Essence of Geographically Correlated Failure Events in Communication Networks

Balázs Vass*†, János Tapolcai*
*MTA-BME Future Internet Research Group, High-Speed Networks Laboratory (HSN*Lab*),
Budapest University of Technology and Economics, {vb,tapolcai}@tmit.bme.hu
† EIT Digital Industrial Doctoral School

***Abstract*—Modeling and listing the joint device failures of telecommunication optical backbone networks caused by large-scale regional disasters is the aim of my dissertation [1] digested in the following. The use-cases of these failure lists include helping the operators of modern telecommunication networks to meet the predefined Quality-of-Service (QoS) conditions. Informally speaking, the task tackled is translating the composed geometric problem of protecting telecommunication networks against regional failures to small-sized purely combinatorial and probabilistic problems, respectively. The development of this translation framework relied on the following pillars: 1) constructing failure models that make the best use of the data available, 2) giving fast algorithms for determining the resulting failure lists, 3) providing a theoretical and practical analysis of the complexity of the algorithms and the properties of the failure lists. The offered failure lists can be leveraged for enhancing network preparedness against disasters.**

I. Introduction

The Internet has become a topmost critical infrastructure. Due to the importance of telecommunication services (as a base for stock market, telesurgery, etc.), improving the preparedness of networks to regional failures is becoming a key issue [2]–[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 on the problem of how to prepare networks to survive regional failures, where the first solutions have assumed that fibers 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 the different types of disasters (e.g., seismic hazard maps for earthquakes) and identifying the hotspots of the disasters [3], [5], [6], [8], [10]. 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 are purely devoted to this particular problem and adapt computational geometric based approaches to capture all of the regional failures and represent them in a compact way [9], [15]–[20], where the major challenge is that regional failures can have arbitrary locations, shapes, sizes, effects, etc. Unfortunately, regional failures are not self-discoverable in practice [21]; this, together with the high number of severe network outages witnessed in the last decades [22]–[26][1] present clear evidence that selecting the proper list of regional failures is still a challenging problem to solve [3], [5]–[10], [27]. To fill this gap in reliable network design, the presented research is devoted to enhancing the state of the art and suggests unified definitions, notions, and terminology.

The output of the approaches discussed in this dissertation 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 (or link set) 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)[2] are defined by the network operators, where each SRLG is a set of links whose joint failure the recovery mechanism should be prepared for. The first part of the dissertation purely focuses on how to define and enumerate SRLGs that cover all types of disasters. In the second part of the dissertation, the question of defining a realistic and applicable Probabilistic SRLG (PSRLG) failure model is addressed.

It turns out that, surprisingly, in practice, only a small number of SRLGs or PSRLGs are needed to serve as inputs for the higher-layer network management tools. Informally speaking, methods offering lists of SRLGs and PSRLGs translate the composed geometric problem of protecting telecommunication networks against regional failures to small-sized purely combinatorial and probabilistic problems, respectively. These findings open up the possibility of leveraging regional (P)SRLG lists for enhancing network preparedness against disasters.

J.T. and B.V. are partially supported by the National Research, Development and Innovation Fund of Hungary (grant No. 124171, 128062, 134604, 135606 and 137698). Supported by the ÚNKP-21-4 New National Excellence Program of the Ministry of Innovation and Technology from the source of the National Research, Development and Innovation Fund.

[1] A recent example is a few days long telecom outage during Cyclone Amphan in West Bengal in May of 2020 as a result of around 100 fiber cuts due to the falling of trees by the wind speeding up to 190km/h.

[2] An SRLG is a set of links that are considered to have a significant chance of failing together.

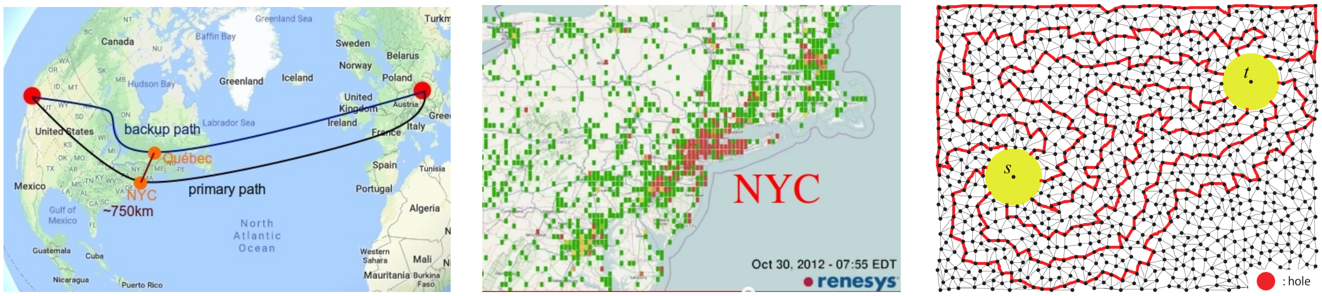

(a) (b) (c)

Figure 1. (a) To avoid most disasters, ensuring several hundred km distance between the primary and the backup paths is enough. (b) The status of the routers during Hurricane Sandy, 2012. Most of the routers in NYC are not functioning, Boston also has problems. (c) For disaster-disjoint routing, storing the disasters and the geometric embedding of the network can be replaced by a short list of SRLGs indicating the link sets that can be hit by the same disaster. Picture credits to [28]–[30].

II. Example Use-Cases of SRLG and PSRLG Lists

Two basic use-cases of SRLG and PSRLG lists are the resilient routing [31], [32], and determining service availabilities depicted in Fig. 1 and 2, respectively.

In Fig. 1a, we can see a pair of imagined primary and backup paths stretching between Central Europe and California. By demanding a distance of several hundred km between the two paths (except their endpoints), we ensure they have a negligible probability of failing together. Fig. 1b depicts the state of the routers during Hurricane Sandy that was considered a severe disaster. In Fig. 1c, a maximal number (here, 7) of s-t paths are shown, such that there are no two paths that are hit at the same time by any position (outside of the yellow regions) of the red disk depicted in the bottom. Here, in the input, instead of storing the possible disasters and the geometric embedding of the network, one can simply use a list of SRLGs indicating the link sets that can be hit by the same disaster: if path p_1 goes through SRLG S, then path p_2 is forbidden to do so.

The example depicted in Fig. 2 underlines difficulty of estimating service availabilities. There, user U reaches her data either in cloud C_1 or in cloud C_2. At the next disaster, the connections to C_1 and C_2 may fail in regions V_1 and V_2, respectively, with an equal chance of $P(V_1) = P(V_2) = 0.001$. If V_1 and V_2 are far from each other (as in Fig. 2a), we may suppose the connections fail independently, meaning an unavailability of $P(V_1) \cdot P(V_2) = 0.000001$ of the cloud. If V_1 and V_2 are at the same place (same bridge, valley, etc., Fig. 2b), the unavailability of the cloud will be $P(V_1) = P(V_2) = 0.001$. If V_1 and V_2 are 'close' to each other, but not in the same place, the availability of the cloud under the next disaster is difficult to estimate. Easing the service availability queries demands the investigation of probabilistic extension of the SRLGs, and designing a realistic probabilistic regional failure model.

III. Problem Statement

In this dissertation, I study both the deterministic and probabilistic versions of the problem of representing the effect of regional disasters on telecommunication networks. In the first part, I purely focus on how to define and enumerate the most lifelike (deterministic) SRLGs that cover all types of disasters. Fig. 3 depicts the most natural strategies for guaranteeing a level of separation between the primary and the backup path in the absence of the simultaneous presence of both a precise knowledge of the physical positions of the network elements and expertise on possible disasters.

Without any requirements, there might be *no separation* at all between these paths. A common practice is to ensure *link-disjointness* on the paths via enumerating all the single link failures as SRLGs. Compared to this, *node-disjointness* (except for the source and destination nodes s and t) ensures resiliency to any single element failure. An SRLG list providing node-disjointness consists of link sets incident to each node.

To enhance the preparedness granted by node-disjointness, one has to leverage some background information on the geographical embedding of the network. Typically, communication networks have few edge crossings, and links are a few hundred kilometers long. Thus it makes sense to grant a given h hops

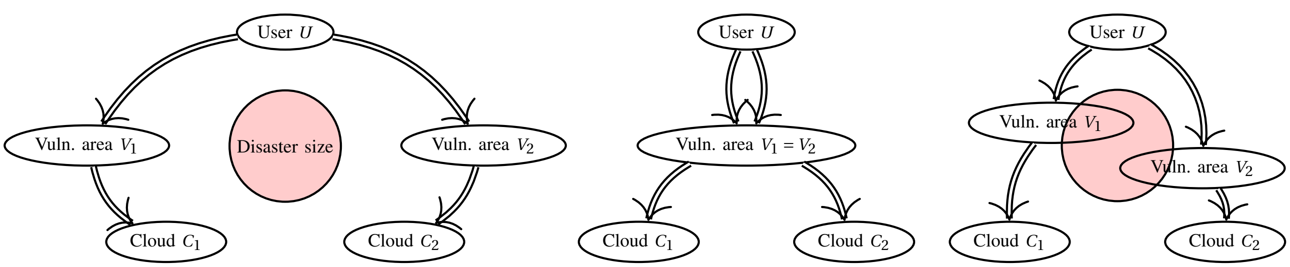

(a) Availability: 99.9999% (b) Availability: 99.9% (c) Availability: hard to tell

Figure 2. User U reaches his data either in cloud C_1 or in cloud C_2. At the next disaster, the connections to C_1 and C_2 may fail in regions V_1 and V_2, respectively, with an equal chance of $P(V_1) = P(V_2) = 0.001$. If V_1 and V_2 may be hit by the same disaster, but are not co-located, the cloud availability under the next disaster is hard to estimate.

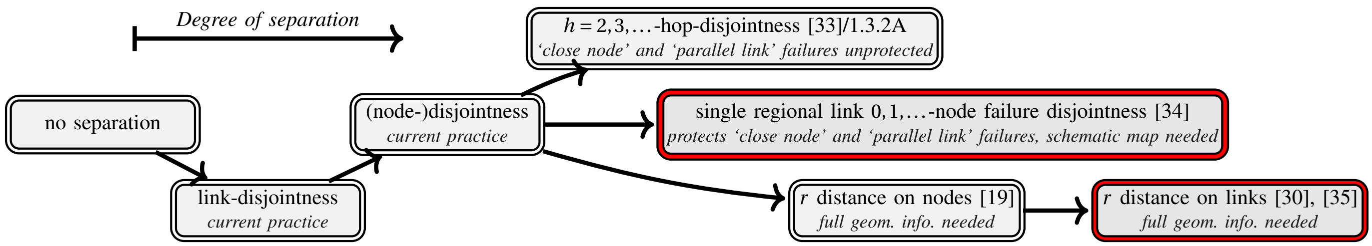

Figure 3. Strategies for separating the primary and backup paths in increasing strength (the more right the better the separation is). By utilizing only a limited geometric information on the network topology, this paper offers an SRLG list M_k that ensures single regional link $k=0,1,\ldots$-node failure disjointness. SRLG list M_k fills the gap between h-hop disjointness and r-distance disjointness. Arcs are directed from weaker separations towards the stronger ones [34].

distance between the primary and backup paths. For this, one may list the links in the vicinity of every network link or node as SRLGs [33]. Unfortunately, a distance of h hops does not necessarily protect the failure of links crossing the same bridge or a bunch of close nodes.

Knowing the exact geographic embedding of the network topology solves this issue: supposing that a disaster may damage the network equipment within a radius r around its epicenter (and the rest of the network is left intact), one only has to list all the maximal link sets which can be hit by a circular disk with radius $r/2$ in a list $M_{r/2}$. Here the challenge is giving fast polynomial algorithms for determining $M_{r/2}$ and showing that $M_{r/2}$ has a manageable size so that we can provide r-disjointness for large network topologies too. My first thesis group (and Subsec. IV-A) is devoted to this issue.

In many cases, one has only a rough idea of the physical embedding of the network, e.g., when the topology is rented from a Physical Infrastructure Provider [36]. In other words, they have a schematic map of the network, where the scale is not necessarily preserved over the area, and routes of links are only known to be within certain areas. In such circumstances, one can provide a separation which is weaker compared to M_r, but still better than relying only on hop-count (cf. Fig. 3): in a list M_k, one can gather the maximal link sets which can be hit by a circular disk hitting k nodes. In the second thesis group (Subsec. IV-B), I provide a model to handle this case together with theoretical and experimental upper bounds on the size and construction time of M_k.

Regarding the prior state of the art, there was no PSRLG model, which would take into count that link failures are not independent when a disaster happens. Also, they did not represent the possible disasters as accurately as possible. In the second part of the dissertation (third thesis group, and Subsec. IV-C, resp.), the aim was to define a realistic and applicable Probabilistic SRLG failure model, which takes into count the failures' correlation. In the evaluation, we use a seismic hazard representation, which preserves more information on possible future earthquakes than usual hazard maps.

A. Related Works and Charting the Landscape of (P)SRLG Enumerating Problems

To have a better overview of the problem versions tackled by both other researchers and our group, in the following, a charting of the (PSRLG) enumerating problems is given based on the input data quality/precision.

Informally speaking, the most important input information parts are the 1) geometric embedding of the network, and the 2) (probabilistic) disaster effects. Unfortunately, in practice, it is far not obvious that this information is available with high precision. As depicted in Fig. 4, we might distinguish three levels of information quality both on the geographic embedding and on the disaster effects and classify the offered (P)SRLG approaches according to these. In the following, the problems appearing in related works are briefly depicted and referenced.

In case of good information on embedding, no information of diasasters, instead of computing list M_r (Subsec. IV-A), one may assume that the disaster has a fixed shape, e.g., an equilateral triangle of any orientation, and calculate list $M_{\mathtt{shape}}$ of maximal link failures caused by this shape [37]. Another possibility is that a set $\mathscr{D}$ of disaster areas is given[3], and computes the list $M_{\mathscr{D}}$ of maximal SRLGs caused by these disasters (see, e.g., a non-probabilistic version of [38]). Having detailed probabilistic disaster data coupled with a precise map of the network allows us to compute PSRLG lists. Here the challenge is to create a model that correctly captures the joint failure probabilities of network elements while producing an output of affordable size (see Chapter IV-C or [15]). Furthermore, a list of PSRLGs enables collecting those SRLGs that have a failure probability above a threshold T. From among these, one can collect the maximals in a list M_T [39]. We note that in the natural condition when the

[3]Where in the interior of the disaster area, everything is damaged, while in the exterior, no failure happens.

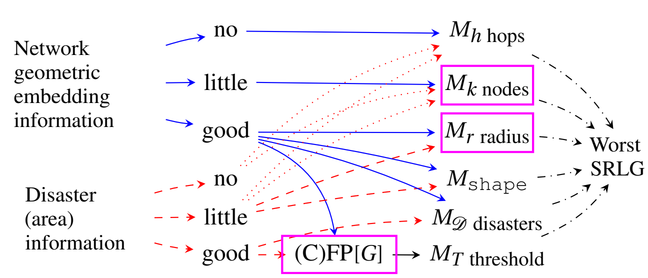

Figure 4. A mind map of SRLG and PSRLG problems related to the quality of input data. For a graph G, (C)FP[G] stores PSRLGs, while lists M_* consists of SRLGs. Problems studied in this dissertation ($M_{k\text{ nodes}}, M_{r\text{ radius}}$, (C)FP[$G$]) are highlighted with purple rectangles.

vulnerability metric or a protection mechanism is monotone[4], the worst SRLG fulfilling a given criteria c (e.g., SRLGs that can be hit by circular disks with a range r or hitting k nodes) will be part of the set of exclusion-wise maximal SRLGs fulfilling c. Thus a worst SRLG can be found by simply searching for the worst SRLG in the list of maximal SRLGs fulfilling c (e.g., in M_r, M_k) [33]. This way, the results presented in the dissertation are firm base ground for solving a whole family of related problems.

IV. Contributions of the Dissertation

The contribution of the dissertation is two-fold. Firstly, it offers provably short lists of SRLGs covering all the failures caused by regional disasters. For this, both a model where the exact geographical embedding of the network is known and another model where only a schematic map of the topology is available is given. Fast polynomial algorithms calculating the above lists are offered.

On the other hand, this dissertation proposes a unified terminology on the Probabilistic SRLGs, along with a model for PSRLG enumeration that produces realistic failure probabilities, where the computed data structure can be stored in provably small space in case of circular disasters, and it handles the correlation of link failures better than the prior state-of-art. The contributions are detailed as follows.

A. Maximal SRLGs Induced by Disks of Radius r

As the first thesis point, in papers [17], [35], [40], [41], polynomial algorithms were proposed 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. Theoretical upper bounds were given on the cardinality of both M_r^p and M_r^s. It is proved that the proposed algorithm for planar embeddings has a computational complexity that is tight in the number of network nodes. Finally, the similarity of M_r^p and M_r^s in practice is compared.

More precisely, in [17], [35], an algorithm was proposed, which, in case of representing a connected network topology $G(V,E)$ 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 $O\big((|V|+x)\big(\log|V|+\phi_r^2\rho_r^5\big)\big)$, 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, and finally, ϕ_r is the maximum number of nodes in the $3r$-neighborhood of a link. The complexity of the proposed algorithm is tight in $|V|$. It is also proved that the cardinality of M_r^p is $O\big((|V|+x)\,\rho_r\big)$, and that this bound is tight. It is proved that $\left|\bigcup_{0<r'<r} M_{r'}^p\right|$ is $O((|V|+x)\rho_r^2)$.

We were also curious about the ipmreciseness caused by the distortion when the network topologies are projected from the Earth surface to the plane. To this end, in [40], [41], a heuristic algorithm is proposed, which, considering a connected network topology $G(V,E)$ on a sphere with links considered as chains of geodesics, and considering a related sufficiently dense set $\mathscr{P}$ of disaster center points, computes list M_r^s of maximal link sets hit by a circular disk with radius r in $O(|\mathscr{P}|[(|V|+x)\gamma+|\mathscr{P}|\rho_r])$, where x is the number of link crossings, γ is the maximal number of geodesics a link stands of, and ρ_r is the maximum number of links which are hit by a circular disk with radius r. Our simulations 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 causes less inaccuracy than the lack of knowledge on the disaster characteristics. There, the planar representation can serve the purpose of vulnerable region detection well enough.

B. Maximal SRLGs Caused by Circular Disks Hitting k Nodes

As the second thesis point, in [20], [34], [42], [43], to ensure geographic distance between primary and backup paths when the geographical embedding of the network topology is approximate, a model for enumerating regional SRLGs relying only on a schematic map of the network topology was proposed. For networks described in this model, [34] 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. Theoretical upper bounds on the cardinality of M_k were given. Evaluating the model and data structure, we showed that in the case of real network topologies as input combined with practical k values, M_k is reasonably short.

More detailed, [34] introduced the Limited Geometric Information Failure Mode, which goes as follows. To ensure geographic distance between primary and backup paths when the geographical embedding of the network topology is approximate, we proposed the following model. The (not necessarily planar) network is modelled as an undirected connected geometric graph $G=(V,E)$ 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 center 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*.

An algorithm was given, which, in case of representing a network topology $G(V,E)$ in the Euclidean plane with each link $e\in E$ 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 $O\big(|V|^2\big(\big(k^2+1\big)\rho_k^3+\rho_k\gamma+\big(k+1+\log(n\rho_0)\big)\rho_0\gamma\big)\big)$, where ρ_k denotes the maximal number of links hit by a circular disk hitting at most k nodes. List M_k has $O\big(n\,(k+1)\,\rho_k\big)$ elements, this bound being tight in these parameters for $k=O(1)$.

Regarding to the simulation results, in case of real network topologies, with their edges considered polyginal chains and

[4]Vulnerability metric or protection mechanism μ is *monotone*, if, according to μ, for any $E_1\subseteq E_2$, the failure of E_2 is *worse* than the failure of E_1.

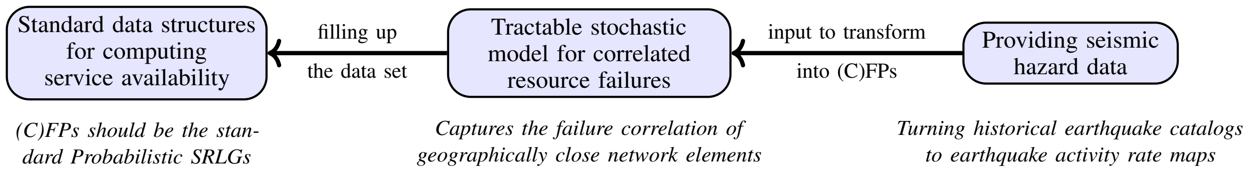

Figure 5. Main contributions of thesis group 3: there is offered a 1) standard data structures (for graph G, CFP[G] and FP[G]) for storing joint failure probabilities of link sets, 2) a tractable stochastic model of network element failures caused by disasters, and finally 3) providing the seismic hazard data represented it in a more precise way than the usual hazard maps. Note that our stochastic model can handle the combined inputs of an arbitrary number of disaster families (e.g., tornadoes, earthquakes, tsunamis, etc.). Structures CFP[G] and FP[G] could be established using other models too.

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 the investigated networks for $k=0,1$, and grew to only to <25 for $k=5$. The conclusion is that list M_k has a reasonably small size for practical k values.

C. PSRLGs Modeling Correlated Link Failures Caused by Disasters

In the third and last thesis point, in papers [18], [44], [45], a stochastic model of link failures caused by disasters was defined (cf. Fig. 5), which considers the correlation between failures of links that are geographically close to each other. To unify the notions and terminology on Probabilistic SRLGs, we proposed standard data structures for containing the disaster probabilities. In the case of circular disk shaped disasters, for the size and query time of these data structures, theoretical upper bounds were given. Evaluation the model and data structures showed that in the case of taking real seismic data as input, these data structures have a manageable size.

More precisely, inspired by earthquake behaviours, we defined a stochastic model of link failures caused by disasters. This model is the first to explicitly consider the correlation between failures of links that can be subject to the same disaster. To unify the notions and terminology linked to probabilistic extensions of Shared Risk Link Groups, two standard data structures for describing the disaster probabilities were proposed. Namely, for a graph G, these structures are called FP[G] and CFP[G], respectively. In FP[G], for each link set S, the probability that *exactly* S will fail is stored as FP(S) , while in CFP[G], the probability that *at least* S will fail is stored as CFP(S) .

On the theoretical side, in case of disasters having shapes of circular disks in a given L_p metric, representing the network topology $G(V,E)$ 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, and the maximum number of links which are hit by one of the disasters by ρ, the followings were found. There are $O((|V|+x)\rho^2\gamma^4)$ FPs with nonzero probability. The number of CFPs with positive probability is lower bounded by $\Omega(2^\rho)$ and upper bounded by $O(2^\rho(|V|+x)\rho^2\gamma^4)$. Storing all the positive CFPs in a balanced binary tree, the worst-case query time of the CFP of a given link set is $O(\rho\log((|V|+x)\rho\gamma))$. Storing all the positive FPs in a list, the query time of the CFP of a given link set is $O((|V|+x)\rho^2\gamma^4)$.

Using real-world seismic hazard data combined with Italian, European, and contiguous US network topologies, we found the following. Assuming network equipment fails only at shaking of intensity VIII 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 ($\leq\sim 50$ 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.

V. Possible Future Work

Possible future directions of this research include but not restrict to:

- better integration of failure modeling into disaster resilience approaches (FRADIR [39], [46], [47]-like studies),
- proving our conjecture that the regional SRLG-disjoint routing problem is in solvable in polynomial time,
- evaluating our probabilistic failure model with more complex real-world inputs,
- as a side-track of a future SRLG list comparing study, creating the 'SRLG-Zoo', a webpage similar to Topologyzoo [48], from where one could download network topologies and related (P)SRLG lists.

VI. Conclusion

The dissertation digested here is dedicated to proving that the effect of regional disasters (natural on man-made) can be modeled with a low number of SRLGs or PSRLGs. These carefully constructed lists of (P)SRLGs can be used as input for e.g., network recovery/planning mechanisms.

The presented lists of vulnerable regions can be used as input of various problems arising in the field of network resiliency. Some of these problems are resilient (geodiverse) routing, (k-)content connectivity, network failure detection, service availability queries, resilient backbone network planning, disaster avoidance control, resilient SDN, etc.

REFERENCES

[1] B. Vass, "Models and algorithms for enumerating geographically correlated failure events in communication networks," dissertation, Budapest University of Technology and Economics, 2021. [Online]. Available: http://hdl.handle.net/10890/16654

[2] 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)*, March 2019, pp. 21–28.

[3] F. Dikbiyik, M. Tornatore, and B. Mukherjee, "Minimizing the risk from disaster failures in optical backbone networks," *J. Lightw. Technol.*, vol. 32, pp. 3175–3183, 2014.

[4] T. Gomes, J. Tapolcai, C. Esposito, D. Hutchison *et al.*, "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)*, Sept 2016, pp. 11–22.

[5] M. F. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of disaster-resilient optical datacenter networks," *J. Lightw. Technol.*, vol. 30, 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," *IEEE Commun. Mag.*, vol. 52, 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] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of the fiber infrastructure to disasters," *IEEE/ACM Trans. Netw.*, vol. 19, pp. 1610–1623, 2011.

[10] R. Souza Couto, S. Secci, M. Mitre Campista, K. Costa, and L. Maciel, "Network design requirements for disaster resilience in IaaS clouds," *IEEE Commun. Mag.*, vol. 52, pp. 52–58, 2014.

[11] A. Xie, X. Wang, and S. Lu, "Risk Minimization Routing Against Geographically Correlated Failures," *IEEE Access*, vol. 7, 2019.

[12] Jacek Rak, et al., "COST Action Resilient communication services protecting end-user applications from disaster-based failures (RECODIS)," http://www.cost.eu/COST_Actions/ca/CA15127.

[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, pp. 248–258, 2003.

[15] P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay *et al.*, "The resilience of WDM networks to probabilistic geographical failures," *IEEE/ACM Trans. Netw.*, vol. 21, pp. 1525–1538, 2013.

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

[17] 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, USA, may 2017.

[18] J. Tapolcai, B. Vass, Z. Heszberger, J. Biró *et al.*, "A tractable stochastic model of correlated link failures caused by disasters," in *Proc. IEEE INFOCOM*, Honolulu, USA, Apr. 2018.

[19] 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.

[20] 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, march 2017.

[21] J. Strand, A. L. Chiu, and R. Tkach, "Issues for routing in the optical layer," *IEEE Commun. Mag.*, vol. 39, pp. 81–87, 2001.

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

[23] Y. Nemoto and K. Hamaguchi, "Resilient ICT research based on lessons learned from the Great East Japan Earthquake," *IEEE Commun. Mag.*, vol. 52, pp. 38–43, 2014.

[24] A. Kwasinski, W. W. Weaver, P. L. Chapman, and P. T. Krein, "Telecommunications power plant damage assessment for Hurricane Katrina – site survey and follow-up results," *IEEE Systems Journal*, vol. 3, pp. 277–287, 2009.

[25] J. S. Foster Jr, E. Gjelde, W. R. Graham, R. J. Hermann *et al.*, "Report of the commission to assess the threat to the united states from electromagnetic pulse (EMP) attack: Critical national infrastructures," DTIC Document, Tech. Rep., 2008.

[26] R. Wilhelm and C. Buckridge, "Mediterranean fibre cable cut – a RIPE NCC analysis," 2008. [Online]. Available: https://www.ripe.net/data-tools/projects/archive/mediterranean-fibre-cable-cut

[27] B. Mukherjee, M. Habib, and F. Dikbiyik, "Network adaptability from disaster disruptions and cascading failures," *IEEE Commun. Mag.*, vol. 52, pp. 230–238, 2014.

[28] Google maps. Accessed: 2021. [Online]. Available: https://www.google.com/maps/

[29] Renesys. Accessed: 2021. [Online]. Available: http://www.renesysgroup.com/

[30] Y. Kobayashi and K. Otsuki, "Max-flow min-cut theorem and faster algorithms in a circular disk failure model," in *IEEE INFOCOM*, 2014.

[31] J. P. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar *et al.*, "Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines," *Computer Networks*, vol. 54, 2010.

[32] J. Rak, *Resilient routing in communication networks*. Springer, 2015.

[33] B. Vass, J. Tapolcai, D. Hay, J. Oostenbrink, and F. Kuipers, "How to model and enumerate geographically correlated failure events in communication networks," in *Guide to Disaster-Resilient Communication Networks*. Springer, 2020, pp. 87–115.

[34] B. Vass, J. Tapolcai, and E. R. Bérczi-Kovács, "Enumerating Maximal Shared Risk Link Groups of Circular Disk Failures Hitting k Nodes," *IEEE/ACM Transactions on Networking*, vol. 29, pp. 1648–1661, 2021.

[35] 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*, vol. 28, pp. 2421–2434, 2020.

[36] "Shaping Europe's digital future, Actors in the broadband value chain," 2019, European Commission, https://ec.europa.eu/digital-single-market/en/actors-broadband-value-chain.

[37] 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 Trans. Netw.*, vol. 23, pp. 908–921, 2015.

[38] J. Oostenbrink and F. Kuipers, "Computing the impact of disasters on networks," *ACM SIGMETRICS Performance Evaluation Review*, vol. 45, pp. 107–110, 2017.

[39] A. Pašić, R. Girao-Silva, B. Vass, T. Gomes *et al.*, "FRADIR-II: An Improved Framework for Disaster Resilience," in *2019 11th International Workshop on Resilient Networks Design and Modeling (RNDM)*, 2019.

[40] B. Vass, L. Németh, A. de Sousa, M. Zachariasen, and J. Tapolcai, "Vulnerable regions of networks on sphere," in *Int. Workshop on Resilient Networks Design and Modeling (RNDM)*, Longyearbyen (Svalbard), Norway, Aug. 2018.

[41] B. Vass, L. Németh, and 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.

[42] 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)*, 2016.

[43] 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, Sep. 2016.

[44] A. Valentini, B. Vass, J. Oostenbrink, L. Csák *et al.*, "Network resiliency against earthquakes," in *2019 11th International Workshop on Resilient Networks Design and Modeling (RNDM)*, Oct 2019, pp. 1–7.

[45] B. Vass, J. Tapolcai, Z. Heszberger, J. Bíró *et al.*, "Probabilistic Shared Risk Link Groups Modeling Correlated Resource Failures Caused by Disasters," *IEEE Journal on Selected Areas in Communications*, vol. 39, pp. 2672–2687, 2021.

[46] A. Pašić, R. Girao-Silva, B. Vass, T. Gomes, and P. Babarczi, "FRADIR: A Novel Framework for Disaster Resilience," in *2018 10th International Workshop on Resilient Networks Design and Modeling (RNDM)*, 2018.

[47] A. Pašić, R. Girão-Silva, F. Mogyorósi, B. Vass *et al.*, "eFRADIR: An Enhanced FRAmework for DIsaster Resilience," *IEEE Access*, vol. 9, pp. 13 125–13 148, 2021.

[48] "The internet topology zoo," accessed: 2021. [Online]. Available: http://www.topology-zoo.org/dataset.html