BUDAPEST UNIVERSITY OF TECHNOLOGY AND ECONOMICS
DEPT. OF TELECOMMUNICATIONS AND MEDIA INFORMATICS

Modelling and Enumerating Geographically Correlated Failure Events in Communication Networks

Balázs Vass

Summary of the Ph.D. Dissertation

Supervised by

János Tapolcai
High Speed Networks Laboratory
Dept. of Telecommunications and Media Informatics
Budapest University of Technology and Economics

Budapest, Hungary
2020

1 Introduction

The Internet is a critical infrastructure. Due to the importance of telecommunication services, improving the preparedness of networks to regional failures is becoming a key issue [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]. The majority of severe network outages happen because of a disaster (such as an earthquake, hurricane, tsunami, tornado, etc.) taking down a lot of (or all) equipment in a given geographical area. Such failures are called *regional failures*. Many studies have touched the problem of how to prepare networks to survive regional failures, where the first solutions have assumed that fibres in the same duct or within 50 km of every network node fail simultaneously (namely, in a single regional failure) [13, 14]. These solutions were further improved by examining the historical data of different type of disasters (e.g., seismic hazard maps for earthquakes) and identifying the hotspots of the disasters [2, 5, 6, 8, 9, 11]. The weak point of these approaches is that, during network equipment deployment, many of the risks are considered and compensated (e.g., an earthquake-proof infrastructure in areas with larger seismic intensity), implying that the historical data does not represent the current deployments, and therefore, not the current risks. Thus, it may be more realistic to assume that any physically close-by equipment has a higher chance to fail simultaneously. More recent studies (including those forming the base of my Ph.D. theses) are purely devoted to this particular problem and adapt combinatorial geometric based approaches to capture all of the regional failures and represent them in a compact way [15, 16, 10, 17, C8, C9, C10], where the major challenge is that regional failures can have arbitrary locations, shapes, sizes, effects, etc. This collection of Ph.D. theses presents parts of the state of the art and suggests unified definitions, notions and terminology. For a more comprehensive survey, I kindly refer the reader to our book chapter [B1].

The output of the approaches discussed in the followings can serve as the input of the network design and management tools. Currently, network recovery mechanisms are implemented to protect a small set of pre-defined failure scenarios. Each recovery plan corresponds to the failure of some equipment. Informally speaking, when a link fails, the network has a ready-to-use plan on how to recover itself. Technically, a set of so-called *Shared Risk Link Groups (SRLGs)* are defined by the network operators, where each SRLG is a set of links whose joint failure the recovery mechanism should be prepared for. Service availability queries on the other hand require probabilistic refinements of the SRLG model, called collectively as *Probabilistic SRLGs (PSRLGs)*. More concretely, to evaluate the availability of network services, besides knowing the disaster frequencies, for every set S of links of the network, we need to store the probability that *exactly/at least* the links of S will fail simultaneously during the next disaster.

During my Ph.D. studies, I was purely focusing on how to define and enumerate succinct lists of SRLGs that cover all types of disasters. I have also addressed the issue of determining PSRLG lists that capture the correlated nature of link failures caused by a disaster and have tractable sizes.

2 Research objectives

The objective of this dissertation is creating the missing models and related algorithms translating the composed geometric problem of protecting telecommunication networks against regional failures to purely combinatorial and probabilistic problems, respectively. This translation has to provide a (probabilistic) list of possible failures as its output. The best technique to choose for enumerating the vulnerable regions of a network varies on (1) the available geometric information on the network topology, (2) (probabilistic) information on the effects of possible disasters in the network area, and (3) the desired output structure (SRLG/PSRLG). Clearly, there are some cases, when it is trivial what we can do as best as this translation. For example, when we have absolutely no information on the geographical embedding of our network topology, the best we can do as an enumeration of vulnerable regions is to list the $k = 1, 2, \ldots$ neighbourhoods of each network node or link as lists of SRLGs (cf. [B1]/1.3.2A). In

contrast to these cases, this dissertation tackles some non-trivial real-life problem versions discussed in the followings.

Theses 1 and 2 aim creating models and related algorithms providing succinct lists of SRLGs that cover all disasters in case of two different setting of quality of input data. In the first part of the dissertation (Thesis 1) it is assumed that (1) we have a precise knowledge on the geographical embedding of the network, but (2) only a maximum destruction radius is known about the nature of future disasters, and based on these information, (3) we want SRLG lists as output. My main goal here was to provide theoretical upper bounds on the number and computation time of SRLGs.

In Thesis 2 one is given only a schematic map of the network (what happens e.g., when the Internet Sevice Provider rents parts of its network from a Physical Infrastructure Provider), and wants a set of regional SRLGs. In this part of my dissertation the objective is creating a model which can handle this 'blury' representation of the topology, and provides a succinct list of SRLGs in polynomial time. I wanted to provide both theoretical and practical evaluations of this model.

In the last part of my dissertation (3), I desired to provide realistic PSRLGs. To be able to provide realistic PSRLGs, one needs (1) a precise geometric representation of the network, (2) detailed knowledge on the possible future disaster events, and lastly, a model representing these informations as precise yet succint lists of PSRLGs. In Thesis 3, I wanted to design a) standard PSRLG data structures, b) a state-of-the-art probabilistic failure model, and finally, c) an extensive evaluation of the model based on real-world network topology and disaster data.

3 Methodology

While modelling and elaboration of algorithms, I extensively used the tools of combinatorial optimisation and geometry. My results mainly rely on graph theory, computational geometry, complexity theory, and probability theory. While my results are primarily analytical, in some cases, the complexity of the problem demanded the extensive usage of *simulation* tools.

I implemented my proposed methods in Python 3.5, then I verified their effectiveness and pertinence through simulations.

4 New Results

When several network elements may fail together as a result of a single event, they are often characterized by *Shared Risk Group*s (SRGs). Each SRG has a corresponding failure event (or events); when such an event occurs, all elements in the SRG fail together. Specifically, the communication network is modelled as a graph $\mathcal{G} = (V, \mathcal{E})$, whose vertices are routers, PoPs[1], optical cross-connects (OXC), and users, while the edges are communication links (mostly optical fibres). SRGs are then defined as subgraphs $\langle V', E' \rangle$, where $V' \subseteq V$ and $E' \subseteq E'$.

In many cases, it is sufficient to consider only *links* in SRGs, and in this case, these groups are called *Shared Risk Link Groups* (SRLGs). For example, an SRLG may contain one edge (to capture a single-link failure) or all edges that touch one vertex (to capture a single-node failure). SRLGs may be more complex and represent simultaneous failures of multiple network elements. In particular, in this dissertation, we focus on geographically-correlated failures in which links within a specific region fail together.

A set $\mathcal{S}$ of SRLGs can be used as an input to network design and network recovery/protection mechanisms to ensure these mechanisms withstand the failures corresponding to these SRLGs. For example,

[1]A point of presence (PoP) is an artificial demarcation point or interface point between communicating entities.

to ensure connectivity between a specific pair of nodes, protection mechanisms may construct two edge-disjoint paths when $\mathcal{S} = \{\{e\}|e \in E\}$, two node-disjoint paths when $\mathcal{S} = \{\{(u,v) \in E\}|v \in V\}$, or two paths that do not traverse the same geographical region when $\mathcal{S}$ corresponds to all sets of links that are physically close-by.

The following definition captures the notion of SRLG introduced by regional failures, such as a natural disaster or an attack. For ease of presentation, we will call these failure events *disasters*, regardless of their cause.

Definition 4.1 (SRLG) *A set of links $S \subseteq E$ is an SRLG if we may assume there will be a disaster that can cause all edges in S to fail together. If the disaster can be characterized by a bounded geographical area in the two-dimensional plane $D \subset \mathbb{R}^2$, and S is the set of edges that intersect with D, then S is called the* regional SRLG that represents *D, and is denoted by $S = SRLG(D)$. If D is a circular disk, we call $SRLG(D)$ a* circular *SRLG.*

Circular SRLGs, which are the most common in literature, can also be characterized by the failure epicentre $p \in \mathbb{R}^2$ and the failure radius $r \in \mathbb{R}$. In this case $S = \{e \in E|d(e,p) \leq r\}$, where $d(e,p)$ is the Euclidean distance between edge e and point p.

4.1 Maximal SRLGs Induced by Disks with Radius r

Thesis 1 *[C6, C9, J2, J4] I proposed polynomial algorithms for enumerating lists M_r^p and M_r^s of maximal link sets (SRLGs) which can be hit by a disaster overestimated by a shape of a circular disk with an arbitrary given radius r, in case of embedding the network in the Euclidean plane and on the sphere, respectively. I gave theoretical upper bounds on the cardinality of both M_r^p and M_r^s. I proved that the proposed algorithm for planar embeddings has a computational complexity which is tight in the number of network nodes. Finally, compared the similarity of M_r^p in M_r^s in practice.*

Thesis 1.1 *[C9, J2] I proposed an algorithm, which, in case of representing a connected network topology $\mathcal{G}(V,\mathcal{E})$ in the Euclidean plane with links considered as line segments, computes the list M_r^p of maximal link sets hit by a circular disk with radius r in $O\left((|V|+x)\left(\log|V| + \phi_r^2\rho_r^5\right) + x'\log|V|\right)$, where x is the number of link crossings, ρ_r is the maximum number of links which are hit by a circular disk with radius r, ρ_r is the maximum number of links which is hit by a circular disk with radius r, ϕ_r is the maximum number of nodes in the $3r$-neighborhood of a link, and finally, x' is the number of link crossings if the links are elongated with $3\sqrt{2}r$ in both directions. I proved that the complexity of the proposed algorithm is tight in $|V|$. I proved that the cardinality of M_r^p is $O\left((|V|+x)\,\rho_r\right)$, and that this bound is tight. I proved that $\left|\bigcup_{0<r'<r} M_{r'}^p\right|$ is $O((n+x)\rho_r^2)$.*

I emphasize that although the same problem was investigated in my master thesis [T1], both the algorithm for enumerating M_r^p with its related theoretic worst case time bound and the theoretic upper bound on the cardinality $|M_r^p|$ were improved signifigantly. In addition, I proved that, in certain conditions, the complexity of the improved algorithm presented in [J2] is optimal in the number of network nodes $|V|$. In the followings, I present the main ideas behind these results.

In paper [C9], I presented a low-polynomial algorithm for computing M_r when links are considered as line segments (and the network is embedded in the plane). I showed that the number of elements of M_r is linear in the number of nodes in the network n, and its calculation can be done in a squared complexity of n (Theorem 6 of [C9]). Simulations indicate that this list has a size of $\approx 1.2n$ in practice.

To be more precise with the theoretical results, Corollary 4 of [C9] tells that the number of SRLGs in M_r is at most proportional to the product of 1) the number of nodes n plus the number of link intersections x, and 2) in the cardinality ρ_r of the biggest link set contained. The computing time needed

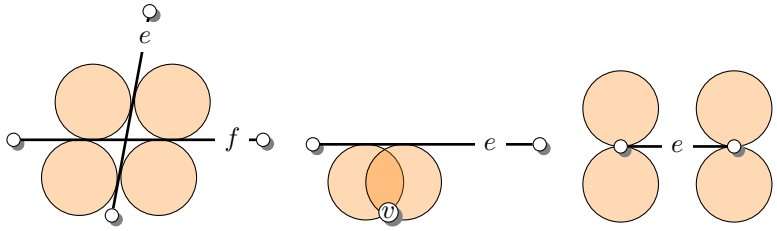

(a) $\forall e \in E$ and $\forall f \in E$ **(b)** $\forall e \in E$ and $\forall v \in V$ **(c)** $\forall e \in E$

Figure 1: The disk failures examined.

is $O\left((n+x)^2\rho_r^5\right)$ [C9, Thm. 6]. I note that x is 0 or a small number, and according to simulation results, ρ_r increases linearly with r, suggesting an $O(n^2r^5)$ runtime for $r > 0$.

Algorithm 1: Sketch of algorithm proposed in [C9]

Input: graph $G = \langle V, E\rangle$ embedded in plane, radius r
Output: List M_r of maximal SRLGs of disasters being circular disks with radius r
begin
1 $M'_r := \emptyset$
2 Calculate $X :=$ {`points of edge crossings`}
3 **for** $w \in V \cup X$ **do**
4 Determine $E_w :=$ {`edges not further from` w `than` $3r$}
5 **for** $e_1, e_2 \in E_w$ **do**
6 Calculate circles c_i described in Fig. 1/(a)
end
7 **for** $e \in E_w$ **do**
8 Calculate circles c_j described in Fig. 1/(b) with w as point
9 Calculate circles c_k described in Fig. 1/(c)
end
10 Refresh[2] M'_r with link sets hit by circles c_i, c_j, c_k (1 circle at a time)
end
11 **return** M'_r as M_r
end

In the followings, I give an overview of the proposed algorithm (Alg. 1), which relies on a series of geometric considerations. The most important one is Theorem 1 of [C9], which leverages that the link sets possibly hit by any of the infinite number of possible disaster locations can be determined via checking the effect of a quadratic number of disks on the network edges. In particular, for a positive real r, and a non-empty set of edges H which is hit by a circular disk of radius r, there exists a disk c of radius r which hits the edges of H such that at least one of the following holds (see Fig. 1 for illustrations): (a) There are two non-parallel links in H such that c intersects both of them in a single point. These two points are different. (b) There are two links 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 a link $e \in H$ at an endpoint of e.

Intuitively, there is no reason for checking for the circles described in Fig. 1 in case of two network elements which are much further apart than the disaster radius r. Indeed, according to [J2], one can build up the solution of the global problem based on some local calculations, as follows. Let X be the set of link intersection points. After determining X, one has to collect edges not further from w than $3r$ into a set E_w, for all $w \in V \cup X$, then determine the maximal failures of sets E_w, and finally, get the result by collecting the maximal elements of the resulting lists. This can be done in $O\left((|V|+x)\left(\log|V| + \phi_r^2\rho_r^5\right) + x'\log|V|\right)$ as proposed in Thesis 1.1.

I gave a lower bound on the computational complexity leveraging Lemma 4 of [18]. Namely, I proved that reporting that there are no intersecting line segments takes $\Omega(|V|\log|V|)$ or, in other words, that

[2]This means that M' is the set of maximal failures among which are already checked, and if f is maximal amongst them, it is added to M' and all f's subsets are eliminated from M'; or if f is not maximal in M', nothing happens.

computing M_r^p in the special case of $r = 0$ needs $\Omega(|V|\log|V|)$ time. This, together with the upper bound on time complexity means that the best time complexity in the number of nodes for calculating M_r^p is $\Theta(|V|\log|V|)\cdot f(x,\phi_r,\rho_r,x')$, where $f()$ is a function of the parameters.

The upper bound on the cardinality of M_r^p can be proved through some geometric considerations showing that there are at most $O((|V|+x)\rho_r)$ pairs of geometric objects which needed to be checked for possibly causing maximal failures in all three cases of Fig. 1. $\left|\bigcup_{0<r'<r} M_{r'}^p\right|$ results similarly.

Thesis 1.2 *[C6, J4] I proposed an algorithm, which, in case of representing a connected network topology $\mathcal{G}(V,\mathcal{E})$ on a sphere with links considered as chains of geodesics, computes list M_r^s of maximal link sets hit by a circular disk with radius r in polynomial time. I proved that the cardinality of M_r^s is $O\left(|\mathcal{E}|^3\right)$. Furthermore, $\left|\bigcup_{0<r'<\infty} M_{r'}^s\right|$ is also $O(|E|^3)$.*

In the followings, I sketch claims and an algorithm on calculating either M_r^s or M_r^p, thus everything holds in the case of the spherical representation of the networks. Lists M_r^s and M_r^p are collectively called M_r^g, where g stands for the geometry type p or s, standing for planar or spherical, respectively.

Let us make the following definitions for the sake of clarifying the intuition. (1) Let a disk c be *smaller* than disk c_0, if c has a smaller radius than c_0, or if they have an equal radius and the center point of c is lexicographically smaller than the center point of c_0. Among a set of circles S_c, let c be the smallest if it is smaller than any other circle in S_c. (2) Let $F \subseteq E$ be a finite nonempty set of edges (not necessarily a failure). We denote the smallest disk among the disks enclosing the polylines of F by c_F and we say c_F is the *smallest enclosing disk* of F.

It is not difficult to see that c_F always exists for line segments or geodesics (depending on the geometry), and thus, by mapping the corresponding segments/geodesics together we can deduct that the definition is correct for polylines too. The key idea of our approach is that we can limit our focus only on the smallest enclosing disks c_F.

The following observation is the key in bounding the cardinality of M_r^s. Let H be a nonempty set of polylines of edges with smallest enclosing disk c_H. Then there exists a subset $H_0 \subseteq H$ with $|H_0| \leq 3$ such that $c_H = c_{H_0}$. $|M_r^s| \leq \binom{|\mathcal{E}|}{3} + \binom{|\mathcal{E}|}{2} + |\mathcal{E}| = O(|\mathcal{E}|^3)$.

$|M_r^s|$ being $O(|\mathcal{E}^3|)$ means that if the smallest covering disk of at most three links can be calculated in polynomial time, one can design a polynomial algorithm for calculating M_r^s. A non-trivial lemma of

Algorithm 2: Determining maximal r-range SRLG lists

```
Input: G(V,E), r, geometry g, coordinates of nodes and
  edge polylines
Output: M_r^g
begin
    M_r^g := ∅
    Store E as a list,
    for i_1 ∈ {1,...,m} do
        for i_2 ∈ {i_2,m} do
            for i_3 ∈ {i_3,m} do
                c_{i_1,i_2,i_3} := c_{E[i_1],E[i_2],E[i_3]}
                if radius of c_{i_1,i_2,i_3} is ≤ r then
                    f := F(c_{i_1,i_2,i_3})
                    refresh M_r^g with f // as in
                        Algorithm 3
                end
            end
        end
    end
    return M_r^g
end
```

Algorithm 3: Refreshing SRLG list M with failure f

```
Input: SRLG list M, failure f
Output: M refreshed with f
begin
    maximal:=True
    for f_M ∈ M do
        if f ⊆ f_M then
            maximal:=False
        end
    end
    if maximal then
        M := M ∪ {f}
        for f_M ∈ M do
            if f ⊃ f_M then
                M := M \ {f_M}
            end
        end
    end
    return M
end
```

[C6] is the following. If H is a set of line segments in the plane or geodesics on the sphere with $|H| \leq 3$, then its smallest covering disk c_H can be determined in $O(1)$ time.

This means, that if H is a set of polylines of edges, with $|H| \leq 3$, then c_H can be determined in $O(\gamma^3)$ time. For this, one first have to unpack each polyline into the $\leq \gamma$ line segments/geodesics it is consisting of. Then, for each element h_i in H, pick a segment s_i. For each triplet (couple) of segments calculate the smallest enclosing disk (which can be done in $O(1)$), and lastly chose the smallest from among the resulting disks.

Based on these, Alg. 2 computes M_r^g in $O(|\mathcal{E}|^3(\gamma^3 + |\mathcal{E}|^4))$. In the following, I sketch the proof of the complexity. There are $O(|\mathcal{E}|^3)$ smallest enclosing disks to calculate, each in constant time. We claim that for each disk the calculation time of refreshing M_r^g with the resulting failure (as in Alg. 3) is $O(|\mathcal{E}|^4+\gamma^3)$ in case of each disk, because after the computation of the smallest enclosing disk in $O(\gamma^3)$ and determining f in $O(|\mathcal{E}|)$ there has to be done $O(|\mathcal{E}|^3)$ comparisons of link set, and each can be done in $O(|\mathcal{E}|)$.

Thesis 1.3 *[C6, J4] Through simulations, I showed that M_r^s and M_r^p can differ in practice, thus it is more precise to compute the SRLG lists with the spherical representation. However, in many of the cases, the distortion yielding from representing the network in the plane is causes less inaccuracy than the lack of knowledge on the disaster characteristics. In those cases, the planar representation can serve the purpose of vulnerable region detection well enough.*

An important question is that, in practice, under which geographic extension of the network can one say that, in the viewpoint of SRLG enumeration, it is practically indifferent whether we consider a spherical or a planar representation of the network. In other words, focusing now only on lists M_r, the question is that under which size of the physical network will M_r^p and M_r^s (maximal link sets which can be hit by a single circular disk with radius r, in the plane and on the sphere, resp.) be the precisely the same. The answer depends not only on the physical size, but also on the characteristics of the network itself: it can represent a dense metropolitan backbone network with multiple nodes close to each other, but it can also be geographically very sparse. Let τ be the distance of the closest non-adjacent and non-intersecting link of the network, and let D be the diameter of the smallest enclosing disk of the network $\mathcal{G}$. We can see that there can be any difference between M_r^s and M_r^p only if $2r \in [\tau, D]$ for the either the spherical or the planar representation. Practical radii of circular disasters range from couple of kilometers to couple of hundreds of kilometers, which mean, they might so small that there cannot be any difference between the SRLG sets (i.e. $2r < \tau$ means $M_r^s = M_r^p$). If τ is smaller than the disaster diameter, then it is easy to find settings, where $M_r^s \neq M_r^p$.

To study the phenomenon more in details, I used two similarity metrics of the SRLG lists: 1) the ratio of SRLGs, which are present in only one of M_r^p and M_r^s, i.e., $\mathcal{M}(r) := |M_r^p \triangle M_r^s|/(|M_r^p| + |M_r^s|) \in [0, 1]$, and 2) the average and maximal Hamming distance of an SRLG from M_r^s to its closest counterpart in M_r^p. I depicted the values of these metrics in Figures 2a-2c, respectively. As a base of evaluation, I took an Italian topology (Fig. 2d, with a diameter $D = 1180$km), and its magnified versions such that the resulting networks have diameters $D = 100, 200, \ldots, 1500$ km on the sphere. It can be seen that, in most cases, all of these metric values are 0 (i.e., $M_r^p = M_r^s$), but one can witness high spikes of big ratios of different SRLGs (Fig, 2a), or spherical SRLGs which have a symmetric difference of 3 links with their closest planar counterpart. This latter phenomenon happens when there are some nodes $u, v \in V$ such that $d(u, v) \leq 2r$ exactly in one of the spherical and planar representations.

The small and inconsistent ratio of different SRLGs in the two studied SRLG list is due to the fact that though the Earth surface is curved, this curvature is not practically significant in case of a backbone topology of a small to medium size country. For example, the maximum distance distortion of the Ortographic projection over Hungary and Italy (having diameters $<$ 530km and $<$ 1250km) is

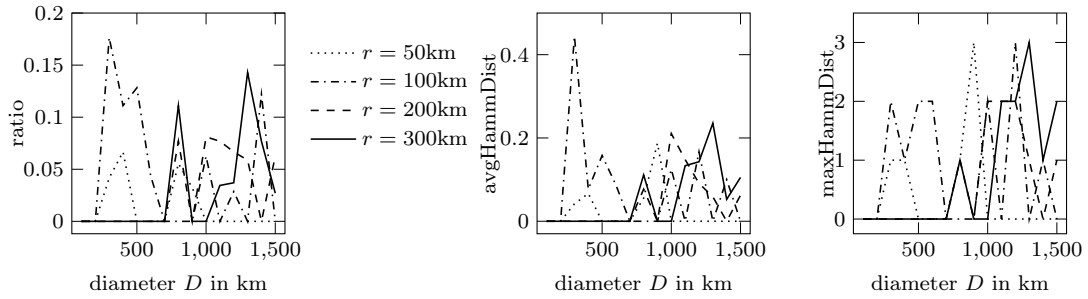

(a) Ratios of different SRLGs

(b) Average Hamming distance to closest SRLG

(c) Maximum Hamming distance to closest SRLG

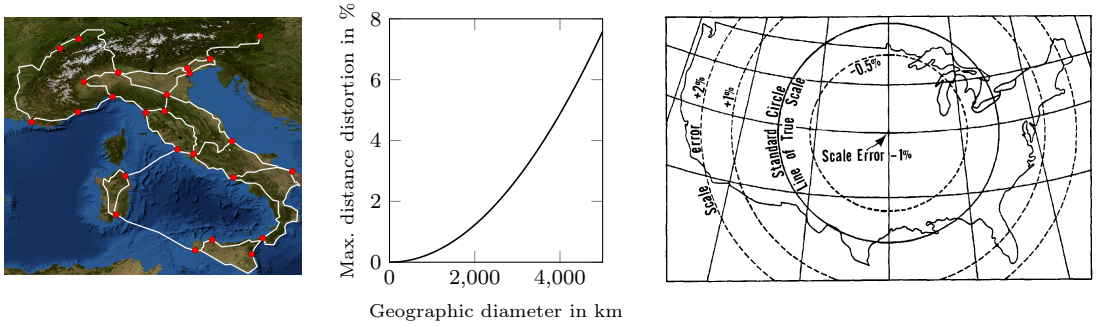

(d) The Italian topology (diameter $\simeq$ 1180km)

(e) Scale error of the Orthographic projection

(f) Scale error of the Oblique stereographic projection over the US is $\leq 4\%$ [19]

Figure 2: Caption

$< 0.1\%$ and $< 0.5\%$, respectively (Fig. 2e). Even the contiguous US can be mapped with $< 4\%$ of distance distortion (Fig. 2f [19]).

Since the calculation time of M_r^s was approximately twice of the M_r^p in my experience, I concluded as follows. M_r^s and M_r^p can differ, thus it makes sense to compute the SRLG lists with the more precise spherical representation. However, in many of the cases, the distortion yielding from representing the network in the plane is causes less inaccuracy than the lack of knowledge on the disaster characteristics (e.g., there can be as much as 10% inaccuracy in determining the disaster radius), thus the planar representation can serve the purpose of SRLG listing well enough.

4.2 Enumerating Maximal SRLGs Caused by Circular Disk Shaped Disasters Hitting k Nodes

The current best practice is to ensure that the primary and backup paths assigned to a connection are node disjoint. This way operators ensure that the distance between the nodes of the primary and backup paths (except at the terminal nodes) are in at least 1-hop-distance from each other. The root of the outages is usually because: (1) *close nodes* when two nodes are placed close to each other; for example, in highly populated areas. (2) *parallel links* when two links are placed close to each other because of some geographic reasons. Unfortunately, handling the geometric information with the network topology is not part of the current best practice. Furthermore, the Internet Service Providers usually hire the links as a service from an independent company, called the Physical Infrastructure Provider, and thus, operators have limited information about the route of the links, or the physical coordinates of the intermediate routing nodes. Thesis 2 proposes a way to ensure geographic distance between primary and backup paths.

Thesis 2 *[C12, C11, C10, J3] To ensure geographic distance between primary and backup paths when the geographical embedding the network topology is approximate, I proposed a model for enumerating regional*

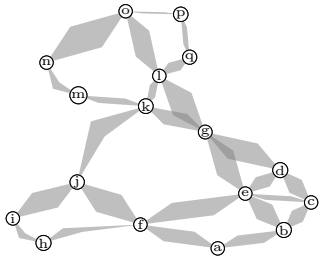

Figure 3: Input graph $\mathcal{G} = (V, \mathcal{E})$ with containing polygons, $|V| = 17$, maximum number of sides of the containing polygons $\gamma = 4$.

SRLGs relying only a schematic map of the network topology. For networks described in this model, I proposed a polynomial algorithm for enumerating list M_k of maximal link sets (SRLGs) which can be hit by a disaster overestimated by a shape of a circular disk hitting an arbitrary number k of nodes. I gave theoretical upper bounds on the cardinality of M_k. Evaluating the model and data structure, I showed that in case of real network topologies as input combined with practical (small) k values, M_k is a reasonably short list of link sets.

When defining the limited geometric information failure model presented in Thesis 2.1, I had the following design goals: 1) Do *not underestimate* the set of links involved in a possible regional failure, 2) Handle networks embedded in a *schematic map* with the exact route of the cables are unknown, 3) Allowing the design a *fast and space efficient* way of calculating the set of SRLGs (presented in Theses 2.2, 2.3).

Thesis 2.1 (The Limited Geometric Information Failure Model) *[J3] To ensure geographic distance between primary and backup paths when the geographical embedding the network topology is approximate, I proposed the following model. The (not necessarily planar) network is modelled as an undirected connected geometric graph $\mathcal{G} = (V, \mathcal{E})$ with $|V| \geq 3$ nodes. The nodes of the graph are embedded as points in the Euclidean plane, and their exact coordinates are considered to be known. In contrast to this, precise positions of edges are not known, instead, it is assumed that for each edge e there is a* containing polygon *(or simply* polygon*) e^p in the plane in which the edge lies. The disasters are assumed to have a shape of a circular disk with an arbitrary radius and centre position, but hitting at most k nodes for $k \in \{0, |V|-2\}$. The failures caused by these disasters are called* regional link k-node failures*.*

I argue this failure model can reasonably represent the possible regional failures, without actually requiring to know the scaling of the topology map.

Based on our output, operators can generate SRLG-disjoint primary and backup paths to protect the connection against natural disasters[3]. The distance between the primary and backup paths is a straightforward metric to compare the failure models. Based on the logical topology, the conventional approach to defining the distance is the hop-distance between the nodes traversed by the primary path and the nodes traversed by the backup path, except the terminal nodes. Based on this definition, we can list the failure models in increasing order of their strength [J3]: single link failures ($\geq$ 0-hop-distance[4], but link-disjoint routing), single node failures ($\geq$ 1-hop-distance), single regional link 0-node failures, single regional link 1-node failures, single regional link 2-node failures, etc.

I believe the proposed approach well captures the possible regional network failures based on the little geographic information available at network devices.

[3]The routing algorithms modify the SRLGs, whose failure isolates the source and destination nodes: those SRLGs are replaced with a smaller non-isolating SRLG according to the failure model.

[4]The minimum distance between the nodes of the working path and the nodes of the SRLG-disjoint protection path, except the terminal nodes.

Thesis 2.2 *[J3] I proposed an algorithm, which, in case of representing a network topology $\mathcal{G}(V,\mathcal{E})$ in the Euclidean plane with each link $e \in \mathcal{E}$ being part of a related polygonal region e^p having at most γ sides, computes the list M_k of maximal link sets which can be hit by a circular disk hitting at most k nodes in $O\left(|V|^2\left(\left(k^2+1\right)\rho_k^3+\rho_k\gamma+(k+1+\log(n\rho_0))\,\rho_0\gamma\right)\right)$, where ρ_k denotes the maximal number of links hit by a circular disk hitting at most k nodes. I proved that list M_k has $O\left(n\left(k+1\right)\rho_k\right)$ elements, this bound being tight in these parameters for $k=O(1)$.*

The proposed method is based on a set of (computational) geometric considerations. The key observation is that for any element of M_k there exists a circular disk-shaped disaster having k nodes in the interior which has (1) two nodes on its boundary, or else (2) only one node u on its boundary and having an infinite radius. This allows us to enumerate all possible maximal failures using a sweep surface method as follows.

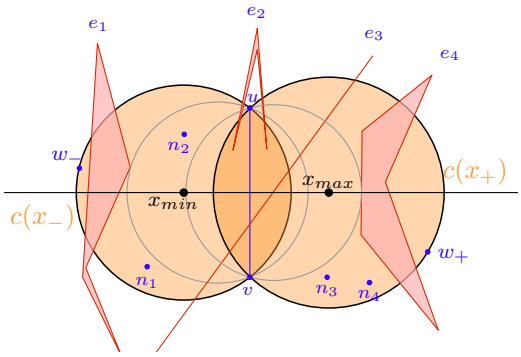

Figure 4: Illustration of an apple with $k=2$. Apple $A_k^{u,v}$ consists of specific ordered lists of links and nodes which can be hit by a disk from $\mathcal{C}_k^{u,v}$. For more details, please check [J3].

Let $\{u,v\} \subseteq V$ be two nodes for which the set $\mathcal{C}_k^{u,v}$ of circles which have k nodes in its interior and u and v on its boundary is not empty. These $\{u,v\}$ pairs are part of the set E_k of *k-Delaunay* edges, and their set can be determined in low-polynomial time [20, Thm. 2.4]. I defined a data structure *apple* $A_k^{u,v}$, which contains ordered lists of links and nodes which can be hit by a circle from $\mathcal{C}_k^{u,v}$. Suppose u and v are positioned as in Fig. 4. With the help of $A_k^{u,v}$, one can sweep through circles of $\mathcal{C}_k^{u,v}$ ordered by the abscissas of their centre points allowing to collect the set $M_k^{u,v}$ of maximal hit link sets by disks from $\mathcal{C}_k^{u,v}$.[5] Then the globally maximal elements of all lists $M_k^{u,v}$ are collected in M_k^2.

In the second case, the set of maximal failures M_1 from M_k for which exist a half-plane going through a node and hitting them can be calculated similarly via turning a half-plane around every node while

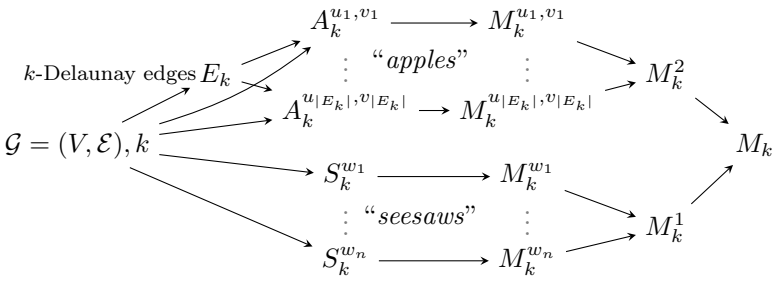

Figure 5: Sketch of algorithm from [J3] for enumerating set M_k of maximal link sets which can be hit by a circular disk hitting k nodes.

[5] The proposed method is similar to the sweep line (sweep surface) algorithms in computational geometry, e.g. Fortune's algorithm for computing the Voronoi diagram of a point set [21].

checking the set of hit links and the number of hit nodes. Finally, M_k can be obtained by collecting the maximal elements of M_k^1 and M_k^2. The process is sketched in Fig. 5. Unfortunately, discussing the technical details would exceed the limits of this thesis booklet.

In the following, I present numerical results that validate my model on some realistic physical networks. The network topologies with the obtained list of SRLGs for various k are available online[6].

Thesis 2.3 *In case of real network topologies, with their edges considered polyginal chains and line segments between their endpoints, respectively, list M_k of maximal link sets which can be hit by a circular disk hitting at most k nodes has $\approx 1.2 \cdot |V|$ and $\approx 2.2 \cdot |V|$ elements for $k = 0$ and $k = 1$, respectively. Additionally, $|M_k|$ increases sublinearly in function of k. Parameter ρ_k representing the maximal number of hit links by a disaster hitting k nodes was ≤ 10 for all networks for $k = 0, 1$, and grew to only to < 25 for $k = 5$. I concluded that list M_k has a reasonably small size for practical k values.*

As noted in the thesis, in the simulations, the input topologies are interpreted in two ways: 1) *polygon*, where links are polygonal chains, and 2) *line*, where the corner points of the polygonal links are substituted with nodes (of degree 2); here the shapes of links are line segments.

To have a better impression on the SRLGs on M_k, Fig. 6 visualises example results for both interpretations of the US ATT-L1 network, which has 126 nodes and 208 links as polygonal chains, and 162 nodes and 244 links after transforming it into a network of line segments. The 126-node US (ATT-L1) network was covered with 190 SRLGs, which is less than listing every single node and link as an SRLG. Fig. 6 shows these SRLGs, intuitively each corresponds to a mid-size regional failure. The SRLGs meet our intuition that there are more network nodes in the crowded areas, and thus it generates more SRLGs for them, while in the less crowded areas are covered with SRLGs corresponding to bigger areas.

In practice, it is important to have small SRLGs because it strongly influences the performance of the survivable routing algorithms. On Fig. 6 the SRLGs are relatively small, each SRLG contains a bit less than 3 links on average.

Based on Table 1, the number of SRLGs is roughly $\approx 1.2 \cdot |V|$ for $k = 0$, and $\approx 2.2 \cdot |V|$ for $k = 1$. Fig. 7a depicts this linear growth in n for $k = 0$. Fig. 7b shows the increase in the number of SRLGs for each network independently for the same range of k. Here we can experience a slightly sub-linear increase in $|M_k|$ in function of k. Fig. 7c shows that the edge density ρ_k increases linearly with k. Parameters $\rho_{k=0}$ and $\rho_{k=1}$ were less or equal to 10 in case of all investigated networks, and $\rho_{k=5}$ was also below 25. In practice, it is important to have small SRLGs because it strongly influences the performance of

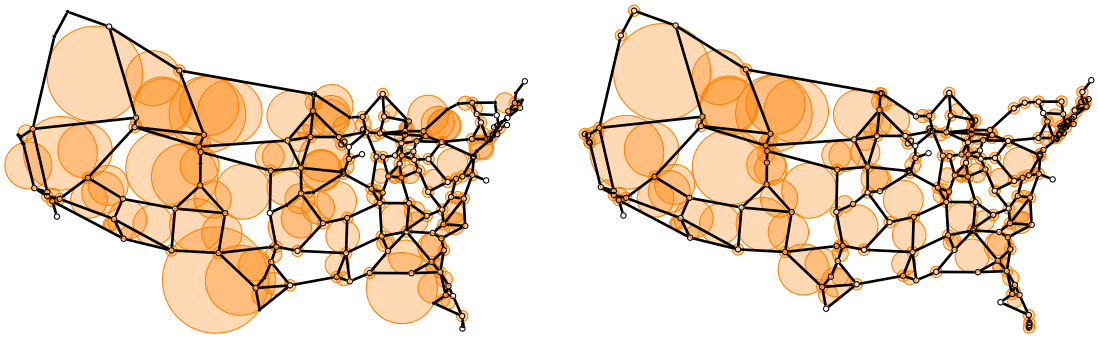

(a) *Polygon*: there are 190 SRLGs with average of 2.98 links and $\rho_0 = 5$.

(b) *Line*: there are 216 SRLGs with average of 2.79 links and $\rho_0 = 5$.

Figure 6: The SRLGs of $k = 0$ are visualized for the two cases (a) links are polygonal chains, and (b) the corner points of the polygonal line segments are treated as degree two nodes and all links are line segments. In order to have a perspicuous illustration, each SRLG is drawn with the smallest possible circular disk that covers all of its links, even if the disk has nodes interior.

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

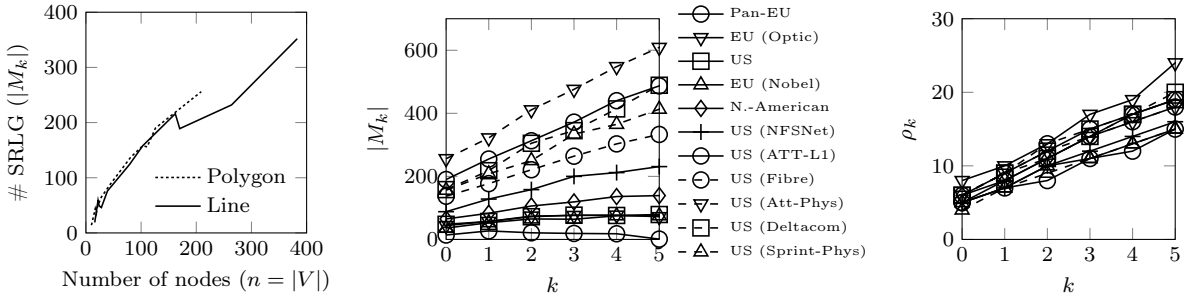

(a) The increase of the number of SRLGs with respect to the number of nodes, for $k = 0$.

(b) The number of SRLGs ($|M_k|$)

(c) The max. size ρ_k of SRLGs in M_k

Figure 7: Comparison of the two interpretation of the input topologies: links are polygonal line segments, or the corner points of the polygonal line segments are treated as degree two nodes and all links are line segments.

Name	$\|V\|$		$\|\mathcal{E}\|$		# SRLG $k=0$		# SRLG $k=1$	
	Polygon	Line	Polygon	Line	Polygon	Line	Polygon	Line
Pan-EU	10	16	16	22	14	19	27	35
EU (Optic)	17	22	40	45	44	59	57	71
EU (Nobel)	19	28	32	41	36	46	53	81
US [2]	21	24	39	42	48	49	57	64
N.-American	28	39	50	61	65	76	83	97
US (NFSNet)	44	79	73	108	88	124	128	172
US (Fibre)	81	170	141	230	137	189	177	249
US (Deltacom)	103	103	302	302	158	158	218	218
US (Sprint-Phys)	111	264	160	313	156	232	208	307
US (ATT-L1)	126	162	208	244	190	216	255	285
US (Att-Phys)	209	383	314	488	256	352	322	457

Table 1: Results of physical backbone topologies of [22].

the survivable routing algorithms. In our case, $\rho_{k=0}$ ranged between 4 and 8 depending on the network topology.

We can conclude that M_k has a reasonably small size for practical k values ($k <\approx 5$).

4.3 Unifying the Terminology on Probabilistic SRLGs and a Tractable Stochastic Model of Correlated Link Failures Caused by Disasters

The likelihood of a disaster to occur is not the same at all points of the plane. For example, earthquakes are more likely to occur in rupture zones than in other places, and regions with lower altitude are more likely to suffer from floods. Thus, the probability of an event to occur is important. Thesis 3 tackles this question by proposing standard terminology, containers and algorithms to speak about, store and compute Probabilistic SRLGs (PSRLGs).

Thesis 3 *[C3, C8, J1] I defined a stochastic model of link failures caused by disasters, which considers the correlation between failures of links which are geographically close to each other. To unify the notions and terminology on Probabilistic SRLGs, I proposed standard data structures for containing the disaster probabilities. In case of circular disk shaped disasters, for the size and query time of these data structures, I proposed theoretical upper bounds. Evaluating the model and data structures, I showed that in case of real seismic data as input, these data structures have a manageable size.*

Fig. 8 visualises Thesis 3. Namely, on the highest level, I propose standard data structures for storing failure probabilities of link sets. These data structures are universal, and thus independent of the model, which is used to fill them up. On the second level, I created a stochastic model of link failures caused by disasters, which grasps the correlation between the link failures caused by disasters more precisely

compared to former models. This model is a general one, as it can handle any kind of disaster events as input. On the third level, I evaluated my model using real seismic data. This way, the results summarised in Thesis 3 demonstrate how the raw disaster data alongside with the geographic embedding of the network topology can be translated to PSRLGs, i.e. simple link sets with an associated probability. PSRLGs then on their course, are the inputs for service availability queries, stochastic optimization problems, etc.

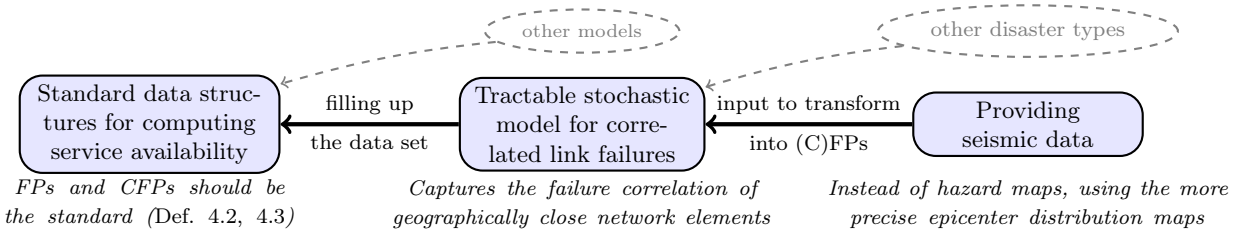

Figure 8: Components of Thesis 3: 1) standard data structures (CFPs and FPs) for storing joint failure probabilities of link sets, 2) a tractable stochasic model of link failures caused by disasters, and finally 3) evaluating using seismic data represented more precisely than the usual.

I divided Thesis 3 in three Subtheses:

Thesis 3.1 *[C3, C8, J1] Inspired by earthquake behaviours, I defined a stochastic model of link failures caused by disasters. This model is the first to explicitly consider the correlation between failures of links which can be subject to the same disaster. To unify the notions and terminology linked to probabilistic extensions of Shared Risk Link Groups, I proposed two standard data structures for containing the disaster probabilities, namely, the FP and the CFP: for a link set S, FP(S) and CFP(S) is the probability that* exactly*, and* at least *S will fail, respectively.*

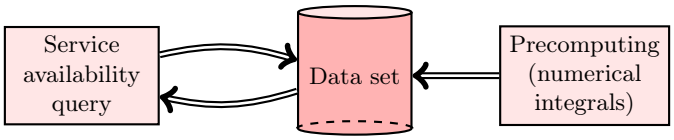

Figure 9: Framework to compute service availability

Standard PSRLG terminology Let us concentrate on the seismic hazard as an example for a moment. The hazard induced by earthquakes sometimes given in the form of an *epicentre distribution map*, which gives for each location $p \in \mathbb{R}^2$, the probability that a disaster happened with epicentre p. Moreover, the size (or radius) of the disaster can also be a random variable (e.g., earthquakes with a larger magnitude are less likely to happen than earthquakes with smaller magnitude, even if their epicentres are the same). Thus, it is customary to consider a set $\mathcal{D}$ of disasters $D \subseteq \mathbb{R}^2$ (that can be of infinite size), and attach a probabilistic measure to this set. For simplicity, let's assume that $\mathcal{D}$ is finite, and let $p_D = \Pr[\textit{disaster } D \in \mathcal{D} \text{ occurs}]$[7]. We note that an SRLG S can represent more than one disaster in $\mathcal{D}$; thus, we denote by the support$(S) = \{D \in \mathcal{D} | S = SRLG(D)\}$.

In the literature, there is a variety of non-equivalent definitions of PSRLGs [8], thus a unified terminology is desired. Definitions 4.2 and 4.3 propose standard PSRLG notions. An FP (Def. 4.2) tells the probability that the failed link set will be *exactly* S, while a CFP (Def. 4.3) tells the probability that *at least* S will fail:

[7]For infinite sets, one can use discretization and consider only finite number of sets, albeit with a small error.

[8]E.g. in the first paper considering probabilistic extensions SRLGs (which was [23]), each SRLG event $r \in R$ occurs with probability π_r, and once an SRLG event r occurs, link (i, j) will fail independently of the other links with probability $p^r_{i,j} \in [0, 1]$. Thus, we could call the [23]-PSRLGs as 'two stage PSRLGs'.

Definition 4.2 (FP) *Given a set $\mathcal{D}$ of disasters $D \subseteq \mathbb{R}^2$, a probability p_D for each disaster in $\mathcal{D}$, and a link set $S \subseteq E$, the* Link Failure State Probability (FP) *of S is $FP(S) = \sum_{D \in support(S)} p_D$. We note that if a disaster in support(S) actually occurs, then all links in S fail (with probability 1).*

Definition 4.3 (CFP) *Given a set $\mathcal{D}$ of disasters $D \subseteq \mathbb{R}^2$, a probability p_D for each disaster in $\mathcal{D}$, and a link set $S \subseteq E$, the* Cumulative Link Failure Probability (CFP) *of S is $CFP(S) = \sum_{T \supseteq S} \sum_{D \in support(T)} p_D$. We note that if a disaster in $\bigcup_{T \supseteq S} support(T)$ occurs, then all links in S fail (with probability 1).*

In a sense, FPs are like probability density functions (PDFs), while CFPs are like their cumulative distribution functions (CDFs). Collectively, we call FPs, CFPs as *Probabilistic SRLGs (PSRLGs).*

Stochastic model of correlated link failures caused by disasters First of all, my model for regional failures caused by a disaster assumes the following:

Assumption 1 *In the investigated time period, there will be at most one disaster.*[9]

The model has the following parameters with randomly chosen values: **epicenter** p, which is a point in the plane $\mathbb{R}^2$, **shape (and size)** s, which is a real value in $[0, 1]$. Each point $p \in \mathbb{R}^2$ is assigned with a **hazard** $h(p)$ representing the probability that p becomes the epicenter of the next disaster (see Fig. 10a). Specifically, $h(p)$ is a probability density function on the area $\mathbb{R}^2$, and therefore,

$$\int_{p \in \mathbb{R}^2} h(p) \mathrm{d}p = 1 \ . \tag{1}$$

After a regional disaster of the examined type (e.g. EMP attack, natural disasters, such as solar flares, earthquakes, hurricanes, and floods) the physical infrastructure (such as optical fibers, amplifiers, routers, and switches) in some area is destroyed. The possible shapes for this area are defined by a set $r(p, s)$ that represents a closed region on the plane (the actual shape of the destroyed area) as a function of epicenter p and the shape/size parameter s. This is a general disaster model, where several possible damage areas can be defined by $r(p, s)$ (see Fig. 10b).

Assumption 2 *A* regional disaster *of epicenter p and shape/size s will result in the failure of exactly those links of network G which have a point in $r(p, s)$.*

The model assumes that $r(p, s)$ is monotone increasing in s (see Fig. 10b for an example)[10], or more formally we assume that

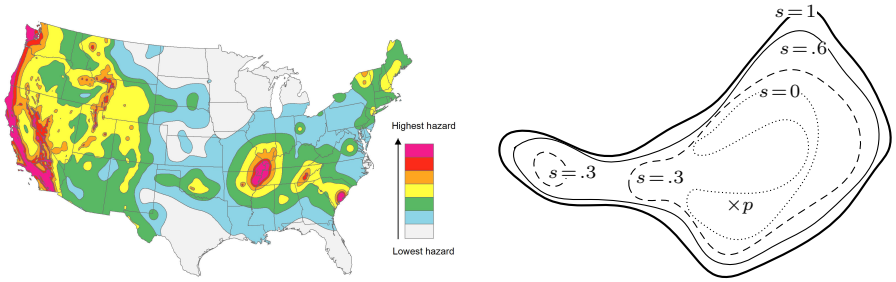

(a) Hazard map $h(p)$ for earthquakes as function of epicenter p. [25]

(b) Shape of regional disaster $r(p, s)$ for epicenter p and different sizes $s = 0, 0.3, 0.6, 1$.

Figure 10: Example of real-world inputs.

[9] The case when more disasters are expected to happen simultaneously can be handled by defining a new mixed disaster type, see also [24].

[10] Various failure shapes were studied so far [10, 16, 15, 26, 5, 7, 2, 6, 8, 9, 11, 27, 28, 29, 30, 31, C9], mainly in the form of circular regional disasters or line-segment failures, but in some cases also for arbitrary geometric objects [15, 26]. All of these models meet Assumption 3.

Assumption 3

$$r(p,s) \subseteq r(p,s') \text{ if } s < s' \quad \forall p \in \mathbb{R}^2, 0 \leq s, s' \leq 1 \ , \tag{2}$$

$r(p,s)$ for a given p is a result of uniform sampling of damage areas. Namely, for a given p the probability of the failure to be of size smaller than s is exactly s. Thus, s is called *relative size* in the remainder of the paper.

It is important to notice that given the disaster epicenter and relative size, the outcome of the attack is deterministic. In other words, any link e within $r(p,s)$ fails with probability 1, if a failure event with parameters p and s occurs. Let us denote the set of failed links by $R(p,s)$. Assumption 2 implies that, given a point p, $R(p,s) \subseteq R(p,s')$ if $s \leq s'$. Let $s(p,e)$ denote the corresponding smallest size s for which a failure at point p can cover link e. Furthermore, we denote by ρ the maximum number of links that can be affected by a single failure (of maximum size $s = 1$):

$$\rho = \max_{p \in \mathbb{R}^2} R(p,1) \ . \tag{3}$$

Let $f(e,p)$ be the **probability** that link e fails if a disaster with epicenter p happens. Note that $f(e,p) > 0$ can occur iff $e \in R(p,1)$. $f(e,p)$ can be computed from $R(p,s)$, where s is in the range $[0,1]$. Hence,

$$f(e,p) = \int_{s=0}^{1} I_{R(p,s)}(e) \mathrm{d}s \ , \tag{4}$$

where the indicator function $I_{R(p,s)}(e)$ indicates whether $e \in R(p,s)$.

We now extend our notation to capture the probability of the failure of link e in the next disaster:

$$P(e) := \int_{p \in \mathbb{R}^2} h(p) f(e,p) \mathrm{d}p. \tag{5}$$

We denote the probability that a set of links $S \subseteq E$ fail simultaneously, given that the disaster epicenter is $p \in \mathbb{R}^2$:

$$f(S,p) := \int_{s=0}^{1} \prod_{e \in S} I_{R(p,s)}(e) \mathrm{d}s \ . \tag{6}$$

In other words, if the sequence of links is $S = (e_1, e_2, \ldots, e_{|S|}) \subseteq R(p,1)$ and $s(p,e_1) \leq s(p,e_2) \leq \cdots \leq s(p,e_{|S|})$, then $\prod_{e \in S} I_{R(p,s)}(e) = 1$ iff $s \geq s(p,e_{|S|})$, otherwise the product is 0. This implies that

$$f(S,p) = f(e_{|S|},p) = \min_{e \in S} f(e,p) \ . \tag{7}$$

Finally, using the above results, we can compute CFP(S) as follows[11]:

$$\mathrm{CFP}(S) = \int_{p \in \mathbb{R}^2} h(p) f(S,p) \mathrm{d}p = \int_{p \in \mathbb{R}^2} h(p) \min_{e \in S} f(e,p) \mathrm{d}p \ . \tag{8}$$

For example, in Fig. 11, the results of applying the formula to the 5-node network are shown for all the non-zero joint link failure probabilities. In this example, $r(p,s)$ is always a circular disk of radius $s \cdot 50$km. Potentially there are exponentially many joint failure events in terms of the network size; however, links far from each other have zero probability to fail jointly because of a single disaster. This holds, for example, for links f and e, whose smallest distance is 200km.

Other works (e.g., [15, in the proof of Lemma 8]) expressed the joint failure probability of a set S by multiplying the failure probabilities of the links in S, thus implicitly assuming these failures are

[11] FP(S) can be computed based on similar observations.

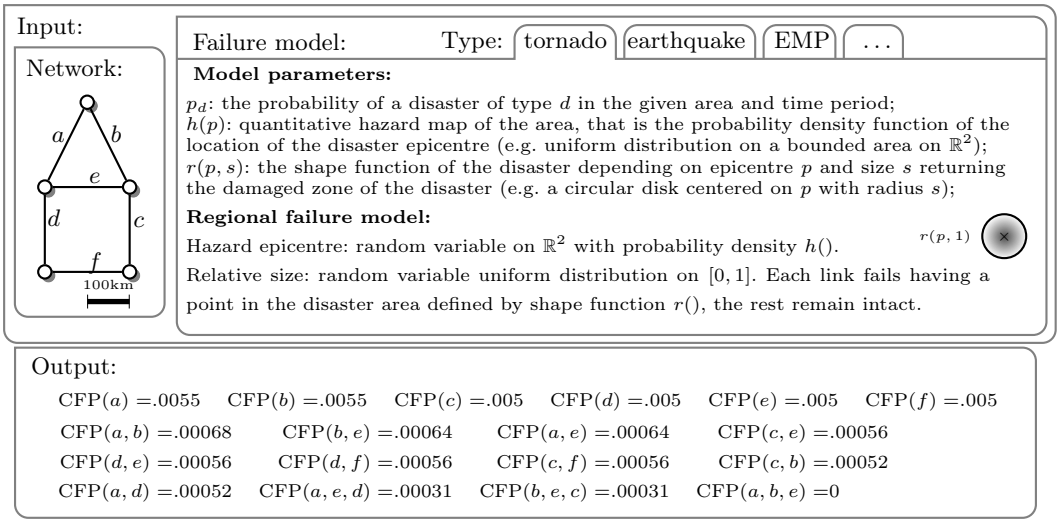

Figure 11: An illustration of the problem inputs and outputs.

independent. Unlike [15], my model assumes deterministic failure outcome (once its epicenter and shape are set). This implies that, in my model, failures are dependent. For example, two lines in the same location (e.g., within the same conduit) always fail together (e.g., when the conduit is cut).

Thesis 3.2 *[C3, C8, J1] In case of circular disk shaped disasters, representing the network topology $\mathcal{G}(V, \mathcal{E})$ in the Euclidean plane with links considered as polygonal chains consisting of at most γ line segments, denoting the number of link crossings by x, an the maximum number of links which are hit by one of the disasters by ρ_r, I proved the followings. There are $O((n + x)\rho^2\gamma^2)$ FPs with nonzero probability. The number of CFPs with nonzero probability is lower bounded by $\Omega(2^\rho)$ and upper bounded by $O(2^\rho(n + x)\rho^2\gamma^2)$. Storing all the nonzero CFPs in a balanced binary tree, the worst-case query time of the CFP of a given link set is $O(\rho \log((n + x)\rho\gamma))$. Storing all the nonzero FPs in a list, the query time of the CFP of a given link set is $O((n + x)\rho^2\gamma^2)$.*

Data set name	Space complexity	Query time
CFP	$\Omega(2^\rho)$ and $O(2^\rho(n + x)\rho^2\gamma^2)$	hashing: *constant* with high prob. balanced binary tree: $O(\rho \log((n + x)\rho\gamma))$ worst-case
FP	$O((n + x)\rho^2\gamma^2)$	$O((n + x)\rho^2\gamma^2)$

Table 2: Computing service availability: Trade-off between space complexity and query time

Cardinality of structures FP and CFP In case of the failure model presented before, the number of FPs can be nicely upper bounded as follows.

Proposition: In case of circular disk shaped disasters (i.e. $r(p, s)$ is circular), there are $O((n+x)\rho^2\gamma^2)$ FPs with nonzero probability.

Proof: Let us concentrate for line segment links for a moment. According to [C9, Claim 2], m is $O(n + x)$ for line segment links. We know from [32, Thm. 5] that the number of k-Voronoi cells for line segments is $O(k(m-k)+x)$, or alternatively, $O(k(n+x-k)+x)$ thus disasters hitting k links can hit at most this many link sets. Since a circular disk can hit at most ρ links, this sums up to $O(\rho^2(n+x-\rho)+x)$, which is $O(\rho^2(n + x))$.

If links can be polygonal chains consisting of at most γ line segments, there can be $O(k\gamma(n+x)\gamma)$ of k-Voronoi regions, yielding an upper bound of $O((n + x)\rho^2\gamma)^2)$ for the number of FPs needed. □

While CFP(S) can be queried directly to obtain the joint failure probability of a link-set, the number of CFPs needed to describe the stochastic effect of the next disaster can be very large:

Proposition: A lower and an upper bound on the number of CFPs with nonzero probabilities is $\Omega(2^\rho)$ and $O(2^\rho(n+x)\rho^2\gamma^2)$, respectively.

Proof: By the definition of ρ, there is a link set S with CFP(S) > 0 and $|S| = \rho$. As, for any $S' \subseteq S$, CFP(S) > 0 implies CFP(S') > 0, implying the lower bound. Speaking of the upper bound, by the former Proposition, there are at most $O((n+x)\rho^2\gamma^2)$ nonzero FPs, each having at most 2^ρ subsets, yielding the upper bound. □

Query time of structures FP and CFP It can be shown that, in my model, for any given link set $S \subseteq E$, CFP(S) $= \sum_{T \supseteq S}$ FP(T). Based this, in case of list FP and a link set S, to return CFP(S), one must go through the list and sum up the values FP(T) for all $T \supseteq S$. According to the former Propositions, this can be done in $O((n+x)\rho^2\gamma^2)$.

If structure CFP is stored as a list, the query time of CFP(S) is $\Omega(\rho)$. Note that if S is not stored in CFP, then CFP(S) $= 0$. The query time of sets can be reduced to a constant with very high probability (with the help of hashing). Using a balanced binary tree, its worst-case query time is $O(\rho \log((n+x)\rho\gamma))$, by the former Propositions, which is still very impressive. The drawback of structure CFP is that it has an $\Omega(2^\rho)$ space complexity, which makes it very inefficient for bigger network densities.

Before stating Thesis 3.3, which presents simulations based on seismic hazard, I spend some word on a standard scale for measureing the intensity of shaking produced by an earhquake, the Mercalli-Cancani-Sieberg (MCS) scale [33]. The scale ranges from I to XII: an intensity $I \leq$V does not cause structural damage, at $I =$IX, soil liquefaction starts, and finally, $I =$XII means total damage.

Thesis 3.3 *[C3, C8, J1] Using real-world seismic hazard data combined with Italian, European and contiguous US network topologies, I found the followings. Assuming network equipment fails only at a shaking of intensity IX of the MCS scale, there is no significant difference in the cardinality of CFPs and FPs with positive probability. The number of CFPs becomes unacceptably large and slow to compute only at the combined presence of strong earthquakes (with $M_w \geq 8$), short network links ($\leq\sim$ 70 km), and network resources poorly resistant to ground shaking (failing at intensity VI). Structure FP has a low cardinality and can be computed in some minutes in these circumstances too, even on a commodity laptop. Finally, listing CFPs with at most l links rarely yields a list equivalent to keeping some of the most probable CFPs.*

In this thesis, I present numerical results that validate our model and demonstrate the use of the proposed algorithms on some real backbone networks (taken from [34] and [C3], resp.) accompanied with real seismic hazard inputs. The algorithms were implemented in Python 3.6., using its various libraries, respecting the regional failure model presented at Thesis 3.3, while discretizatizing the problem. Runtimes were measured on a commodity laptop with core i5 CPU at 2.3 GHz with 8 GiB of RAM.

Seismic Hazard Representation During the evaluation, we are investigating the failures caused by the next earthquake within a given geographic area, thus we assume there is exactly one earthquake in the investigated time period. The earthquake is identified with its epicenter and moment magnitude [35]: **epicenter** $c_{i,j}$, which represents a latitude-longitude cell on the Earth's surface, taken from a grid of cells over the network area, and **moment magnitude** $M_w \in \{4.6, 4.7, \ldots, 8.6\} =: \mathcal{M}$. We index the cell grid such that $i \in \{1, \ldots, i_{max}\} =: \mathcal{I}_i, j \in \{1, \ldots, j_{max}\} =: \mathcal{I}_j$.

Let E_{i,j,M_w} denote the set of earthquakes with centre point in $c_{i,j}$ and magnitude in $(M_w - 0.1, M_w]$. As cells and magnitude intervals are considered small enough that the failures caused by each earthquake

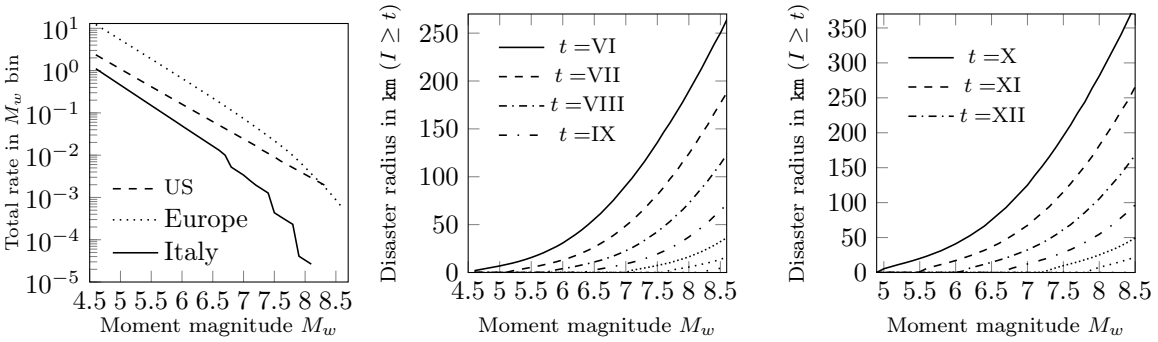

(a) Rate vs. magnitude

(b) Disaster ranges $R(M_w, t)$ for Italy/Europe

(c) Disaster ranges $R(M_w, t)$ for the USA

Figure 12: Seismic input data

in E_{i,j,M_w} will often be the same[12], we will represent all E_{i,j,M_w} with a single earthquake having a center point in the center of $c_{i,j}$ and a magnitude of M_w. Let the probability that the next earthquake is in E_{i,j,M_w} be p_{i,j,M_w}. Note, that these probabilities add up to 1, i.e. $\sum_{i,j \in \mathcal{I}_i \times \mathcal{I}_j} \sum_{M_w \in \mathcal{M}} p_{i,j,M_w} = 1$.

The initial input is the activity rates r_{i,j,M_w} of earthquake types (see Fig. 12a) instead of the p_{i,j,M_w} values, so we have to translate these rates to probabilities. We claim that under the assumption that each kind of earthquake E_{i,j,M_w} is arriving according to independent Poisson arrival processes[13] with parameters r_{i,j,M_w}, the rates of earthquakes E_{i,j,M_w} are straightforward to translate into the probabilities p_{i,j,M_w} of being the next earthquake to occur

$$p_{i,j,M_w} = r_{i,j,M_w} \Big/ \sum_{i,j \in \mathcal{I}_i \times \mathcal{I}_j} \sum_{M_w \in \mathcal{M}} r_{i,j,M_w}. \tag{9}$$

Each network element e has an associated **intensity threshold** $t(e)$ meaning that if the intensity I of the ground shaking reaches this threshold ($I \geq t(e)$) at any point of the physical embedding of e, the element fails. In our simulation, every network element has the same threshold $t(e) := t$, where $t \in \{\text{VI,VII,VIII,IX,X,XI,XII}\} := T$ according to Mercalli-Cancani-Sieberg (MCS) scale [33].

After each earthquake E_{i,j,M_w}, the physical infrastructure (such as optical fibers, amplifiers, routers, and switches) in an area $\mathit{disk}(c_{i,j}, R(M_w, t))$ of a circular disk is destroyed. The center point of $\mathit{disk}(c_{i,j}, R(M_w, t))$ is the center of $c_{i,j}$, while its radius $R(M_w, t)$ is monotone increasing in the magnitude M_w, and decreasing in the intensity threshold t (see Fig. 12b). As earthquakes can occur anywhere in the cell, we increase the radius by the distance between the center of the cell and its outer corners. This way, the disk is always an overestimate of the damaged area of any earthquake in cell $c_{i,j}$ with magnitude M_w.

Calculating activity rates and disaster radii is not subject of my theses.

Simulation results I did our experiments on six topologies, of which one is Italian (Fig. 13), two are Europeans, and three are US topologies. On Table 3, one can see the node and link count of the networks, along with the number of related CFPs and FPs with nonzero probability in case of an intensity tolerance t =VI. Regarding the number of CFPs, one can observe that the networks group by geographic location: the US topologies have roughly 200 to 600 CFPs, the Europeans have more: $\sim 5 \times 10^4$ to $\sim 3 \times 10^5$, and finally, the Italian topology has a lot, $\sim 10^7$. This result is due to the shorter links in Europe, and

[12]The sides of grid cells used in our simulations were 0.05° long in the Italian rate map, and 0.1° in case of the EU and the USA, meaning 4km to 10km of cell side length.

[13]Although earthquakes can be clustered in time and space with their distribution that is over-dispersed if compared to the Poisson law , a common way to treat this problem (i.e. cluster in time and space) is to decluster the earthquake catalog, as we did: before using the catalog, we removed all events not considered main shocks via a declustering filter [36].

Network name	n	m	# CFPs at t =VI	# FPs at t =VI
Italian	25	34	8358809	764
US	26	43	229	144
Nobel EU	28	41	51812	256
EU	37	57	295235	362
N.-American	39	61	348	208
NFSNET	79	108	550	327

Table 3: The investigated network topologies

Figure 13: Italian topology

especially in Italy, and refrains the theoretical complexity of $\Omega(2^\rho)$, i.e. the exponential number of CFPs in function of the maximal number of hit links ρ. The number of FPs ranges in $[144, 764]$, which is significantly smaller than the number of CFPs, and aligns with the theoretical polynomial upper bound. The list of FPs was calcualted in some minutes in all the invesigated settings on a commodoty laptop.

Fig. 14a gives a more detailed explanation on the number of CFPs seen on Table 3 through the values of ρ, the maximum number of links hit by an earthquake. At t =VI, the Italian has a $\rho = 22$, for Nobel EU and EU, $\rho = 15$, and finally, for the US, N.-American and NFSNET, ρ equals 5, 5 and 6, respectively. Apparently, networks covering similar geographic areas have similar ρ values, thus, similar orders of magnitude of CFPs.

On Fig. 14b, we can see the space requirement of structures CFP() and FP() in function of the intensity threshold t. The sudden drop in the number of CFPs dies down after t =IX, and starting from here, the number of CFPs is not significantly higher than the number of FPs.

Remaining at the cardinality of the failed link set, Fig. 14c investigates in details the dependency between CFP(S) and $|S|$, for $|S| = 1, 2$ and 3. There are 34 single link failures in the Italian network whose CFPs range between $[0.0029, 0.12]$, there are 428 dual link failures with non-zero probabilities between $[5\times10^{-9}, 0.0035]$, there are 3030 triple link failures with non-zero probabilities between $[5\times10^{-9}, 0.0015]$. Here we can see that the least probable CFPs with size l is less probable with size $l+1$, thus listing CFPs with at most l links rarely yields a list equivalent to keeping some of the most probable CFPs.

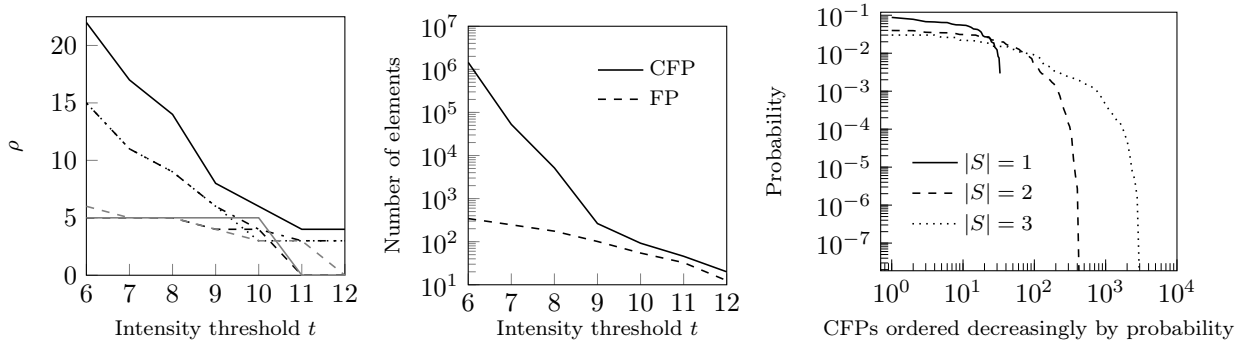

(a) Link density ρ

(b) Average CFP/FP number vs. intensity threshold

(c) CFPs of Italian for t =VI

Figure 14: Simulation results for my PSRLG model

5 Applicability of New Results

Results presented in my Theses constitute state of the art algorithms and models for enumerating regional Shared Risk Link Groups (SRLGs) and regional Probabilistic SRLGs (PSRLGs), which structures are key in translating the composed geometric problem of protecting telecommunication networks against regional failures to purely combinatorial and probabilistic problems, respectively.

Regional SRLG lists offered by Theses 1 and 2 are key inputs for creating geodiverse routing algorithms. Since routing methods are not aware of the geometrical embedding of the topology, providing the list of link sets which may fail simultaneously allows the routing to become disaster-disjoint. SRLG lists M_r and M_k in Theses 1 and 2 are designed for the case when the Internet Service Provider do or do not know the exact geometry of the network topology, respectively.

The results of Theses 1 and 2 are also vital for determining monitoring trails for regional failure detection. This is simply because without knowing the sets of links which are considered to have a high chance of failing jointly, one would have to use an infeasibly large number of costly m-trails for detecting regional failures.

The PSRLG model and data structures (CFP() and FP()) presented in Thesis 3 enable fast service availability queries, which are crucial in optimizing these availabilities. Their effectiveness was demonstrated through converting seismic hazard data to PSRLGs, and could be used for other disaster types too. I believe the proposed data structures CFP() and FP() should be used as standard containers of disaster hazards.

Acknowledgements

I would like to express my deep gratitude to my advisor János Tapolcai, for his patient guidance, enthusiastic encouragement and useful critiques along the past years. His support was indispensable in understanding the research methodologies and becoming a researcher in the field of telecommunications.

I owe many thanks to those who contributed to this thesis work: Erika Bérczi-Kovács, Lajos Rónyai, David Hay, Alessandro Valentini, and all my co-authors. I would also thank Stefan Schmid, David Hay and Costin Raiciu for hosting me as a guest researcher in 2019.

The work presented here was done at the High Speed Networks Laboratory, Budapest University of Technology and Economics (BME). The instruments of COST CA15127 Action "Resilient Communication Services Protecting End-user Applications from Disaster-based Failures" (RECODIS) greatly helped my research.

References

[1] A. de Sousa and D. Santos, "The minimum cost d-geodiverse anycast routing with optimal selection of anycast nodes," in *2019 15th International Conference on the Design of Reliable Communication Networks (DRCN)*, pp. 21–28, March 2019.

[2] 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.

[3] 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.

[4] T. Gomes, J. Tapolcai, C. Esposito, D. Hutchison, F. Kuipers, J. Rak, A. de Sousa, A. Iossifides, R. Travanca, J. André, L. Jorge, L. Martins, P. O. Ugalde, A. Pašić, D. Pezaros, S. Jouet, S. Secci,

and M. Tornatore, "A survey of strategies for communication networks to protect against large-scale natural disasters," in *8th International Workshop on Resilient Networks Design and Modeling (RNDM)*, pp. 11–22, Sept 2016.

[5] 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.

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

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] A. Xie, X. Wang, and S. Lu, "Risk minimization routing against geographically correlated failures," *IEEE Access*, vol. 7, pp. 62920–62929, 2019.

[13] A. Kushwaha, D. Kapadia, A. Gumaste, and A. Somani, "Designing multi-layer provider networks for circular disc failures," in *International Conference on Optical Network Design and Modeling (ONDM)*, (Dublin, Ireland), May 2018.

[14] H. Zang, C. Ou, and B. Mukherjee, "Path-protection routing and wavelength assignment (RWA) in WDM mesh networks under duct-layer constraints," *IEEE/ACM Transactions on Networking (TON)*, vol. 11, no. 2, pp. 248–258, 2003.

[15] 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.

[16] M. T. Gardner and C. Beard, "Evaluating geographic vulnerabilities in networks," in *IEEE Int. Workshop Technical Committee on Communications Quality and Reliability (CQR)*, pp. 1–6, 2011.

[17] S. Trajanovski, F. Kuipers, P. Van Mieghem, *et al.*, "Finding critical regions in a network," in *IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 223–228, IEEE, 2013.

[18] B. Chazelle and H. Edelsbrunner, "An optimal algorithm for intersecting line segments in the plane," *Jurnal of the ACM (JACM)*, 1992.

[19] J. P. Snyder, *Map projections: A working manual*. U.S. Government Printing Office, 1987.

[20] T.-H. Su and R.-C. Chang, "The k-Gabriel graphs and their applications," in *Algorithms. Lecture Notes in Computer Science, vol.450.*, (Berlin, Heidelberg), 1990.

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

[22] S. Orlowski, R. Wessäly, M. Pióro, and A. Tomaszewski, "Sndlib 1.0: survivable network design library," *Networks*, vol. 55, no. 3, pp. 276–286, 2010.

[23] H.-W. Lee, E. Modiano, and K. Lee, "Diverse routing in networks with probabilistic failures," *IEEE/ACM Transactions on Networking*, vol. 18, no. 6, pp. 1895–1907, 2010.

[24] M. Rahnamay-Naeini, J. E. Pezoa, G. Azar, N. Ghani, and M. M. Hayat, "Modeling stochastic correlated failures and their effects on network reliability," in *IEEE Int. Conf. on Comp. Comm. and Networks (ICCCN)*, pp. 1–6, 2011.

[25] "US National Seismic Hazard Maps."

[26] S. Trajanovski, F. A. Kuipers, A. Ilić, J. Crowcroft, and P. Van Mieghem, "Finding critical regions and region-disjoint paths in a network," *IEEE/ACM Transactions on Networking (TON)*, vol. 23, no. 3, pp. 908–921, 2015.

[27] O. Gold and R. Cohen, "Coping with physical attacks on random network structures," in *IEEE ICC*, pp. 1166–1172, 2014.

[28] X. Wang, X. Jiang, A. Pattavina, and S. Lu, "Assessing physical network vulnerability under random line-segment failure model," in *IEEE High Performance Switching and Routing (HPSR)*, pp. 121–126, 2012.

[29] H. Saito, "Analysis of geometric disaster evaluation model for physical networks," *IEEE/ACM Trans. Netw.*, vol. 23, no. 6, pp. 1777–1789, 2015.

[30] H. Saito, "Spatial design of physical network robust against earthquakes," *J. Lightwave Technol.*, vol. 33, no. 2, pp. 443–458, 2015.

[31] F. Iqbal and F. Kuipers, "Spatiotemporal risk-averse routing," in *IEEE INFOCOM Workshop on Cross-Layer Cyber Physical Systems Security (CPSS)*, IEEE, 2016.

[32] E. Papadopoulou and M. Zavershynskyi, "The higher-order Voronoi diagram of line segments," *Algorithmica*, vol. 74, no. 1, pp. 415–439, 2016.

[33] A. Sieberg, "Erdebeben," *Handbuch der Geophysic*, vol. 4, pp. 552–554, 1931.

[34] "Network library." Available at: `https://github.com/jtapolcai/regional-srlg/tree/master/psrlg`.

[35] H. Kanamori, "The energy release in great earthquakes," *Journal of Geophysical Research (1896-1977)*, vol. 82, no. 20, pp. 2981–2987, 1977.

[36] J. K. Gardner and L. Knopoff, "Is the sequence of earthquakes in southern california, with aftershocks removed, poissonian?," *Bulletin of the Seismological Society of America*, vol. 64, pp. 1363–1367, 1974.

Publication of New Results

Book Chapters

[B1] **B. Vass**, J. Tapolcai, D. Hay, J. Oostenbrink, F. A. Kuipers, "How to Model and Enumerate Geographically Correlated Failure Events in Communication Networks", in Guide to Disaster-Resilient Communication Networks, Springer, 2020

[B2] T. Gomes, L. Martins, R. Girao-Silva, D. Tipper, A. Pašić, **B. Vass**, L. Garrote, U. Nunes, M. Zachariasen, J. Rak, "Enhancing Availability for Critical Services", in Guide to Disaster-Resilient Communication Networks, Springer, 2020

[B3] T. Gomes, D. Santos, R. Girão-Silva, L. Martins, B. Nedic, M. Gunkel, F. Dikbiyik, **B. Vass**, J. Tapolcai, J. Rak, "Disaster-Resilient Routing Schemes for Regional Failures", in Guide to Disaster-Resilient Communication Networks, Springer, 2020

Journal Papers

[J1] **B. Vass**, J. Tapolcai, Z. Heszberger, J. Bíró, D. Hay, F. A. Kuipers, J. Oostenbrink, A. Valentini, L. Rónyai, "Probabilistic Shared Risk Link Groups Modelling Correlated Resource Failures Caused by Disasters", IEEE JSAC Special Issue OptNet (submitted)

[J2] J. Tapolcai, L. Rónyai, **B. Vass**, and L. Gyimóthi, "Fast Enumeration of Regional Link Failures Caused by Disasters with Limited Size", IEEE-ACM Transactions on Networking, 2020

[J3] **B. Vass**, E. Bérczi-Kovács, and J. Tapolcai, "Enumerating Maximal Shared Risk Link Groups of Circular Disk Failures Hitting k Nodes", IEEE-ACM Transactions on Networking (submitted)

[J4] **B. Vass**, L. Németh, J. Tapolcai, "The Earth is Nearly Flat: Precise and Approximate Algorithms for Detecting Vulnerable Regions of Networks in Plane and on Sphere", Networks, Wiley, 2020

Conference and Workshop Papers

[C1] B. Németh, Y.-A. Pignolet, M. Rost, S. Schmid, **B. Vass**, "Cost-Efficient Embedding of Virtual Networks With and Without Routing Flexibility", IEEE IFIP Networking, Paris, France, 2020

[C2] D. Haja, **B. Vass**, L. Toka, "Towards making big data applications network-aware in edge-cloud systems", IEEE 8th International Conference on Cloud Networking (CloudNet), Coimbra, Portugal, 2019

[C3] A. Valentini, **B. Vass**, J. Oostenbrink, L. Csák, F. A. Kuipers, B. Pace, D. Hay and J. Tapolcai, "Network Resiliency Against Earthquakes", IEEE Int. Workshop on Resilient Networks Design and Modeling (RNDM), Nicosia, Cyprus, 2019

[C4] A. Pašić, R. Girao-Silva, **B. Vass**, T. Gomes, F. Mogyorósi, P. Babarczi, J. Tapolcai, "FRADIR-II: An Improved Framework for Disaster Resilience", IEEE Int. Workshop on Resilient Networks Design and Modeling (RNDM), Nicosia, Cyprus, 2019

[C5] D. Haja, **B. Vass**, L. Toka, "Improving Big Data Application Performance in Edge-Cloud Systems", IEEE 12th International Conference on Cloud Computing (CLOUD), Milan, Italy, 2019

[C6] **B. Vass**, L. Németh, A. de Sousa, M. Zachariasen and J. Tapolcai, "Vulnerable Regions of Networks on Sphere", IEEE Int. Workshop on Resilient Networks Design and Modeling (RNDM), Longyearbyen (Svalbard), Norway, 2018.

[C7] A. Pašić, R. Girão-Silva, **B. Vass**, T. Gomes, and P. Babarczi, "FRADIR: A Novel Framework for Disaster Resilience", IEEE Int. Workshop on Resilient Networks Design and Modeling (RNDM), Longyearbyen (Svalbard), Norway, 2018.

[C8] J. Tapolcai, **B. Vass**, Z. Heszberger, J. Biró, D. Hay, F. A. Kuipers, and L. Rónyai, "A Tractable Stochastic Model of Correlated Link Failures Caused by Disasters," in Proc. IEEE INFOCOM, Honolulu, HI, USA, 2018.

[C9] J. Tapolcai, L. Rónyai, **B. Vass**, and L. Gyimóthi, "List of Shared Risk Link Groups Representing Regional Failures with Limited Size," in Proc. IEEE INFOCOM, Atlanta, GA, USA, 2017

[C10] **B. Vass**, E. Bérczi-Kovács, and J. Tapolcai, "Enumerating Shared Risk Link Groups of Circular Disk Failures Hitting k nodes," in Proc. International Workshop on Design Of Reliable Communication Networks (DRCN), Munich, Germany, 2017.

[C11] **B. Vass**, E. Bérczi-Kovács, and J. Tapolcai, "Enumerating Circular Disk Failures Covering a Single Node," in Int. Workshop on Resilient Networks Design and Modeling (RNDM), Halmstad, Sweden, 2016.

[C12] **B. Vass**, E. Bérczi-Kovács, and J. Tapolcai, "Shared Risk Link Group Enumeration of Node Excluding Disaster Failures," in Int. Conference on Networking and Network Applications (NaNA), Hakodate (Hokkaido), Japan, 2016. Winner of Best Paper Award

Posters

[P1] **B. Vass**, "Shared Risk Link Groups of Disaster Failures," in IEEE Conference on Computer Communications Student Poster (INFOCOM Poster), 2016.

Theses

[T1] **B. Vass**, "A Combinatorial Geometric Approach for Network Attack Analysis," M.Sc. Thesis (Advisors: E. Bérczi-Kovács and A. Kőrösi), ELTE, Budapest, Hungary

Other Publications

[O1] **B. Vass**, "Hálózatok legfennebb r sugárban ható hibáinak felsorolása" 1st prize at the Scientific Students' Associations conference 2015 organized by the Mathematical Institute of the Eötvös Loránd University (in Hungarian)

Based on Google Scholar, the number of my independent citations is ≥ 24 as of August 17, 2020.