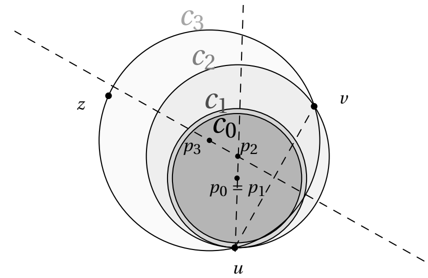

Fig. 1: For disk $c_0 \in C_i$ $(i \in \{0, 1\})$ there exists a $c_2 \in C_i$ such that $c_0 \subseteq c_2$ and c_2 has at least 2 nodes on its boundary. Note that $c_3 \in C_i$ going through nodes u, v and z and E_{c_3} not necessarily contain c_0 and E_{c_0}, respectively.

Let denote $\mathscr{M}_1$ and $\mathscr{M}_{1,\{u,v\}}$ the exclusion-wise maximal elements of $\mathscr{F}_1$ and $\mathscr{F}_{1,\{u,v\}}$, respectively. Our goal is to determine $\mathscr{M}_1$.

Proposition 1. *For any $c_1, c_2 \in C$, $c_1 \subseteq c_2$ it holds that $E_{c_1} \subseteq E_{c_2}$.* □

Let C_0 denote the set of disks that have no node of V in their interior. Clearly, $|C_i|$ is infinite for every $i \in \{0, 1\}$.

Claim 2. *For any $c_1(x, y, r) \in C_i$, $i \in \{0, 1\}$ with $E_{c_1} \neq \emptyset$ there exists a $c_2 \in C_i$ such that $c_1 \subseteq c_2$ (so $E_{c_1} \subseteq E_{c_2}$) and c_2 has 2 nodes of V on its boundary.*

Proof. Disk c_2 is generated as follows. We start with circle c_0 (having centre point p_0). and start to increase its radius until it reaches a node $u \in V$. Let the new disk be c_1 (see Fig. 1). We can further blow the circle larger without loosing any covered area by moving the centre point along the line $(x, y) - (u_x, u_y)$ while keeping u on the boundary.

Assume indirectly that it never reaches a second node. We get a contradiction because c_1 intersects line ab and $a, b \in V$.

Thus the circle will reach a second node $v \in V$. Let $c_2 \supseteq c_1$ be this circle having $u, v \in V$ on its boundary and $E_{c_2} \supseteq E_{c_1}$. Clearly, if $c_1 \in C_i$, then $c_2 \in C_i$.

□

Definition 3. *Graph $D_{\Delta 1} = (E_{\Delta 1}, V)$ is called* Delaunay-1 graph, *where $\{u, v\} \in E_{\Delta 1}$ iff there exists a $c \in C_1$ such that u and v are on the boundary of c.*

Claim 4. $\mathscr{M}_1 \subseteq \bigcup_{\{u,v\} \in E_{\Delta 1}} \mathscr{M}_{1,\{u,v\}}$.

Proof. Clearly, for all $f \in \mathscr{M}_1$ there exists a $c_1 \in C_1$ such that $f = E_{c_1}$. According to Claim 2, there exists a $c_2 \in \cup_{\{u,v\} \in E_{\Delta 1}} C_{1,\{u,v\}}$ for which $E_{c_2} \supseteq E_{c_1}$. This implies $f \subseteq E_{c_2}$. Since f is an exclusion-wise maximal element of $\mathscr{F}_1$ by definition of $\mathscr{M}_1$, this is possible only if $f = E_{c_2}$.

We get that for every $f \in \mathscr{M}_1$ there exists a $c_2 \in \cup_{\{u,v\} \in E_{\Delta 1}} C_{1,\{u,v\}}$ such that $f = E_{c_2}$. This implies $\mathscr{M}_1 \subseteq \cup_{\{u,v\} \in E_{\Delta 1}} \mathscr{M}_{1,\{u,v\}}$. □

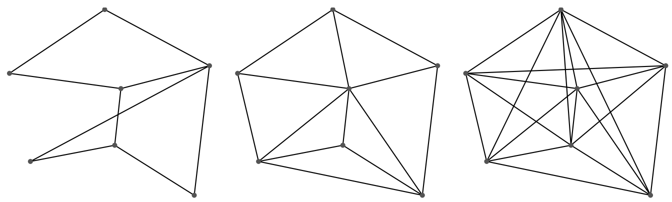

(a) Network topology (b) Delaunay graph E_Δ (c) Delaunay-one $E_{\Delta 1}$

Fig. 2: The network topology can be very different from the Delaunay an Delaunay-one graphs. $E_\Delta \subseteq E_{\Delta 1}$ (Claim 6).

It is useful to introduce graph $D_\Delta = (E_\Delta, V)$ too, which is called Delaunay triangulation, where $\{u, v\} \in E_\Delta$ iff there exists a $c \in C_0$ such that u and v are on the boundary of c. On Figure 2 we show an example on a triplet of G, D_Δ and $D_{\Delta 1}$ on a given point set V. Connections between these graphs will be discussed in Section III.

Before presenting the algorithm for determining $\mathscr{M}_1$, we bound its size. It turns out that $|\mathscr{M}_1|$ is $O(nm)$, and in case of some artificial network families $|\mathscr{M}_1|$ is $\Theta(n^3)$ [13].

This gives us the idea to use some parameters which are in relation with the size of $\mathscr{M}_1$ and with the computational complexity.

We use parameters θ_0 and θ_1 for the maximum number of edges in E crossing the circumcircle of a Delaunay triangle, and a disk $c \in C_1$, respectively. Clearly, these parameters are small in case of backbone networks.

Claim 5. $\theta_0 \leq \theta_1$.

Proof. Enlarging the circumcircle c_0 of a Delaunay triangle u, v, w while keeping u and v on the boundary gives a disk c_1 in C_1 (see proof of Claim 6). Note that c_1 intersects at least as many edges in E as c_0, which proves the claim. □

Since $|\mathscr{M}_1|$ can be asymptotically large, it is not possible to give an algorithm which is "really fast" on all graphs. On the other hand, our algorithm computes $\mathscr{M}_1$ in $O(n^2\theta_1^3)$ time (Thm. 12), giving $O(n^2)$ if θ_1 is constant, which is a natural assumption for many types of networks.

III. Structure of graph Delaunay-1 ($D_{\Delta 1}$)

In this section we clarify the size, structure and construction complexity of graph $D_{\Delta 1}$.

First let us ignore the edges of the network and focus only on the nodes. For a moment assume that we are searching for disks of maximum size from the set C_0 of circles that have no node in interior. As it turns out, finding these circles can be done using disks of maximum size from the set C_0, thus we should first find these disks. Clearly, each disk of maximum size from C_0 passes through at least three nodes, otherwise its size could be further increased. By simplicity we assume that the nodes are in general position i.e. no four nodes are on the same cycle and no three nodes are on the same line. In this case by connecting the three nodes we get non-degenerate triangles. The problem was deeply investigated in the past and it was shown the union of these triangles results a triangulation of the graph, called Delaunay triangulation [14]. Let $D_\Delta = (E_\Delta, V)$ denote the Delaunay triangulation on the set of nodes, where E_Δ denotes the edges of the triangulation, which can be very different form the edges of the network. In Delaunay triangulation no circumcircle of any triangle contains node in interior. Another interesting property is that the dual graph of the Delaunay triangulation is the so-called called Voronoi diagram [15].

The high level idea of this study is to generalise the above idea for the set C_1 of circles that have exactly one node in interior. We have defined graph Delaunay-one, which is similar to the Delaunay triangulation, but corresponding to cycles with exactly one node in interior. Note that the Delaunay-one graph is typically not planar any more.

Claim 6. $E_\Delta \subseteq E_{\Delta 1}$.

Proof. If $\{u, v\} \in E_\Delta$, then there exists a common neighbour w of u and v in D_Δ. Because of the general position of the nodes there exists a $c \in C_1$ such that u and v are on the boundary of C_1, and w is in its interior, thus $\{u, v\} \in E_{\Delta 1}$. □

It will be shown that the cardinality of $E_{\Delta 1}$ is linear in n, $D_{\Delta 1}$ having fewer than two times as many edges as a planar graph can have (Lemma 10). On the other hand, the time complexity of constructing $E_{\Delta 1}$ is $O(n \log n)$ (Lemma 9).

In this section we use the following notations.

For every $v \in V$ let $\Gamma(v)$ denote the neighbourhood of v in the Delaunay triangulation, and let $d_\Delta(v) := |\Gamma(v)|$, i.e. the degree of v in the Delaunay triangulation. Let $u_1, u_2, \ldots, u_{d_\Delta(v)}$ be the nodes of $\Gamma(v)$ enumerated in clockwise order.

Thus if v is not in the set V_H of nodes on the convex hull of V, $\{u_1, u_2\}, \ldots, \{u_{d_\Delta(v)}, u_1\} \in E_\Delta$ and if node $v \in V_H$ then we can assume w.l.o.g. that $\{u_1, u_2\}, \ldots, \{u_{d_\Delta(v)-1}, u_{d_\Delta(v)}\} \in E_\Delta$.

If node $v \notin V_H$, then let $D^v_{\Delta 1} = (E^v_{\Delta 1}, \Gamma(v))$ be a constrained Delaunay triangulation of $\Gamma(V)$ such that $\{w, z\} \in E^v_{\Delta 1}$ iff $\{w, z\}$ is an edge of the Delaunay triangulation of $\Gamma(V)$ and does not have any part exterior to polygon P_v defined by nodes $u_1, \ldots u_{d_\Delta(v)}, u_1$.

If node $v \in V_H$, then let $D^v_{\Delta 1} = (E^v_{\Delta 1}, \Gamma(v))$ be a constrained Delaunay triangulation of $\Gamma(V)$ such that $\{w, z\} \in E^v_{\Delta 1}$ iff $\{w, z\}$ is an edge of the Delaunay triangulation of $\Gamma(V)$ and does not have any part exterior to polygon P_v defined by nodes $u_1, \ldots u_{d_\Delta(v)}, v, u_1$.

Claim 7. *$C^v_{1,\{w,z\}}$ is not empty iff $\{w, z\} \in E^v_{\Delta 1}$.*

Proof. First we prove that if $\{w, z\} \in E^v_{\Delta 1}$, then $C^v_{1,\{w,z\}}$ is not empty.

If $v \in V_H$ and $\{w, z\} \in E_{\Delta 1}$ or $v \in V \setminus V_H$ and $\{w, z\}$ lies on polygon P_v, then $\{w, z\} \in E_\Delta$, and v is a common neighbour of w and z in D_Δ. This means that the circle c_0 going through w, z, v does not contain any node in interior. We construct c_1 by increasing infinitesimal the radius of c_0 while keeping w and z on the boundary. Since the nodes are in general position, there is only v in interior of c_1, and thus $c_1 \in C^v_{1,\{w,z\}}$.

If $v \in V \setminus V_H$ and $\{w,z\} \in E^v_{\Delta 1}$, and is in interior of polygon P_v, consider the followings. Since $\{w,z\} \in E^v_{\Delta 1} \setminus E_\Delta$, there exist circles in $C_{1,\{w,z\}}$, which cover v, therefore $C^v_{1,\{w,z\}}$ is not empty.

Now we prove that if $C^v_{1,\{w,z\}}$ is not empty, then $\{w,z\} \in E^v_{\Delta 1}$.

Assume indirectly that there exists a node pair $\{w,z\} \in E_{\Delta 1} \setminus E^v_{\Delta 1}$ for which set $C^v_{1,\{w,z\}}$ is not empty. Take a disk $c_1 \in C^v_{1,\{w,z\}}$. We denote the centre point of c_1 by p_1. Let c_2 be the disk with centre point on line $p_1 v$ and having v and w on the boundary. Since $c_1 \in C^v_{1,\{w,z\}}$, $c_2 \in C_{0,\{v,w\}}$, which means $\{v,w\} \in E_\Delta$. Similarly, $\{v,z\} \in E_\Delta$.

This means that both w and z are element of $\Gamma(v)$. We can assume w.l.o.g. that $w = u_k$ for some $k \in \mathbb{N}$ and there exists a $k < l \in \mathbb{N}$ for which $z = u_l$.

For a line ab and point c denote the half plane bounded by ab and containing c by p^c_{ab}. Since $\{u_k, u_l\} \subset \Gamma(v)$ and $\{u_k, u_l\} \notin E^v_{\Delta 1}$, $l - k \geq 2$ and $\{u_{k+1}, \ldots, u_{l-1}\} \subset p^v_{u_k,u_l}$. This implies $u_{k+1} \in p^w_{u_k,u_l}$ and $u_{k+1} \in p^z_{u_k,v}$, because $\{\{v,u_l\},\{u_k,v\},\{u_k,u_{k+1}\}\} \in E_\Delta$ and D_Δ has no edge crossings. This way $u_{k+1} \in int(\Delta_{u_k v u_l})$, which implies that if a $c \in C_{1,\{u_k,u_l\}}$ contains v, it contains u_{k+1} too, which is contradiction.

□

Corollary 8. $E_{\Delta 1} = \cup_{v \in V} E^v_{\Delta 1}$.

Proof. By Claim 7, for every $v \in V$, $E^v_{\Delta 1}$ contains exactly those $\{w,z\}$ node pairs, for which $C^v_{1,\{w,z\}}$ is not empty. This implies $E_{\Delta 1} = \cup_{v \in V} E^v_{\Delta 1}$. □

Lemma 9. *$E_{\Delta 1}$ can be constructed in $O(n \log n)$ time.*

Proof. The edge set E_Δ can be constructed in $O(n \log n)$ [14]. Sets $\Gamma_{E_\Delta}(v)$ of neighbours can be computed in $O(n)$ for all $v \in V$.

Since constrained Delaunay triangulation also can be computed in $O(n \log n)$ time [14], present complexity is $O(\sum_{v \in V} |\Gamma_{E_\Delta}(v)| \log |\Gamma_{E_\Delta}(v)|)$, which is $O(n \log n)$, since

$$\sum_{v \in V} |\Gamma_{E_\Delta}(v)| \log |\Gamma_{E_\Delta}(v)| \leq \sum_{v \in V} d_{E_\Delta}(v) \log n \leq (3n-6) \log n.$$

It can be shown that an $e \in E^v_{\Delta 1}$ is in E_Δ iff e is part of P_v, thus $E_{\Delta 1}$ can be constructed by adding to E_Δ the edges of $E^v_{\Delta 1}$ interior to P_v for all $v \in V$. Finally the duplications have to be eliminated. This can be done in $O(n \log n)$.

Thus the overall complexity is $O(n \log n)$.

□

Let h denote the number of nodes in V on the convex hull of V.

Lemma 10. $|E_{\Delta 1}| \leq 6n - 2h - 9$.

Proof. First observe that $|E^v_{\Delta 1}| \leq 2d_\Delta(v) - 3$, since it can be proven by induction that every maximal triangulation on $d_\Delta(v)$ nodes with every node bounded by the infinite region has $2d_\Delta(v) - 3$ edges.

On the other hand, $|E_\Delta| = 3n - 3 - h$, and thus $\sum_{v \in V} d_\Delta(v) = 2|E_\Delta| = 6n - 2h - 6$. Finally by Corollary 8 during summing up the $|E^v_{\Delta 1}|$ values, there are counted only the edges of $E_{\Delta 1}$, every edge once or twice. The edges of E_Δ not on the convex hull of V are counted twice. Therefore

$$|E_{\Delta 1}| \leq \sum_{v \in V} |E^v_{\Delta 1}| - (|E_\Delta| - h) \leq$$

$$\leq \sum_{v \in V} (2d_\Delta(v) - 3) - (3n - 3 - 2h) =$$

$$= 2 \sum_{v \in V} d_\Delta(v) - 3n - 3n + 3 + 2h =$$

$$= 2(6n - 2h - 6) - 6n + 3 + 2h = 6n - 2h - 9.$$

□

Corollary 11. *$|E_{\Delta 1}| \leq 2|E_\Delta| - 3$, and thus $D_{\Delta 1}$ has less than 2 times as many edges as many a maximal planar graph has on point set V.*

Proof. Since D_Δ is a planar graph on the point set V, it is enough to prove $|E_{\Delta 1}| \leq 2|E_\Delta| - 3$, which is true by Lemma 10. □

IV. The Algorithm

Before presenting the algorithm itself, the main idea beyond will be highlighted. As seen before, all the failures from $\mathcal{M}_{1,\{u,v\}}$ can be covered with a disk from $C_{1,\{u,v\}}$ for some $\{u,v\} \in E_{\Delta 1}$, thus it is enough to examine the maximal edge sets covered by disks in $\cup_{\{u,v\} \in E_{\Delta 1}} C_{1,\{u,v\}}$. This gives the idea to generate the 'locally maximal' covered edge sets $\mathcal{M}_{1,\{u,v\}}$ for all $\{u,v\} \in E_{\Delta 1}$, then merge them in $\mathcal{M}_1$ by eliminating globally non-maximal elements.

The challenging part is determining sets $\mathcal{M}_{1,\{u,v\}}$. For solving the problem, consider the followings. For simplicity, assume for a moment that for an edge $\{u,v\} \in E_{\Delta 1}$ there exist at least 2 nodes of V both on the left and right side of line uv. Clearly, $C_{1,\{u,v\}}$ has a c_1 leftmost and a c_2 rightmost element. Circle c_1 is transformed imaginary continuously into c_2 while keeping the shape of disk and keeping $\{u,v\}$ on the boundary. The key observation that during this procedure edges from the left side are getting out from the covered area, while edges from the right side are getting in. Based on this idea one can determine sets $\mathcal{M}_{1,\{u,v\}}$ in polynomial time. But first some additional definitions have to be made.

Consider graph $D_{\Delta 1} = (V, E_{\Delta 1})$. By definition, $C_{1,\{u,v\}}$ is not empty for every $\{u,v\} \in E_{\Delta 1}$. If the set of nodes V_1 left from line uv has at least 2 elements, then let w_1 be the node with the second highest value of angle $\angle(uwv)$, where $w \in V_1$; if $|V_1| \leq 1$, let w_1 be a point infinitely far leftwards from line uv. We determine w_2 similarly on the right side of line uv.

Let $c^1_{u,v}$ and $c^2_{u,v}$ be the disks with u, v, w_1 and u, v, w_2 on the boundary. It is easy to see that $c^1_{u,v}$ and $c^2_{u,v}$ are the leftmost and rightmost disks from line uv in $C_{1,\{u,v\}}$, respectively. This way all disks in $C_{1,\{u,v\}}$ are covered by $c^1_{uv} \cup c^2_{u,v}$.

Let z_i be the node in interior of $c^i_{u,v}$ if there exist any, else $z_i \leftarrow w_i$ for $i \in \{1,2\}$.

Let $e \in E^3_{u,v}$ iff $e \in E$ intersects $c^1_{u,v} \cap c^2_{u,v}$, $e \in E^1_{u,v}$ iff $e \in E \setminus E^3_{u,v}$ intersects $c^1_{u,v} \setminus c^2_{u,v}$ and $e \in E^2_{u,v}$ iff $e \in E \setminus E^3_{u,v}$ intersects $c^2_{u,v} \setminus c^1_{u,v}$ (see Figure 3).

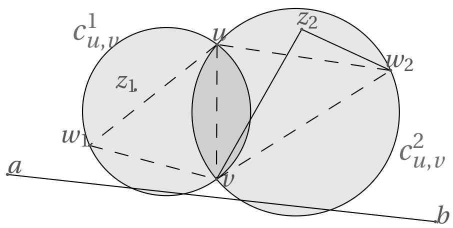

Fig. 3: Example on a Delaunay-one-edge $\{u,v\}$ with w_1, w_2, $C^1_{u,v}$ and $C^2_{u,v}$. Here $E^1_{u,v} = \{\{a,b\}\}$, $E^2_{u,v} = \{\{a,b\},\{z_2,w_2\}\}$, $E^3_{u,v} = \{\{z_2,v\}\}$.

Using the previously presented plan, Algorithm 1 computes $\mathcal{M}_1$ in the following way. First it generates the Delaunay-1 graph $D_{\Delta 1} = (E_{\Delta 1}, V)$ (see Algorithm 2). After that for every $\{u,v\} \in E_{\Delta 1}$ it generates sets $\mathcal{M}_{1,\{u,v\}}$ (Algorithms 3 and 4). Finally it computes $\mathcal{M}_1$ by gathering the globally maximal elements of sets $\mathcal{M}_{1,\{u,v\}}$ (Algoithm 5).

We proved the next theorem.

Theorem 12. *$\mathcal{M}_1$ can be computed in $O(n^2\theta_1^3)$) using Algorithm 1, and has $O(n\theta_1)$ elements, each of them consisting of $O(\theta_1)$ edges.*

Proof. At line 1 the set $E_{\Delta 1}$ is computed in $O(n \log n)$ time.

In line 3 we calculate sets $E^i_{u,v}$ for an edge $\{u,v\} \in E_{\Delta 1}$ in $O(n\theta_1)$ time according to Lemma 17. In line 4 we generate sets $\mathcal{M}_{1,\{u,v\}}$ ($\{u,v\} \in E_{\Delta 1}$), each in $O(\theta_1^2)$ time (Lemma 18). For all edges these give a running time of $O(n\theta_1(n+\theta_1))$ for lines 2-4.

Finally, in line 5 $\mathcal{M}_1$ is calculated from lists $\mathcal{M}_{1,\{u,v\}}$ in $O(n^2\theta_1^3)$ time (Lemma 21). □

Corollary 13. *Assuming θ_1 is upper bounded by a constant, $\mathcal{M}_1$ can be computed in $O(n^2)$ time, and its total length is $O(n)$.*

A. On the number of edges of G and size of θ_1

Lemma 14. *The number of edges $|E|$ is $O(n\theta_0)$, more precisely $m \le (2n-5)\theta_0$.*

Proof. The Delaunay triangulation is a planar graph and thus $|E_\Delta| \le 3n-6$. Since every Delunay triangle has 3 Delaunay edges, a Delaunay edge is the edge of at most 2 Delaunay triangles, and there are at least 3 Delaunay edges on the convex hull of V, the number of Delaunay triangles is at most

$$\frac{2|E_\Delta|-3}{3} \le \frac{2}{3}(3n-6)-1 = 2n-5.$$

Since every $c \in C_0$ intersects at most θ_0 edges of the network, and contains a Delaunay triangle and every edge intersects at least one triangle, the number of edges $|E|$ cannot be larger than θ_0 times the number of Delaunay triangles. We get $m \le (2n-5)\theta_0$. □

Algorithm 1: Main algorithm: Generating the maximal regional link single node failures

---

Input: $G = (V,E)$, (v_x, v_y) for all $v \in V$
Output: Set $\mathcal{M}_1$ of maximal regional link single node failures.
begin
1 $E_{\Delta 1} \leftarrow$1-DELAUNAY(V);
2 **for** $\{u,v\} \in E_{\Delta 1}$ **do**
3 $E^1_{u,v}, E^2_{u,v}, E^3_{u,v}, w_1, w_2, z_1, z_2 \leftarrow$ GETEDGESETS$(G, \{u,v\})$;
4 $\mathcal{M}_{1,\{u,v\}} \leftarrow$ GENERATE $(G, \{u,v\}, E^1_{u,v}, E^2_{u,v}, E^3_{u,v}, w_1, w_2, z_1, z_2)$;
5 $\mathcal{M}_1 \leftarrow$ ELIMINATEREDUNTANTS$(\mathcal{M}_{1,\{u,v\}}, \forall \{u,v\} \in E_{\Delta 1})$;
6 **return** $\mathcal{M}_1$

Combining the lemma with Claim 5 we get the following.

Corollary 15. *The number of edges in G is $O(n\theta_1)$.* □

A graph family may have $O(n^3)$ maximal link single node failures. However, we are convinced that θ_1 is small in case of typical backbone networks and there exists a small constant c that it never exceeds, thus $|\mathcal{M}_1| \le cn$.

B. Algorithm 2 (Method Delaunay-1)

Lemma 16. *Algorithm 2 computes $E_{\Delta 1}$ in $O(n \log n)$ time from node set V.*

Proof. In Algorithm 2 we realise the plan described in Lemma 9, thus the overall complexity of the algorithm is $O(n \log n)$. □

Algorithm 2: DELAUNAY-1: Finding set $E_{\Delta 1}$ of edges $\{u,v\}$ with $C_{1,\{u,v\}} \neq \emptyset$

Input: V
Output: $E_{\Delta 1}$.
begin
$E_\Delta \leftarrow$DELAUNAYTRIANGULATION(V);
for $v \in V$ **do**
$E^v_{\Delta 1} \leftarrow$CONSTRAINEDDELAUNAY$(\Gamma_{E_\Delta}(v))$;
$E_{\Delta 1} \leftarrow \cup_{v \in V} E^v_{\Delta 1}$;
return $E_{\Delta 1}$

C. Algorithm 3 (Method Getedgesets)

Lemma 17. *Algorithm 3 runs in $O(n\theta_1)$ time.*

Proof. For a fixed edge $\{u,v\} \in E_{\Delta 1}$, points w_1, w_2, z_1 and z_1 can be determined in $O(n)$ time, after that $c^1_{u,v}$ and $c^2_{u,v}$ can be determined in constant time. The proof is completed by applying Corollary 15.

□

Algorithm 3: GETEDGESETS: Finding $E^1_{u,v}$,$E^2_{u,v}$,$E^3_{u,v}$, w_1, w_2, z_1, z_2

Input: $G, \{u,v\} \in E_{\Delta 1}$
Output: $E^1_{u,v}, E^2_{u,v}, E^3_{u,v}, w_1, w_2, z_1, z_2$
begin
  Determine $w_1, w_2, z_1, z_2, c^1_{u,v}$ and $c^2_{u,v}$;
  $E^1_{u,v} \leftarrow \emptyset$;
  $E^2_{u,v} \leftarrow \emptyset$;
  for $\{a,b\} \in E$ **do**
    for $i \in \{1,2\}$ **do**
      if $\{a,b\} \cap c^i_{u,v} \neq \emptyset$ **then**
        $E^i_{u,v} \leftarrow E^i_{u,v} \cup \{a,b\}$
  $E^3_{u,v} \leftarrow E^1_{u,v} \cap E^2_{u,v}$;
  $E^1_{u,v} \leftarrow E^1_{u,v} \setminus E^3_{u,v}$;
  $E^2_{u,v} \leftarrow E^2_{u,v} \setminus E^3_{u,v}$;
  return $E^1_{u,v}$, $E^2_{u,v}$, $E^3_{u,v}$, w_1, w_2, z_1, z_2

D. Algorithm 4 (Method Generate $\mathscr{M}_{1,\{u,v\}}$)

Lemma 18. *Algorithm 4 generates set $\mathscr{M}_{1,\{u,v\}}$ in $O(\theta_1^2)$ time for any given $\{u,v\} \in E_{\Delta 1}$.*

Proof. $\mathscr{M}_{1,\{u,v\}}$ is generated by imaginary transforming continuously the leftmost disk from $C_{1,\{u,v\}}$ into the rightmost disk from $C_{1,\{u,v\}}$ through disks having u and v on the boundary. The number n_z of covered nodes by the disks is not constant throughout this process (its value can be both 0 and 2 too), thus n_z is tracked. Using Proposition 19 it van be shown that Algorithm 4 is correct.

We assume lines 3, 4 and 6 run in constant time. This means that for a given $\{u,v\} \in E_{\Delta 1}$ in lines 5 and 6 we get values $m^i_{a,b}$ in $O(\theta_1)$ time, in lines 7 and 8 we sort them in $O(\theta_1 \log \theta_1)$ time, line 9 runs in constant time, and in lines 10-22 $\mathscr{M}_1^{u,v}$ is calculated in $O(\theta_1^2)$ time by listing edge sets with locally maximal cardinality covered by disks while transforming $c^1_{u,v}$ into $c^2_{u,v}$, then eliminating non-maximal elements of the resulting list. Note that before eliminating sets the cardinality of which is not locally maximal might be listed too.

Thus the overall complexity of the algorithm is $O(\theta_1^2)$. □

Proposition 19. *It can be shown that if a $c \in C_1^{u,v}$ does not contain $L^1[i-1]$ but contains $L^1[i]$, then $L^1[j]$ is covered by c iff $j \geq i$. Similarly, if c contains $L^2[i-1]$ but does not contain $L^2[i]$, then $L^2[j]$ is covered by c iff $j \leq i-1$. Trivially, E^3 is covered by c, and that is also clear that for any $e_1 \in E^1$ and $e_2 \in E^2$ the edges e_1, e_2 are covered by c iff $m^1_{e_1} + m^2_{e_2} \leq \pi$.* □

Corollary 20. *For every $\{u,v\} \in E_{\Delta 1}$, $|\mathscr{M}_{1,\{u,v\}}|$ is $O(\theta_1)$.*

Algorithm 4: Generaring set $\mathscr{M}_{1,\{u,v|}$

Input: $G, \{u,v\} \in E_{\Delta 1}, E^i_{u,v} (i \in \{1,2,3\}), w_1, w_2, z_1, z_2$
Output: Set $\mathscr{M}_{1,\{u,v\}}$
begin
1   $P^1 \leftarrow$ half plane bounded by line uv, containing w_1 ("the half plane on left hand side");
2   $P^2 \leftarrow$ half plane bounded by line uv, containing w_2 ("the half plane on right hand side");
3   $m_{z_1} \leftarrow \angle u z_1 v$;
4   $m_{z_2} \leftarrow \angle u z_2 v$;
5   **for** $i \in \{1,2\}$ *and* $\{a,b\} \in E^i_{u,v}$ **do**
6     $m^i_{a,b} \leftarrow \max\{\angle ucv | c \in [a,b] \cap P^i\}$
7   $L^1 \leftarrow$SORT($E^1 \cup \{z_1\}$ by $m^1_{a,b}$ values and m_{z_1} increasingly);
8   $L^2 \leftarrow$SORT($E^2 \cup \{z_2\}$ by $m^2_{a,b}$ values and m_{z_2} decreasingly);
9   $\mathscr{M}_{1,\{u,v\}} \leftarrow \emptyset$, $i \leftarrow 1, j \leftarrow 1$, $n_z \leftarrow 1$;
10 **repeat**
11   **while** $m^1_{L^1(i)} + m^2_{L^2(j)} \leq \pi$ *and* $j \leq$*length*(L^2) **do**
12      $j \leftarrow j+1$;
13   **if** $L^2(j)$ *isNode* **then**
14        $n_z \leftarrow n_z + 1$;
          if $n_z = 1$ **then**
15   ADDCOVERED($M_{1,\{u,v\}}, L_1, L_2, L_3, i, j$)
      if $n_z = 1$ **then**
16   ADDCOVERED($M_{1,\{u,v\}}, L_1, L_2, L_3, i, j$) ;
17   **while** $m^1_{L^1(i)} + m^2_{L^2(j)} > \pi$ *and* $i \leq$*length*(L^2) **do**
18      $i \leftarrow i+1$;
19   **if** $L^1(i)$ *isNode* **then**
20        $n_z \leftarrow n_z - 1$;
          if $n_z = 1$ **then**
21   ADDCOVERED($M_{1,\{u,v\}}, L_1, L_2, L_3, i, j$)
    until $i >$*length*(L^1) *or* $j >$*length*(L^2);
22 Eliminate the non-maximal elements of $\mathscr{M}_{1,\{u,v\}}$;
23 **return** $\mathscr{M}_{1,\{u,v\}}$
24 **Function** *ADDCOVERED($\mathscr{M}_{1,\{u,v\}}, L_1, L_2, L_3, i, j$)*
25 $\mathscr{M}_{1,\{u,v\}} \leftarrow$ $\mathscr{M}_{1,\{u,v\}} \cup (\{L^1(i), \ldots, L^1(n), L^2(1), \ldots, L^2(j-1)\} \cup E^3) \cap E$

E. Algorithm 5 (Method Eliminateredundants)

Lemma 21. *Algorithm 5 computes $\mathscr{M}_1$ in $O(n^2\theta_1^3)$ using sets $\mathscr{M}_{1,\{u,v\}}$.*

Proof. The correctness proof of Algorithm 5 can be made by checking that it eliminates all globally non-maximal elements of $\mathscr{M}_{1,\{u,v\}}$ and leaves exactly one copy of each globally maximal element of $\mathscr{M}_{1,\{u,v\}}$ in the end.

There are $O(n^2\theta_1^2)$ times two sets of size $O(\theta_1)$ to be compared, thus the overall complexity of Algorithm 5 is $O(n^2\theta_1^3)$. If θ_1 is constant, this means $O(n^2)$.

Algorithm 5: ELIMINATEREDUNTANTS

Input: $\mathcal{M}_{1,\{u,v\}}$ for all $\{u,v\} \in E_{\Delta 1}$
Output: $\mathcal{M}_1$
begin
- **for** $\{u,v\} \in E_{\Delta 1}$ **do**
  - **for** $\{w,z\} \in E_{\Delta 1}, \{w,z\} \neq \{u,v\}$ **do**
    - **for** $f_{u,v} \in \mathcal{M}_{1,\{u,v\}}$ **do**
      - **for** $f_{w,z} \in \mathcal{M}_{1,\{w,z\}}$ **do**
        - **if** $f_{u,v} \supseteq f_{w,z}$ **then**
          - $\mathcal{M}_{1,\{w,z\}} \leftarrow \mathcal{M}_{1,\{w,z\}} \setminus f_{w,z}$
        - **else**
          - **if** $f_{u,v} \subset f_{w,z}$ **then**
            - $\mathcal{M}_{1,\{u,v\}} \leftarrow \mathcal{M}_{1,\{u,v\}} \setminus f_{u,v}$
- $\mathcal{M}_1 \leftarrow \bigcup_{\{u,v\} \in E_{\Delta 1}} \mathcal{M}_{1,\{u,v\}}$;
- **return** $\mathcal{M}_1$

□

V. Related Works

One approach to analyse the network vulnerability against regional failures is using probabilistic failure models, where each link in the SRLG has some probability to fail [1]. The probabilistic failure models can quantify the network protection schemes through evaluating their end-to-end connection availabilities. However often the end-to-end connection availability is not a sufficiently detailed modelling of the failure state, because it ignores the reconfiguration costs, possible traffic changes due to the failure, some limitations in the protocol and failure discovery mechanisms and physical impairments of the network. Therefore during network planing it is preferred to model the network behaviour of each possible failure state independently, and preconfigure the network for fast failure recovery of the failure state. The power of probabilistic failure models is that they implicitly treat a great number of failure states, but this on the other hand strongly limits the applicability of these models for network planing. Instead, the scope of this study is to define a small number of failure states (as SRLGs) to protect.

In [16] the circulant failure model is generalised and the regional failures are modelled by a given elementary geometric figures as ellipse, rectangle, square, or equilateral triangle with a predetermined size. The problem is similar to our one, and it is showed that there is a polynomial number of non-trivial positions for such a figure that need to be considered. In this study for simplicity we stick to failures of cycles only but with arbitrary size. Our main contribution compared to [16] is that this polynomial is basically linear in practice for disks covering a single node.

VI. Conclusions

In this paper we propose a fast and systematic approach to enumerate the list of possible failures deleting a node caused by natural disasters. Our approach assumes the disaster erases the network elements in a shape of a disk of any size which have exactly node interior. Although the number of possible areas affected by such disasters is infinite, we show that the generated list of failures is short, it is basically linear to the network size. We provide a fast polynomial time algorithm for enumerating the corresponding set of Shared Risk Link Groups.

References

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

[2] 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.

[3] 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.

[4] J. Heidemann, L. Quan, and Y. Pradkin, *A preliminary analysis of network outages during hurricane sandy*. University of Southern California, Information Sciences Institute, 2012.

[5] 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.

[6] I. B. 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.

[7] 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.

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

[9] 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.

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

[11] J. P. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, "Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines," *Computer Networks*, vol. 54, no. 8, pp. 1245–1265, 2010.

[12] S. Neumayer, A. Efrat, and E. Modiano, "Geographic max-flow and min-cut under a circular disk failure model," *Computer Networks*, vol. 77, pp. 117–127, 2015.

[13] B. Vass, "Shared risk link groups of disaster failures," in *IEEE Conference on Computer Communications Student Poster (INFOCOM Poster)*, 2016. [Online]. Available: http://lendulet.tmit.bme.hu/lendulet_website/wp-content/papercite-data/pdf/vass16infocomposter.pdf

[14] F. Aurenhammer, "Voronoi diagramsÑa survey of a fundamental geometric data structure," *ACM Computing Surveys (CSUR)*, vol. 23, no. 3, pp. 345–405, 1991.

[15] M. De Berg, M. Van Kreveld, M. Overmars, and O. C. Schwarzkopf, *Computational geometry*. Springer, 2000.

[16] 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.