Probabilistic Shared Risk Link Groups Modeling Correlated Resource Failures Caused by Disasters

Vass, Balazs; Tapolcai, János; Heszberger, Zalan; Biro, Jozsef; Hay, David; Kuipers, Fernando A.; Oostenbrink, Jorik; Valentini, Alessandro; Ronyai, Lajos

DOI

[10.1109/JSAC.2021.3064652](https://doi.org/10.1109/JSAC.2021.3064652)

Publication date

2021

Document Version

Accepted author manuscript

Published in

IEEE Journal on Selected Areas in Communications

Citation (APA)

Vass, B., Tapolcai, J., Heszberger, Z., Biro, J., Hay, D., Kuipers, F. A., Oostenbrink, J., Valentini, A., & Ronyai, L. (2021). Probabilistic Shared Risk Link Groups Modeling Correlated Resource Failures Caused by Disasters. *IEEE Journal on Selected Areas in Communications*, 39(9), 2672-2687. [9373653]. <https://doi.org/10.1109/JSAC.2021.3064652>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Probabilistic Shared Risk Link Groups Modelling Correlated Resource Failures Caused by Disasters

Balázs Vass*, János Tapolcai*, Zsolt Heszberger*, József Bíró*, David Hay†, Fernando A. Kuipers‡, Jorik Oostenbrink‡, Alessandro Valentini§, Lajos Rónyai¶

*MTA-BME Lendület Future Internet Research Group and MTA-BME Information Systems Research Group Budapest University of Technology and Economics (BME), {balazs.vass, tapolcai, heszi, biro}@tmit.bme.hu

†School of Engineering and Computer Science, Hebrew University, Jerusalem, Israel, dhay@cs.huji.ac.il ‡Delft University of Technology, Delft, the Netherlands, {J.Oostenbrink, F.A.Kuipers}@tudelft.nl §DiSPuTer

Department, University of Study “G.d’Annunzio” of Chieti-Pescara, Chieti, Italy, alessandro.valentini@unich.it

¶Institute for Computer Science and Control (ELKH SZTAKI) and BME, ronyai@sztaki.hu

Abstract—To evaluate the expected availability of a backbone network service, the administrator should consider all possible failure scenarios under the specific service availability model stipulated in the corresponding service-level agreement. Given the increase in natural disasters and malicious attacks with geographically extensive impact, considering only independent single component failures is often insufficient. This paper builds a stochastic model of geographically correlated link failures caused by disasters to estimate the hazards an optical backbone network may be prone to and to understand the complex correlation between possible link failures. We first consider link failures only and later extend our model also to capture node failures. With such a model, one can quickly extract essential information such as the probability of an arbitrary set of network resources to fail simultaneously, the probability of two nodes to be disconnected, the probability of a path to survive a disaster. Furthermore, we introduce standard data structures and a unified terminology on *Probabilistic Shared Risk Link Groups* (PSRLGs), along with a pre-computation process, which represents the failure probability of a set of resources succinctly. In particular, we generate a quasilinear-sized data structure in polynomial time, which allows the efficient computation of the cumulative failure probability of any set of network elements. Our evaluation is based on carefully pre-processed seismic hazard data matched to real-world optical backbone network topologies.

Index Terms—Disaster resilience, network failure modeling, probabilistic shared risk link groups, PSRLG enumeration, seismic hazard, Voronoi diagram

I. INTRODUCTION

A crucial part of network management is guaranteeing high availability of network services. For backbone optical networks, the required level of service availability is usually

explicitly defined in a contract between the communication service provider (CSP) and the customer, called a service-level agreement (SLA). A violation of the agreed-upon service availability may lead to a financial penalty for the CSP; hence, CSPs must carefully (under-) estimate the availability of their services and, if necessary, reserve protection resources and implement recovery schemes to meet the availability demands. A typical availability value is “five-nine” (99.999%), which translates to an average of at most 5.26 minutes of downtime per year. However, a recent taxonomy of Internet failures [3] has revealed that big network outages last much longer and are often caused by disasters beyond the protection schemes deployed to protect against single failures. As a first step, this paper focuses on how to take into account the correlations between link failures properly. We provide efficient methods to compute and store the link failure correlation in tightly-coupled systems (instead of limiting the set of disasters to a small number or wrongly assuming link-failure events to be independent [4]–[6]).

The problem of correlated network element failures has become more severe in the last decades due to the increased use of virtual environments, whose physical structure is typically hidden from the user. Nevertheless, networks are built on physical infrastructure and comprise optical cross-connects and fibers, prone to physical failures. While some of these failures are isolated, in many cases, several nodes and links located in a geographic area fail simultaneously, e.g., due to a natural disaster, such as an earthquake, a hurricane, or a tsunami [7], [8]. A recent example is a few-day-long telecom outage during Cyclone Amphan in West Bengal in May of 2020 due to around 100 fiber cuts due to tree falls by a 190km/h wind. Such geographically correlated failure events are also called *regional failures* and, due to their significant impact, are receiving increased attention [4], [8]–[27].

A. Related Work

Computing availability in the presence of independent single-point failures is a well-investigated topic (cf. [3], [28]–[32] and references therein). Also, dealing with correlated

Manuscript received July 3, 2020; revised November 24, 2020; accepted January 5, 2021. An earlier version of the paper appeared at IEEE INFOCOM 2018 [1]. The pre-processing of the seismic input follows the approach of our paper [2] appeared at IEEE RNDM 2019. Part of this work has been supported by COST Action CA15127 (RECODIS), the Hungarian Scientific Research Fund (grant No. OTKA K124171, K115288, FK17 123957, KH18 129589, and K17 124171), and the Hungarian Ministry of Innovation and the National Research, Development and Innovation Office within the framework of the Artificial Intelligence National Laboratory Programme, and the Federmann Cyber Security Research Center at the Hebrew University in conjunction with the Israel National Cyber Directorate in the Prime Minister’s Office. Zsolt Heszberger is also supported by the MTA Bolyai János Research Grant and the UNKP-20-4 Bolyai+ Research Grant.

Fig. 1. Main contributions: We offer 1) standard data structures (for graph G , $\text{CFP}[G]$ and $\text{FP}[G]$) 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 $\text{CFP}[G]$ and $\text{FP}[G]$ could be established using other models too.

failures has a long history in the form of Shared Risk Link Groups (SRLGs) (e.g., [24], [28], [29], [31]–[33]). An SRLG typically comprises a few network components (links or nodes) with considerable risk of failing together. There have been some efforts to attach probability values to an SRLG, called Probabilistic SRLG (PSRLG) [34], [35]. A natural approach is to select a set of disaster scenarios as input [9], e.g., based on historical data. It is mostly assumed that the risk groups are part of the input, and for example, the aim is to find a pair of risk-disjoint paths. There has been some work, e.g., [24], [36], where the risk groups are based on the proximity of links to each other, which may be considered a simplistic form of geographically correlated failures. The terminology on PSRLGs has not been unified yet.

Much of the work on regional failures has assumed a given disaster shape (often a circular disk or even a line segment) and, under that particular model, has addressed specific sub-problems in network planning, like finding the most vulnerable part(s) of the network [10]–[12], [16], studying the impact on the network of a randomly placed disaster [20]–[22], designing a network and its services with disaster resiliency in mind [13], [15], [17], [18], and (re)routing of connections to minimize service impact due to a disaster [14], [23]. Some work has considered probabilities, either in the context of a disaster having a certain probability of disconnecting a link, e.g., [4] or in the context of only having partial (probabilistic) information on the geographical layout of a network, e.g., [19].

While the papers mentioned above considered geographically correlated failures, a common property of the targeted sub-problems is to search for the location(s) where a disaster will cause the maximum *expected* damage to the network. In particular, this is a simple averaging process that is unable to exhibit correlations among many important failure events. The problem of precisely and quickly calculating the correlations between link failures for a more thorough network vulnerability assessment has not been addressed sufficiently.

B. Main Contributions

The main contributions of this paper are the following:

- We provide a general stochastic model of disasters to explicitly capture the correlations between resource failures as a result of regional disasters.
- To unify the terminology, we offer two natural standard definitions of the meaning of the probability involved in Probabilistic Shared Risk link Groups (PSRLGs).
- We devise a pre-computation process to perform the necessary numerical integration off-line. In terms of the

network size, there may be exponentially many joint failure events. However, we construct a concise representation of the joint probability distribution of link failures, which under some practical assumptions has space complexity $O((n+x)\rho^3\gamma^4)$, where n is the number of nodes, x is the number of link crossings (in practice $x \ll n$), ρ represents a density of the topology, which is independent of the network size, and finally, γ stands for the maximum number of line segments a (polyline-shaped) link consists of.

- We provide proof-of-concept implementation and simulations based on real seismic hazard data and network topologies. Our simulations demonstrate how the above-mentioned stochastic model can be efficiently computed, even on commodity computers. This, extended with traditional random failure models, facilitates comprehensive service availability analysis considering disaster failures.

Fig. 1 summarizes the three layers of our contributions. There are two data structures on the left, analogous to CDF and PDF, which we believe should be the standard way of describing the joint failure probability of network resource sets. In the middle, the second layer is a stochastic model that explicitly considers the correlation between the failures of geographically close-by network elements. In the third layer, on the right, is the input to our framework, which might need to be pre-processed to fit the model. As a specific example, we show how to pre-process historical earthquake catalogs to provide proper input for our model. This way, we describe a method of computing PSRLGs of a network from end to end.

This paper is organized as follows: Sec. II presents the framework for computing service availability, Sec. III explains the stochastic model we use to represent regional failures. Sec. IV proposes an offline pre-computation process with performance guarantees. Sec. V extends the previously-defined link failure model to cope with arbitrary network resources, and Sec. VI provides theoretical bounds on the size and query time of the proposed data structures. Sec. VII demonstrates how the data structures can be pre-computed and queried efficiently. Sec. VIII provides a numerical evaluation of the proposed schemes based on seismic hazard data. Finally, Sec. IX concludes our work.

II. NETWORK MODEL AND FRAMEWORK TO COMPUTE SERVICE AVAILABILITY

A. Network Model

The network is modeled as an undirected connected geometric graph $G = (V, E)$, with $n = |V|$ nodes and $m = |E|$

(a) Framework to compute service availability

Data set name	Space complexity	Query time for an arbitrary link set
CFP[G]	$\Omega(2^\rho)$ and $O(2^\rho(n+x)\rho^3\gamma^4)$	hashing: <i>constant</i> with high prob. balanced binary tree: $O(\rho \log((n+x)\rho\gamma))$ worst-case
FP[G]	$O((n+x)\rho^3\gamma^4)$	$O((n+x)\rho^3\gamma^4)$

(b) Trade-off between space complexity and query time in case of circular dish shaped disasters

Fig. 2. Computing service availability via a pre-computed data set: while the disaster hazard can be represented more succinctly using FP[G] for a graph G , with CFP[G] one can achieve lower query times.

links embedded in $\mathbb{R}^2$. The links can be either line segments or polygonal chains (also called ‘polylines’) built up from at most γ adjacent line segments (where γ is a parameter of our model). The number of link crossings is denoted by x . The geometric *density* of the network topology is the maximum number of links that can be hit by a single disaster and is denoted by ρ . The set of links E is lexicographically sorted, any $S \subseteq E$ is stored as a sorted list. Note that our algorithms are mostly linear in the network size.

B. Framework to Compute Service Availability

We aim to develop a service availability computation engine, where the task is basically to translate the compound problem of simultaneous network failures into a scalar. When setting up an SLA between the user and network provider, the availability of a massive number of network services must be evaluated. Therefore, we need to avoid committing resource-intensive computations at every query. Intuitively, there is much redundancy in these queries. The main idea behind our general framework (depicted in Fig. 2a) is to exploit this redundancy by pre-computing some numerical integrals representing failure probabilities of sets of network elements. This, out of the compound geometric and stochastic problem, extracts all the relevant information to a static data set. This data set can address many service availability queries, each of which requiring only lookups and summation.

We propose two standard PSRLG definitions, with different meanings on the probabilities associated with the link sets, to store the failure probabilities of sets of network elements: (1) the Cumulative Failure Probability (CFP), and (2) the Link Failure State Probability (FP). While in this paper we focus on failure probabilities of link sets, if necessary, these structures can store failure probabilities of both links *and* node failures (see Sec. V on extensions of our basic model).

Definition 1 (Cumulative Failure Probability (CFP)): Given a set of links $S \subseteq E$, the cumulative failure probability (CFP) of S , denoted by $\text{CFP}(S)$, is the probability that all links S fail simultaneously (and possibly other links too).

Definition 2 (Link Failure State Probability (FP)): Given a set of links $S \subseteq E$, the link failure state probability (FP) of S , denoted by $\text{FP}(S)$, is the probability that *exactly* the links of S fail simultaneously (and no other links).

Sometimes we will refer as ‘CFP’ to 1) the tuple $(S, \text{CFP}(S))$ for a link set S , or simply, 2) to $\text{CFP}(S)$. For a graph G , we will denote the collection of CFPs with strictly positive probability by $\text{CFP}[G]$. The same applies to the Failure Probabilities (‘FP’s). We note that the reason behind

not referring the tuple of a link set S and $\text{CFP}(S)$ or $\text{FP}(S)$ simply as PSRLGs is that, throughout this paper, we need to make a distinction between these two data structures.

Although for some practical tasks, $\text{FP}[G]$ may be a practical input, in the standpoint of availability queries, we mainly look at $\text{FP}[G]$ as a compact representation of structure $\text{CFP}[G]$ (the space complexity of the proposed structures will be investigated in detail in Sec. VI).

The *space complexity* of our availability computation engine based on either CFPs or FPs is proportional to the number of link sets S with $\text{CFP}(S) > 0$ (resp., $\text{FP}(S) > 0$). The engine’s *time complexity* (namely, its query time) is the time needed to determine the cumulative failure probability of a given link set.

As it turns out, data structures $\text{CFP}[G]$ and $\text{FP}[G]$ present a space-time trade-off: There are more link sets with non-zero CFP than FP, since $\text{FP}(S) > 0$ implies that $\text{CFP}(S') > 0$ for all $2^{|S|}-1$ nonempty sets such that $S' \subseteq S$. On the other hand, availability queries need to address fewer PSRLGs if they are all expressed as CFPs, and computing these from FPs requires iterating over all FPs in the data set. In Sec. VI, we study this trade-off in more detail and give formal bounds on the space complexity and query time for both data structures (see Fig. 2b) when applied to our regional failure model.

C. On Availability Queries when Risk Failures are Correlated

Any availability query can be evaluated by iteratively calling $\text{CFP}(S)$, i.e., the probability of simultaneous failure of all elements in any arbitrary set S . Consider the example network and corresponding CFPs in Fig. 3 (non-listed link sets have CFPs of 0). Suppose we need to establish a high-availability connection from the top right node through a working path c and protection path $f-d-e$. The unavailability of the working path is $\text{CFP}(\{c\}) = 0.0113$, and the unavailability of the protection path is $\text{CFP}(\{f\}) + \text{CFP}(\{d\}) + \text{CFP}(\{e\}) - \text{CFP}(\{f, d\}) - \text{CFP}(\{f, e\}) - \text{CFP}(\{d, e\}) + \text{CFP}(\{f, d, e\}) \simeq 0.0275$, by the inclusion-exclusion principle. The total connection availability is $1 - \text{CFP}(\{c, d\}) - \text{CFP}(\{c, f\}) - \text{CFP}(\{c, e\}) + \text{CFP}(\{c, f, d\}) + \text{CFP}(\{c, f, e\}) + \text{CFP}(\{c, d, e\}) - \text{CFP}(\{c, f, d, e\}) \simeq 0.99872$. We can observe that, based on $\text{CFP}[G]$, the connection availability can be computed with the help of CFPs of subsets of $\{c, d, e, f\}$, that is, the union of the links of the working and protection paths.

In contrast, for computing the total connection availability, the $\text{FP}[G]$ data set requires considering a larger number of data set entries. For example, the availability of working path

Fig. 3. An illustration of the problem inputs and outputs. We note that the earthquake failure model depicted here, detailed in Sec. VIII-A, and used in our simulations, is a special case of our general model presented in Sec. III, that can handle a wide variety of disaster types (including tornadoes, tsunamis, etc.), possibly describing their combined effect.

c can be computed as is $1 - \sum_{\{c\} \subseteq S \subseteq \{a, \dots, e\}} FP(S)$, i.e., we have to subtract the FP of every link set containing c from 1. Furthermore, to compute the total availability of the connection, we need to address all nonempty subsets of $\{a, b, c, d, e\}$. The number of links is not part of neither the working nor the protection path; this means up to exponentially more $FP[G]$ queries than $CFP[G]$ queries. Structure $FP[G]$ has an advantage though: it has provably less elements than $CFP[G]$.

By considering joint failure probabilities, we have found that the total connection availability is < 0.9987 , i.e., below three nines. For comparison, traditional approaches that assume link failures to be independent, would have estimated the total connection availability to be $1 - CFP(\{c\})(CFP(\{d\}) + CFP(\{e\}) + CFP(\{f\}) - CFP(\{d\}) \cdot CFP(\{e\}) - CFP(\{d\}) \cdot CFP(\{f\}) - CFP(\{e\}) \cdot CFP(\{f\}) + CFP(\{d\}) \cdot CFP(\{e\}) \cdot CFP(\{f\})) > 0.99951$, i.e., well above three nines. Even if they correctly compute the availability of each path but assume independent path failures, they estimate the availability by $1 - 0.0113 \cdot 0.0275 > 0.99968$, i.e., even more above three nines. Here, by not considering joint failure probabilities, the traditional approaches significantly overestimate the total connection availability, which can lead to more frequent SLA violations and a financial burden on the CSP.

Unfortunately, (correlated) network failures are hard to compute and predict. Nonetheless, to evaluate the expected availability of a service, a network administrator should consider all possible failure scenarios under the specific service availability model stipulated in the corresponding SLA.

D. Denomination Issues of Probabilistic SRLGs

Probabilistic extensions of SRLGs are called *Probabilistic SRLGs*, PSRLGs. The probabilistic refinement can be defined in multiple ways, thus, in the literature, there are multiple definitions of PSRLGs. E.g., in the first paper considering probabilistic extensions SRLGs (which was [34]), each PSRLG event $r \in R$ occurs with probability π_r , and once a PSRLG event r occurs, link (i, j) will fail independently of

the other links with probability $p_{i,j}^r \in [0, 1]$. Thus, we could call the [34]-PSRLGs as 'two-stage PSRLGs'. In contrast with this paper, [34] does not tackle the issue of computing the PSRLGs.

Since both FPs and CFPs are probabilistic extensions of SRLGs, we say that, collectively, these structures are PSRLGs. Moreover, since any version of probabilistic SRLGs can be described with the help of either CFPs or FPs, and due to their natural simplicity, we believe (C)FPs are the right standard way of defining PSRLGs. In the following, we present a model for calculating $CFP[G]$ and $FP[G]$ describing the correlated failure patterns of networks.

III. THE REGIONAL FAILURE MODEL

To compute service availabilities, we need to answer the following question: **what is the probability that a set of links S fails simultaneously?** In other words, we need to find the *cumulative failure probability* of S , i.e., $CFP(S)$, which has a complicated relationship with the correlation structure of link failures. Links that lie close together more often fail simultaneously, while further apart links rarely do. To find $CFP(S)$, we first propose a general stochastic model of possible network failure events. After some pre-computation, this will allow us to build a succinct representation of the joint probability distribution of link failures described in the previous section.

In our model, failures are considered to come solely from disasters affecting a bounded geographical area. This section focuses only on link failures (node failures can be translated to the joint failure of the set of all links adjacent to the node). We extend our model to incorporate node failures as well in Sec. V.

While traditional approaches focus on single-point failures, which represent hardware/node failures, cable/link cuts, etc., we adopt a model for regional failures and focus on computing the conditional probability $CFP_d(S)$ that, in a given time period, a set of links S fail together under a disaster of type

d (e.g., a tornado, earthquake, Electromagnetic Pulse (EMP), etc.).

Assumption 1: We assume that, in the investigated time period, there will be at most one disaster of any type¹.

In such a case, to obtain the availability values, we need to build a model for each disaster type, and the resulting availability of S can be expressed as $1 - \sum_{d \in D} p_d \cdot \text{CFP}_d(S)$, where D denotes the set of modeled failure types and p_d is the probability of disaster d . From now on, for ease of notation, we will consider a fixed failure type d , and, therefore, the subscript d is omitted hereafter.

A. Stochastic Modeling of Regional Failures

In the remainder of the paper, we will call events that bring down the network in a geographic area simply as *disasters*, indifferent to their cause. We model regional failures caused by a disaster with 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 a **hazard** $h(p)$ representing the probability that p becomes the epicenter of the next disaster (see Fig. 4a). Specifically, $h(p)$ is a probability density function on the area $\mathbb{R}^2$, and therefore,

$$\int_{p \in \mathbb{R}^2} h(p) dp = 1. \quad (1)$$

After a disaster of the examined type, the physical infrastructure (such as optical fibers, amplifiers, routers, and switches) in some areas 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 $r(p, s)$.

Definition 3 (Regional disaster): We assume a *regional disaster* of epicenter p and shape/size s will result in the failure of exactly those links of network G that have a point in $r(p, s)$.

Our next assumption is that $r(p, s)$ is monotone increasing in the relative size s , that is, a more severe version of a disaster hits at least the same region of the network, as a weaker

disaster (see Fig. 4b for an example)². While this assumption holds in general for a variety of disasters, we only use it to achieve ‘nicer’ equations.

Assumption 2:

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

For simplicity, we assume $r(p, s)$ 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.

Note 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 $R(p, s)$. Definition 3 together with Assumption 2 imply 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 $s = 1$):

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

B. The Failure Probability of a Link Set

We first explain how to compute the probability $\text{CFP}(S)$ that a set of links $S \subseteq E$ will fail simultaneously in the next disaster.

Let $f(e, p)$ be the **probability** that link e fails if a disaster with epicenter p happens. Note that by Assumption 2, $f(e, p) > 0$ 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,

$$f(e, p) = \int_{s=0}^1 I_{R(p, s)}(e) ds, \quad (4)$$

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

$$I_{R(p, s)}(e) = \begin{cases} 1 & \text{if } e \in R(p, s), \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

By Assumption 2, if $I_{R(p, s)}(e) = 1$, then $I_{R(p, s')}(e) = 1$, for $s \leq s'$.

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

$$P(e) := \int_{p \in \mathbb{R}^2} h(p) f(e, p) dp. \quad (6)$$

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

$$f(S, p) := \int_{s=0}^1 \prod_{e \in S} I_{R(p, s)}(e) ds. \quad (7)$$

In other words, if the sequence of links is $S = (e_1, e_2, \dots, e_{|S|}) \subseteq R(p, 1)$ and $s(p, e_1) \leq s(p, e_2) \leq \dots \leq$

²Various failure shapes were studied so far [4], [8], [10]–[24], mainly in the form of circular regional disasters or line-segment failures, but in some cases also more general geometric shapes [4], [12]. All of these models meet Assumption 2.

(a) Probabilistic hazard map $h(p)$ for earthquakes as function of epicenter p . [38]

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

Fig. 4. Example of real-world inputs.

$s(p, e_{|S|})$, 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

$$f(S, p) = f(e_{|S|}, p) = \min_{e \in S} f(e, p) . \quad (8)$$

Finally, using the above results³:

$$\text{CFP}(S) = \int_{p \in \mathbb{R}^2} h(p) f(S, p) dp = \int_{p \in \mathbb{R}^2} h(p) \min_{e \in S} f(e, p) dp . \quad (9)$$

For example, on the right of Fig. 3, 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 with a radius computed according to the historical seismic information. Potentially there are exponentially many joint failure events in terms of the network size; however, links far from each other have zero probability of failing jointly because of a single disaster. For example, this holds for links f and b , whose smallest distance is more than the radius of the largest destroyed area.

Former works (e.g., [4, in 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 independent. Unlike [4], our model assumes a deterministic failure outcome (once its epicenter and shape are set). This implies that, in our 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).

C. Example of the Geographical Correlation of Failures

In this section, we first consider a simple linear and discrete model for some of the ideas presented so far. We assume that the ground set of our simplified world is the set of 1000 integer points of a line with coordinates between $z_{\min} = -499$, $z_{\max} = 500$ and we have two links e_0 and e_z , which themselves are integer points from the interval $[-499, 500]$, e_0 is at position 0, and e_z is at position z . Let the probability that i is the location of a disaster be $h_i = 10^{-3}$ for $i = -499, \dots, 500$ so that $\sum_{i=-499}^{500} h_i = 1$. According to Eq. (9), the probability of the failure of link e_0 is

$$P(e_0) := \sum_{i=-499}^{500} h_i f(e_0, i) , \quad (10)$$

where $f(e_0, i)$ is the conditional probability that link e_0 fails if the failure is at position i . According to Eq. (9), the joint probability of the failure of both links e_0 and e_z is

$$P(\{e_0, e_z\}) := \sum_{i=-499}^{500} h_i \min(f(e_0, i), f(e_z, i)) . \quad (11)$$

Let $P(e_z|e_0)$ denote the conditional probability that e_z fails, on the condition that e_0 fails. By definition we have

$$P(e_z|e_0) := \frac{P(\{e_0, e_z\})}{P(e_0)} . \quad (12)$$

³Without Assumption 2, we would have $\text{CFP}(S) = \int_{p \in \mathbb{R}^2} h(p) \int_{s=0}^1 \prod_{e \in S} I_{R(p,s)}(e) ds dp$.

Fig. 5. An example of $f_i(0)$ at different i positions and the corresponding $P(e_z|e_0)$ depending on z . Former models assumed the link failures are independent given an epicenter of the disaster.

This is a function of z in our setting. Intuitively, $P(e_z|e_0)$ is close to 1 if the two links are exactly in the same location (i.e. $z = 0$).

Additionally, $P(e_z|e_0)$ should be a decreasing function of z in the range of $[0, 500]$. See Fig. 5 for an example of $f(e_0, i)$ values and the corresponding $P(e_z|e_0)$.

IV. PRE-COMPUTATION TO SPEED UP QUERIES

In the previous section, we have described a model that generates a regional disaster according to a hazard density $h(p)$ and a failure shape function $r(p, s)$. Recall that our task is to return $\text{CFP}(S)$ for a set of links $S \subseteq E$, which is the probability that links S fail together in case of disaster d .

Unfortunately, the calculation of integrals (9) can be a computationally-intensive process. One solution is to calculate some FPs in advance so that when a query comes on the CFP of an arbitrary set of links S , then the task would be summing up some of the pre-computed FP values.

As Lemma 1 will show, a full list of FPs with non-zero probabilities has $O((n+x)\rho^2\gamma^4)$ items. Every CFP can be derived by summing up

$$\text{CFP}(S) = \sum_{T \supseteq S} \text{FP}(T), \quad \forall S \subseteq E. \quad (13)$$

A. Precomputation of CFPs and FPs

In this subsection, we still rely on Assumption 2 and make the following additional assumptions to apply some computational geometry results. We emphasize that additional specifications 2) and 3) are technical assumptions to avoid lengthy discussions (see the Appendix).

- 1) The shapes $r(p, s)$ are limited to circular disks centered at p . This corresponds to the case where the failure of a link e depends on the Euclidean distance $\text{dist}(p, e)$ of e to the epicenter p of the disaster. In this case, instead of $r(p, s)$, the input is given by radius d as a function of s .
- 2) In our geometric reasoning, we will transform the links of the graph into line segments by slightly shortening them to ensure that no two segments share a common endpoint (see the details of the transformation in Appendix A). We also assume that no more than two links intersect in the same point, and no more than two endpoints lie on the same line.

- 3) The relative size s is a *uniformly Lipschitz continuous function* of radius d . That is, there exists a positive number K such that for every point p in the plane, if we have neighborhoods $r(p, s')$ and $r(p, s)$ with respective radii d' and d , then $|s' - s| \leq K|d' - d|$ holds.

For ease of presentation, we slightly reduce the domain we are integrating over. Let $\mathcal{P}$ denote the set of points p of the plane such that $\text{dist}(p, e) \neq \text{dist}(p, e')$ whenever e and e' are different segments from E . We have that $\mathbb{R}^2 \setminus \mathcal{P}$ is of measure zero, hence in our considerations, integrating over the plane $\mathbb{R}^2$ can be replaced by integrating over $\mathcal{P}$.

Inspired by (8), we can now define the sequence of possible link failures (see Fig. 6), when the epicenter of the disaster is at p :

Definition 4: The *sequence of link failures* for epicenter $p \in \mathcal{P}$ is an ordered set of links $\mathcal{S}(p) = (e_1, e_2, \dots, e_l)$, such that $s(p, e_1) \leq s(p, e_2) \leq \dots \leq s(p, e_l)$, where $l = |R(p, 1)|$. Let $\mathcal{S}^j(p)$ denote the first j links of $\mathcal{S}(p)$, i.e. $\mathcal{S}^j(p) = (e_1, e_2, \dots, e_j)$.

Furthermore, the ordinal number of a set S within $\mathcal{S}(p)$ is defined as follows:

Definition 5:

$$j(S, \mathcal{S}(p)) = \begin{cases} i, & \text{if } S \not\subseteq \mathcal{S}^{i-1}(p) \text{ and } S \subseteq \mathcal{S}^i(p) \\ 0, & \text{otherwise.} \end{cases}$$

Due to Assumption 2 and using also (9), if there is a disaster at point p , the conditional probability of a set of links $S \subseteq \mathcal{S}(p)$ failing together is

$$f(S, p) = f(\mathcal{S}^{j(S, \mathcal{S}(p))}(p), p) = f(e_{j(S, \mathcal{S}(p))}, p). \quad (14)$$

Finally, we use two practical input parameters, x , and ρ , in estimating the space complexity of our approaches. Let x be the number of link crossings in the network G . For backbone networks, x is a small number, as typically, a switch is also installed on each link crossing [39]. The second parameter is ρ , the *link density* of the network, which is defined as the maximal number of links that could fail together (i.e., could be covered by a circle of radius r). The link density ρ , practically, does not depend on the network size. Moreover, ρ is at least the maximal nodal degree in the graph.

Let us divide the plane into disjoint regions $\mathcal{R}_1, \dots, \mathcal{R}_k$, where each point $p \in \mathcal{R}_i$ has the same sequence $\mathcal{S}_i$ of link failures (see Fig. 7, [1] for a more detailed discussion, and [40] for efficient algorithms calculating these regions). Here, k is the number of possible failure sequences. For any point $p \in \mathcal{R}_i$, we introduce notation $\mathcal{S}(p) \equiv \mathcal{S}_i$, $i = 1, \dots, k$.

(a) The sequence of link failures for epicenter p . (b) Bisector curve of e_1 and e_2 , is the boundary of areas with same sequences of link failures.

Fig. 6. Illustration of link failure sequences

(a) Regions with same sequence of link failures. (b) Nearest Neighbor Voronoi Diagram. (c) 2-Voronoi diagram.

Fig. 7. An example of different partitions of the plane into regions used in Lemma 1.

Based on Equation (14), it is sufficient to pre-compute and store the following integrals:

$$P^{i,j} = \int_{p \in \mathcal{R}_i} h(p) f(e_{i,j}, p) dp \quad i = 1, \dots, k, \quad j = 1, \dots, |\mathcal{S}_i|, \quad (15)$$

where $e_{i,j}$ denotes the j -th link in $\mathcal{S}_i$.

Finally, since the regions are mutually disjoint as subsets of $\mathcal{P}$ and cover it entirely, equation (9) can be written as a sum and, according to (14), the failure probability of any link set $S \subseteq E$ can be evaluated as

$$\text{CFP}(S) = \sum_{i=1}^k \int_{p \in \mathcal{R}_i} h(p) f(S, p) dp = \sum_{i=1}^k P^{i,j(S, \mathcal{S}_i)}, \quad (16)$$

where we define $P^{i,0} := 0$ for every $i = 1, \dots, k$. Based on Eq. (13) and (16), one can derive that:

$$\text{FP}(S) = \sum_{i,j} (P^{i,j} - P^{i,j+1}), \quad (17)$$

where the summation is for those pairs (i, j) for which $1 \leq i \leq k$ and $j(S, \mathcal{S}_i) = |S| > 0$. As a default, we set $P^{i, |\mathcal{S}_i|+1} = 0$.

V. MODEL EXTENSIONS

A. Different Link Types

Most optical backbone networks consist of multiple types of links, e.g. *aerial*, *buried* and *submarine*. In case of a disaster, these link types have different failure patterns. For example, in case of an earthquake, the failure regions of aerial cables can be different from the regions for buried cables, while submarine cables tend to be cut at rupture zones. With this in mind, we extend our model as follows. Let L be the set of different link types. For each link type l , disaster zone $r(p, s, l)$ denotes the area where links with type l fail in case of a disaster with epicenter p and relative size s .

In this extension, Assumption 2 ($r(p, s)$ is monotone increasing in relative size s) translates to the following:

$$r(p, s, l) \subseteq r(p, s', l) \quad \text{if } s < s' \quad \forall p \in \mathbb{R}^2, 0 \leq s, s' \leq 1, l \in L. \quad (18)$$

Although their failure regions may differ, this extension still allows links of multiple types to fail due to a single disaster, analogously to many natural settings.

B. Mixed Link Types

Taking the previous extension a step further, we introduce the concept of mixed types. One can imagine that some links

may consist of different “link types”. For example, a link that is mainly buried may need to cross a river above-water. We implement these links by dividing each link into sections with homogeneous types. If a single section fails, the whole link fails. More formally, each link $e \in E$ is partitioned to sections $e^1, \dots, e^M$ with types $l^1, \dots, l^M$, respectively. Section e^i fails if it has a common point with $r(p, s, l^i)$, and link e fails if at least one of its sections fails.

C. Nodes Also Considered Vulnerable

Network nodes have different failure patterns than links, and their probabilistic failures can be represented by PSRLGs as follows. For a node, $v \in V$ that can fail, the edges incident to v have mixed link types, and in a small vicinity of v are considered to have a type $l_v \in L$ specific to the node such that those parts of the links fail exactly then when the node would have failed. This approach translates to CFPs or FPs as follows: the set S of links incident to v fails because the disaster hits every $l \in S$ or the disaster hits node v .⁴

VI. SPACE AND TIME COMPLEXITY OF STRUCTURES CFP[G] AND FP[G]

A. Cardinality of Structures FP[G] and CFP[G]

In our basic model, considering the case of the disaster shapes being circular disks in a given L_p metric, (where, for $p = 2$, we get back the usual Euclidean circles, for $p = 1$ or $p = \infty$, we have a family of parallel-sided squares, and, for $p = 2/3$, astroids, that are specific 4-cornered stars), the number of FPs can be upper bounded as follows.

Lemma 1: In case of a set of circular disk shaped disasters (i.e., $r(p, s)$ is circular) in a given L_p metric, and the edges of the network being in general position,⁵ there are $O((n + x)\rho^2\gamma^4)$ FPs with non-zero probability.

Proof: Let us concentrate on line segment links for a moment. According to [24, Claim 2], the number of links, m , is $O(n + x)$ for line segment links. We know from [41, Thm. 6] that the number of k -Voronoi cells in L_p norm 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 $O(\rho^2(n + x + x))$, which is $O(\rho^2(n + x))$.

If links can be polygonal chains consisting of at most γ line segments, there are $O(\gamma(n + x))$ segments with $O(\gamma^2x)$ crossings, meaning $O(k\gamma^2(n + x))$ k -Voronoi regions. By counting the k -Voronoi regions for $k \in \{1, \dots, \gamma\rho\}$, this yields an upper bound of $O((n + x)\rho^2\gamma^4)$ for the number of FPs. ■

In the same setting, the number of CFPs can be very large:

Lemma 2: The number of CFPs with non-zero probabilities is lower-bounded by $\Omega(2^\rho)$. In case of a set of circular disk shaped disasters in a given L_p metric, and the edges of the

network being in general position, the number of CFPs with non-zero probabilities is upper-bounded by $O(2^\rho(n + x)\rho^2\gamma^4)$.

Proof: By the definition of ρ , there is a link set S with $\text{CFP}(S) > 0$ and $|S| = \rho$. As, for any $S' \subseteq S$, $\text{CFP}(S) > 0$ implies $\text{CFP}(S') > 0$, implying the lower bound. By Lemma 1, there are at most $O((n + x)\rho^2\gamma^4)$ non-zero FPs, each having at most 2^ρ subsets, yielding the upper bound. ■

Every FP and CFP can be stored in $O(\rho)$ space, since it contains a link set of at most ρ links, alongside with a related probability. This way, the space requirement of FP[G] and CFP[G] is upper bounded by $O((n + x)\rho^3\gamma^4)$ and $O(2^\rho(n + x)\rho^3\gamma^4)$, respectively.

B. Query Time of Structures FP[G] and CFP[G]

When storing the non-zero FPs in a list, by Eq. (13), querying the FP[G] structure for $\text{CFP}(S)$ requires iterating over all non-zero FPs and summing up all $\text{FP}(T)$ such that $T \supseteq S$. Thus, S has to be compared with $O((n + x)\rho^2\gamma^4)$ (Lemma 1) other sets, and each comparison can be made in $O(\rho)$. The number of possible additions is also $O((n + x)\rho^2\gamma^4)$, thus the query time of the FP[G] structure is upper-bounded by $O((n + x)\rho^3\gamma^4)$. Alternatively, if we stored the FPs in an ordered balanced binary tree, we would need to lookup all the exponential number of $T \supseteq S$.

The query time of CFP[G] also depends on the data structure used for storing CFPs. For example, if we store all non-zero CFPs in a list, the query time would be $\Omega(2^\rho)$ (Lemma 2). In contrast, by hashing all $\text{CFP}(S)$ on S , we reduce the query time a constant with very high probability. Last, when storing all non-zero CFPs in a self-balancing binary tree, the worst-case query time would be $O(\rho + \log((n + x)\rho\gamma))$ (Lemma 2). Although the CFP structure can achieve impressive query times, this comes at the cost of its space complexity ($\Omega(2^\rho)$), which makes it inefficient for larger network densities.

VII. IMPLEMENTATION ISSUES

The approaches and performance guarantees we gave in Sections IV and VI are valid under the assumption that the shape of a regional failure is always a circular disk. In this section, we propose a heuristic that (1) can accommodate any disaster shape; (2) does not require advanced geometric algorithms; and (3) is more suitable for digital inputs, as it uses discrete functions instead of continuous ones.

We discretize the problem by defining a sufficiently fine grid over the plane such that for each grid cell c , the disaster regions $r(p, s)$ and hit link sets $R(p, s)$ are “almost identical”⁶ for all $p \in c$. This reduces the integration problem from Sec. III to a summation⁷.

We consider $\mathbb{R}^2$ as a Cartesian coordinate system. Let r denote the absolute maximum range of a disaster in km. Let $(x_{\min}, y_{\min})$ be the bottom left corner and $(x_{\max}, y_{\max})$ the top right corner of a rectangular area in which the network lies. It is sufficient to process each c in the rectangle of bottom left corner $(x_{\min} - r, y_{\min} - r)$ and top right corner $(x_{\max} +$

⁴Another possibility is to handle node failures natively, and assume the failure of a node v infers the failure of the links incident to v .

⁵According to the general position assumption, there are no more than three segments touch the same circle and no more than two endpoints lie on the same line. If this assumption is not met, the coordinates of the network could be perturbed.

⁶In particular, we may assume that $f(e, p)$ is independent of p as long as it is in c and denote this common value by $f(e, c)$.

⁷[20] uses a similar grid approach.

$r, y_{max} + r)$, and we denote by $c_{i,j}$ the grid cell in the i -th column and j -th row of this rectangle. We assume we are given the probability $h_{i,j}$ of the next disaster epicenter p lying in cell c : $h_{i,j} = \int_{p \in c_{i,j}} h(p) dp$.

Now, for each c , we can compute the sequence of link failures and store the link sets as follows.

1) *Structure CFP[G]*: For our CFP[G] structure, we use an associative array CFP[G], which can be addressed by a set of links $S = \{\ell_1, \ell_2, \dots, \ell_k\}$ and returns its cumulative failure probability. In the pre-computation process, we have to extract the contribution of $c_{i,j}$ to the failure probability of every subset S of links. To do so, we process the sequence of link failures $\mathcal{S}_{i,j} = (e_1, e_2, \dots, e_l)$ attached to disaster epicenters which are in $c_{i,j}$ ⁸, and increment the CFP[G] values accordingly: $\text{CFP}(\{e_1\})+ = h_{i,j} \cdot f(e_1, c_{i,j})$, $\text{CFP}(\{e_2\})+ = h_{i,j} \cdot f(e_2, c_{i,j})$, $\text{CFP}(\{e_1, e_2\})+ = h_{i,j} \cdot f(e_2, c_{i,j})$, etc. By default, for every link set S , we set initially $\text{CFP}(S) = 0$.

To obtain $\text{CFP}(S)$, we look it up in the associative array. If S is not found, then $\text{CFP}(S) = 0$.

2) *Structure FP[G]*: For our FP[G] structure, we take a similar approach as for the CFP[G] structure and use a list of ' $S, \text{FP}(S)$ ' set-failure probability pairs.

In the pre-computation process, we have to extract the contribution of $c_{i,j}$ to the link failure state probability of every subset S of links. As in the case of the CFPs, we do so by iterating over the sequence of link failures $\mathcal{S}_{i,j} = (e_1, e_2, \dots, e_l)$ and incrementing the FP values accordingly: $\text{FP}(\{e_1\})+ = h_{i,j} \cdot (f(e_1, c_{i,j}) - f(e_2, c_{i,j}))$, $\text{FP}(\{e_1, e_2\})+ = h_{i,j} \cdot (f(e_2, c_{i,j}) - f(e_3, c_{i,j}))$, $\text{FP}(\{e_1, e_2, e_3\})+ = h_{i,j} \cdot (f(e_3, c_{i,j}) - f(e_4, c_{i,j}))$, etc.

To obtain $\text{CFP}(S)$, we sum up $\sum_{T \supseteq S} \text{FP}(T)$.

VIII. MODEL EVALUATION BASED ON SEISMIC HAZARD DATA

In this section, we present numerical results that validate our model and demonstrate the use of the proposed algorithms on real backbone networks (taken from [42] and [2], resp.) accompanied with real seismic hazard inputs. The algorithms were implemented in Python 3.6., using its various libraries⁹, respecting the regional failure model presented in Section III, and following the implementation principles of Section VII. Run-times were measured on a commodity laptop with a Core i5 CPU at 2.3 GHz with 8 GiB of RAM.

As a practical scenario, the simulations presented in this paper focus on transforming the seismic hazard on network topologies to PSRLGs. For a more general proof-of-concept evaluation, we refer the reader to the conference version of our paper [1]. There, we assumed that the epicenter distribution is uniform over the investigated area. The disasters shape is a circular disk with a maximal radius r (at $s = 1$), which is constant over the region.

As a first step, we need to convert the historical seismic hazard data into a regional failure model for our framework. Subsec. VIII-A discusses our earthquake representation, based

on epicenter and moment magnitude. In a nutshell, the model translates the seismic hazard data to a set of circular disk shaped disaster areas with radii depending on the actual moment magnitude (Fig. 8). Note that the epicenter distribution is non-uniform here.

We are taking this probabilistic earthquake set as input, Subsec. VIII-B presents our simulation results validating our PSRLG model.

A. Seismic Hazard Representation

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 period. Each earthquake is uniquely identified by its epicenter and moment magnitude [44]:

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.

moment magnitude $M_w \in \{4.6, 4.7, \dots, 8.6\} =: \mathcal{M}$.¹⁰

We index the grid cells such that $i \in \{1, \dots, i_{max}\} =: \mathcal{I}_i$, $j \in \{1, \dots, j_{max}\} =: \mathcal{I}_j$.

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 small enough that the failures caused by each earthquake in E_{i,j,M_w} will often be identical¹¹, 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. $\sum_{i,j \in \mathcal{I}_i \times \mathcal{I}_j} \sum_{M_w \in \mathcal{M}} p_{i,j,M_w} = 1$.

Our initial input are the activity rates r_{i,j,M_w} of earthquake types (see Fig. 8a) instead of the p_{i,j,M_w} values, so we first have to translate these rates to probabilities. We claim that under the assumption that each kind of earthquake E_{i,j,M_w} arrives according to independent Poisson arrival processes with parameters r_{i,j,M_w} , the rates of earthquakes E_{i,j,M_w} can be transformed to probabilities p_{i,j,M_w} as follows:

$$p_{i,j,M_w} = r_{i,j,M_w} / \sum_{i,j \in \mathcal{I}_i \times \mathcal{I}_j} \sum_{M_w \in \mathcal{M}} r_{i,j,M_w}. \quad (19)$$

We assign each network element e an **intensity threshold** $t(e)$. 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 $t \in \{\text{VI}, \text{VII}, \text{VIII}, \text{IX}, \text{X}, \text{XI}, \text{XII}\} := T$ according to the Mercalli-Cancani-Sieberg (MCS) scale [45]¹².

After each earthquake, E_{i,j,M_w} , the physical infrastructure (such as optical fibers, amplifiers, routers, and switches) in an area $\text{disk}(c_{i,j}, R(M_w, t))$ of a circular disk is destroyed. The center point of $\text{disk}(c_{i,j}, R(M_w, t))$ is the center of $c_{i,j}$, while

¹⁰ $M_w \leq 4.5$ means no damage, while $M_w > 8.6$ has not been experienced in the studied regions.

¹¹The 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.

¹²Intensity $I \leq \text{V}$ does not cause structural damage, while $I = \text{XII}$ means total damage.

⁸Here, we represent $c_{i,j}$ by its center p . According to Def. 4, for $i < j$, link e_i is closer to p than e_j , i.e., $s(p, e_i) < s(p, e_j)$.

⁹The simulation data can be downloaded from [42].

its radius $R(M_w, t)$ is monotone increasing in the magnitude M_w , and decreasing in the intensity threshold t (see Fig. 8b and 8c). 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 an earthquake's damaged area in cell $c_{i,j}$ with magnitude M_w .

1) *Earthquake Activity Rates*: These are the occurrence rates of earthquake events as a function of space, time, and magnitude. To obtain them, we need to define an earthquake source model, defined as an area or an active fault that could host earthquakes as testified by instrumental seismic activity, historical seismicity, geomorphological evidence, and regional tectonics. The choice of the earthquake source model is strongly driven by the available knowledge of the area and by the scale of the problem. It may range from well-defined active faults, especially when working at a local scale, to less understood and wider scale seismotectonic provinces. When the catalog of earthquakes covers a long period, it can be used to compute earthquake activity rates without any information of seismotectonic provinces and/or active faults, via, for example, a smoothed seismicity approach. In this work, we evaluated the earthquake source model for Italy and the USA from the most recent published earthquakes catalogs ([43], and [46], for Italy and the USA, respectively) that cover a long period and can be used to obtain earthquake source model without other information. Although earthquakes can be clustered in time and space with their distribution that is over-dispersed if compared to the Poisson law [47], a common way to treat this problem (i.e., cluster in time and space) is to de-cluster the earthquake catalog, i.e., removing all events not considered mainshocks, via a declustering filter [48]. Here, both catalogs are considered de-clustered. The standard methodology to estimate the density of seismicity in a grid, and used in this work, is the one developed by [49]. The smoothed rate of events in each cell is determined as follows:

$$Sr_i = \frac{\sum_j r_j \exp\left(\frac{-d^2(c_i, c_j)}{d_c^2}\right)}{\sum_j \exp\left(\frac{-d^2(c_i, c_j)}{d_c^2}\right)}, \quad (20)$$

where r_j is the cumulative rate of events with magnitudes greater than the completeness magnitude M_c in each cell

c_i of the grid and computed from the historical catalogue of earthquakes, $d(c_i, c_j)$ is the distance between the centers of grid cells c_i and c_j . The parameter d_c is the correlation distance (for Italy, 30km [50] and for the USA, 75km [51]). Then, the earthquake activity rates for each node of the grid are computed following the Truncated Gutenberg-Richter model [52]:

$$\lambda(M) = \lambda_0 \frac{\exp(-\beta M) - \exp(-\beta M_u)}{\exp(-\beta M_0) - \exp(-\beta M_u)} \quad (21)$$

for all magnitudes M between M_0 (lower or minimum magnitude) and M_u (upper or maximum magnitude); otherwise $\lambda(M)$ is 0. The upper and lower magnitude bounds represent, respectively, the maximum magnitude, or the largest earthquake considered for a particular source model, which depends on the regional tectonic context (in our case, M_w is at most 8.1, 8.6 and 8.3 for Italy, Europe, and the US, respectively), and the minimum magnitude, or threshold value, below which there is no engineering interest or lack of data (in this study, $M_w > 4.5$)¹³. Additionally, λ_0 is the smoothed rate Sr_i of earthquakes at $M_w = 4.5$ and $\beta = \ln(10)$, where b is the b -value of the magnitude-frequency distribution. For Italy, we calculated the b -value of the distribution on a regional basis using the maximum-likelihood method from [53], while for the USA, it comes from [46]. While for Italy and the USA, we computed the earthquake rates (Fig. 8a) following this approach and with the referenced data, for Europe, we used the already published SEIFA model ([54], and [55]), a kernel-smoothed, zonation-free stochastic earthquake rate model that considers seismicity and accumulated fault moment. In this model, activity rates are based on the SHARE European Earthquake Catalogue frequency-magnitude distribution model. The spatial distribution of model rates depends on the density distributions of earthquakes and fault slip rates. A magnitude-frequency distribution indicates the probability that an earthquake of a size within the upper and lower bound of the distribution may occur anywhere inside the source during a specified period.

While this does give us the rates for all combinations of epicenters and magnitudes for Italy, the USA, and Europe (Fig.

¹³Fig. 8a shows that, in the investigated range of magnitudes, the global rate of earthquakes dips exponentially in the function of the magnitude.

Fig. 8. Seismic input data.

8a), we still need the relation between magnitude and disaster area to be able to apply these rates to the network resiliency models.

2) *The Radius of the Damaged Zone*: The only earthquake effect that can be quantified at the scale of the whole country is ground shaking because quantifying any other earthquake effects requires a site investigation. Shaking intensity is localized and is generally diminishing with distance from the earthquake's epicenter. At the scale of a whole country, we can assume that soil and topographic conditions are relatively homogeneous. The seismic intensity only depends on the distance from the earthquake epicenter.

Here, we assume all links (and nodes) inside the area with a given MCS intensity $I \geq t$ (where $t \geq \text{VI}$) are damaged, while all components outside of this area remain functioning. Thus, to obtain all disaster areas, we now only need the disaster area radius for each magnitude $M_w \in \{4.6, 4.7, \dots, 8.6\}$. For this purpose, we used the intensity prediction equation of [56] and [57], for Italy/Europe, and the USA, respectively, where the expected intensity I at a site located at epicentral distance R is:

$$I_{\text{It,EU}} = 1.621 \cdot M_w - 1.343 - 0.0086(D-h) - 1.037(\ln D - \ln h), \quad (22)$$

$$I_{\text{US}} = 0.44 + 1.70 \cdot M_w - 0.0048 \cdot D - 2.73 \cdot \log_{10} D, \quad (23)$$

where $D = \sqrt{R^2 + h^2}$ is a sort of hypocentral distance, and h represents the hypocentral depth, which may be viewed as the average depth of the apparent radiating source [56], h equaling 3.91km and 10km for Italy/Europe and the USA, respectively. In this way, it is possible to compute for each M_w and intensity threshold t the site-distance $R(M_w, t)$ from the epicenter of the desired intensity threshold level. It is worth noting that Eq. (22) has been obtained using only the Italian earthquake historical catalog, and so it is not entirely correct to use it for the entirety of Europe. However, the Italian catalog is one of the more complete catalogs in Europe. There is no similar equation in the literature for the entire continent (to the best of our knowledge), and its development is beyond the paper's scope. We assume that the application of Eq. (22), as a first approximation, can be considered correct for entire Europe.

B. Simulation results

We consider seven topologies: one Italian topology, three other European topologies, and another three US topologies. Unless otherwise stated, we set the intensity tolerance threshold, t , to VI according to the MCS scale. The node and link

counts, as well as the number of CFPs and FPs with non-zero probability, of all topologies are given in Table I both for $t = \text{VI}$ and $t = \text{VII}$.

Interestingly, although the US network has slightly more nodes and links than the Italian network, it has much less CFPs (946 compared to 12106). This difference is easily explainable when we consider our theoretical results from Sec. VI: the number of non-zero CFPs is lower-bounded by $\Omega(2^\rho)$ (Lemma 2), which means an exponential growth with the maximal number of hit links, ρ . Since the Italian network has much shorter links than the American network, its hit link sets tend to be larger. We can observe this same exponential increase with the maximal number of hit links when we decrease the threshold from $t = \text{VII}$ to $t = \text{VI}$. For example, the number of CFPs of NFSNET is 1762 at $t = \text{VII}$, but explodes to 14199 if we decrease this threshold to $t = \text{VI}$. In contrast, the number of FPs makes a much smaller jump, from 523 to 969.

By only storing the x largest CFPs, we can trade in some precision in exchange for a significant reduction in memory usage. Fig. 9a shows the precision of this approach versus x . For the Italian topology, the highest probability among the omitted edge sets is 5.4×10^{-4} or 1.7×10^{-5} if we store only the top 100 or 1000 CFPs respectively. Furthermore, increasing the precision by order of magnitude requires only a bit more than an order of magnitude more CFPs. Similarly, in the case of the other networks, storing the first 100 or 1000 CFPs means that the highest probability among the omitted edge sets is below 5×10^{-4} or 1×10^{-5} , respectively; and increasing the number of CFPs by order of magnitude is more than enough for increasing the precision by a factor of 10.

Speaking of the precision-memory trade, omitting some of the FPs is also possible. In this case, the imprecision in the value of CFP(S) for some S can be upper bounded by the sum of probabilities stored in the omitted FPs. On Fig. 9b, we can see the probability assigned to the x^{th} most probable FP. Fortunately, the highest number of non-zero FPs was low, 969 in our experience, meaning that, most probably, no omission is needed.

As mentioned before, the difference in the number of non-zero CFPs can partly be explained by a difference in hit link set sizes. Fig. 9c shows the maximal number of hit links, ρ , versus the intensity threshold, t . We can confirm that, at $t = \text{VI}$, the Italian network has a much higher density than the US network (13 compared to 7).

We have also investigated the average CFP of a set of links with given cardinality. Fig. 9d shows the average failure

TABLE I
THE INVESTIGATED NETWORK TOPOLOGIES

Network name	n	m	# CFPs at $t = \text{VI}$	# FPs at $t = \text{VI}$	# CFPs at $t = \text{VII}$	# FPs at $t = \text{VII}$
Optic EU	22	45	6377	202	1369	135
Italian	25	34	12106	308	676	200
US	26	43	946	246	260	164
Nobel EU	28	41	3867	149	680	94
EU	37	57	5634	212	745	133
N.-American	39	61	2024	394	556	257
NFSNET	79	108	14199	969	1762	523

Fig. 9. The space and complexity of the data structures for the examined network topologies.

probability concerning the number of links failing together. Single links have an average failure probability between $[4.2 \times 10^{-4}, 2.1 \times 10^{-3}]$, depending on the network topology. The average failure probability for double and triple link failures lies in $[1.2 \times 10^{-5}, 3.9 \times 10^{-4}]$ and $[1.9 \times 10^{-6}, 9.3 \times 10^{-5}]$, respectively. These averages meet our expectations that the correlation between link failures is significant. By our observations, the combination of link failures with the highest CFPs is predominantly the combined failure of links incident to a single node.

Fig. 10 further investigates the relationship between the space requirements of CFP[G] and FP[G]. In Fig. 10a, we show the space requirement of structures CFP[G] and FP[G] as a function of the intensity threshold t . As expected, the number of CFPs drops quickly with the intensity threshold. Our results show that, especially at lower thresholds, choosing the FP structure can significantly reduce space requirements. This phenomenon is even stronger in case of Italy_995, a network with 32 nodes and 70 links over Italy, that we decided to exclude from most of the simulation presentations. The reason for this is its unusually high density: at intensity tolerances of $t = VI$ and $\rho = VII$, it has densities $\rho_{VI} = 31$ and $\rho_{VII} = 19$, yielding $> 10^9$ and 1153294 CFPs, while the number of its FPs is only 2011 and 1090, respectively.

Fig. 10b depicts the number of CFPs and FPs with given cardinality for the Italian. Since there is a link set of cardinality 13 with positive FP, there must be over 1700 subsets of cardinality 6 with non-zero CFP. In comparison, the number of FPs peaks at 71 for cardinality 4.

Continuing our study of the cardinality of failed link sets, Fig. 11a investigates the dependency between CFP(S) and $|S|$ in detail, for $|S| = 1, 2$ and 3. There are 34 single link failures in the Italian network whose CFPs range between $[0.0003, 0.019]$; it has 205 dual link failures with non-zero probabilities between $[7 \times 10^{-8}, 0.0037]$, and there is a number of 648 triple link failures with strictly positive probabilities, ranging between $[7 \times 10^{-8}, 0.0019]$. Here we can see that some CFPs with size l are less probable than some other CFPs containing $l + 1$ links. Thus, only storing CFPs with at most l links rarely yields the same result as only storing the most probable CFPs. Also, we can observe that the CFPs of the most probable triple link sets are not much smaller than the

CFPs of the most probable link pairs. This is another sign that the most probable double and triple link failures are failures of the links incident to the same network node.

Finally, to explore if our data structures behave differently in random-shaped disasters, we have set up the following simulation. For Italian, we took the same grid as we used for the earthquakes and generated four disasters for each grid point by:

- picking between 3 and 6 points on the unit circle, from the uniform distribution,
- connecting these points to form a simple polygon A ,
- choosing 4 random radii $r_1, \dots, r_4$ between 0 and 150km (again, uniformly distributed),
- and, finally, for each r_i , scaling and shifting A such that its circumcircle to have the grid point as center, and a radius of r_i .

With this setting, we found 1401 CFPs and 327 FPs. These results are similar to what we saw for the topology in case of earthquakes and $t = VI$. The difference is that we have slightly more FPs and less CFPs. There are more FPs because, in some areas, the random-shape disaster has larger extensions than the maximum local earthquake disaster regions. The number of CFPs is less because the maximal random shapes in this simulation were smaller than the largest earthquake disaster areas. Thus, the largest hit link sets were smaller in the random scenario, and the smaller ρ translated to fewer CFPs. The charts generated on these structures were very similar to those seen on Fig. 9-11, thus we omit them due to the lack of space.

IX. CONCLUSION

In this paper, we 1) introduced a unified terminology for Probabilistic Shared Risk (Link) Groups, 2) proposed a general stochastic model of regional failures of elements (nodes and links) of the physical network, and finally, 3) evaluated the model after carefully processing raw seismic hazard data. The pre-computation of the proposed PSRLGs is performed offline during network planning by computing numerical integrals using information about the disasters' effects and network equipment resistance to the catastrophes. As a result of the pre-computation, the probability of each set of links' joint failure is stored as *cumulative failure probabilities*. Alternatively, we propose a more space-efficient data structure that

(a) Average CFP/FP number vs. intensity threshold (b) CFP/FP number vs. # links contained for Italian, $t=VI$

Fig. 10. Comparison of space efficiency of structures CFP[G] and FP[G]

stores *link failure state probabilities* instead. Our proposed pre-computation data sets allow us to quickly compute the cumulative failure probability of any arbitrary set of links and can be utilized to more accurately compute the availability of network paths. We have proven that the memory usage of our memory-efficient data structure is upper-bounded by $O((n+x)\rho^3\gamma^4)$ if the failure of a link only depends on the distance to the epicenter of the disaster, where n is the number of nodes, x is the number of link crossings (in practice $x \ll n$), ρ is the maximal number of links subject to a disaster failure, and γ is the maximal number of line segments of a single link.

Our approach facilitates a comprehensive service availability analysis and can be used to answer related questions as well, such as where to place VMs in order to guarantee a certain SLA.

REFERENCES

- [1] J. Tapolcai, B. Vass, Z. Hesberger, J. Biró, D. Hay *et al.*, "A tractable stochastic model of correlated link failures caused by disasters," in *Proc. IEEE INFOCOM*, Honolulu, USA, Apr. 2018.
- [2] A. Valentini, B. Vass, J. Oostenbrink, L. Csák, F. Kuipers *et al.*, "Network resiliency against earthquakes," in *2019 11th International Workshop on Resilient Networks Design and Modeling (RNDM)*, Oct 2019, pp. 1–7.
- [3] G. Aceto, A. Botta, P. Marchetta, V. Persico, and A. Pescapé, "A comprehensive survey on internet outages," *Journal of Network and Computer Applications*, vol. 113, pp. 36–63, 2018.
- [4] P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, "The resilience of WDM networks to probabilistic geographical failures," *IEEE/ACM Trans. Netw.*, vol. 21, no. 5, pp. 1525–1538, 2013.
- [5] P. N. Tran and H. Saito, "Geographical route design of physical networks using earthquake risk information," *IEEE Communications Magazine*, vol. 54, no. 7, pp. 131–137, 2016.
- [6] H. Honda and H. Saito, "Nation-wide disaster avoidance control against heavy rain," *IEEE/ACM Transactions on Networking*, vol. 27, no. 3, pp. 1084–1097, 2019.
- [7] Y. Nemoto and K. Hamaguchi, "Resilient ICT research based on lessons learned from the Great East Japan Earthquake," *IEEE Commun. Mag.*, vol. 52, no. 3, pp. 38–43, 2014.
- [8] J. Heidemann, L. Quan, and Y. Pradkin, *A preliminary analysis of network outages during hurricane Sandy*. University of Southern California, Information Sciences Institute, 2012.
- [9] J. Oostenbrink and F. Kuipers, "Computing the impact of disasters on networks," *ACM SIGMETRICS Performance Evaluation Review*, vol. 45, no. 2, pp. 107–110, 2017.
- [10] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of the fiber infrastructure to disasters," *IEEE/ACM Trans. Netw.*, vol. 19, no. 6, pp. 1610–1623, 2011.

(a) CFPs of Italian for $t=VI$ (b) Italian topology

Fig. 11. CFP comparison of single, double and triple link failures for Italian

- [11] M. T. Gardner and C. Beard, "Evaluating geographic vulnerabilities in networks," in *IEEE Int. Communications Quality and Reliability Workshop (CQR)*, 2011, pp. 1–6.
- [12] 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, no. 3, pp. 908–921, 2015.
- [13] M. F. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of disaster-resilient optical datacenter networks," *J. Lightw. Technol.*, vol. 30, no. 16, pp. 2563–2573, 2012.
- [14] F. Dikbiyik, M. Tornatore, and B. Mukherjee, "Minimizing the risk from disaster failures in optical backbone networks," *J. Lightw. Technol.*, vol. 32, no. 18, pp. 3175–3183, 2014.
- [15] 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, no. 12, pp. 88–95, 2014.
- [16] 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.
- [17] B. Mukherjee, M. Habib, and F. Dikbiyik, "Network adaptability from disaster disruptions and cascading failures," *IEEE Commun. Mag.*, vol. 52, no. 5, pp. 230–238, 2014.
- [18] 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, no. 10, pp. 52–58, 2014.
- [19] O. Gold and R. Cohen, "Coping with physical attacks on random network structures," in *IEEE ICC*, 2014, pp. 1166–1172.
- [20] 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)*, 2012, pp. 121–126.
- [21] H. Saito, "Analysis of geometric disaster evaluation model for physical networks," *IEEE/ACM Trans. Netw.*, vol. 23, no. 6, pp. 1777–1789, 2015.
- [22] —, "Spatial design of physical network robust against earthquakes," *J. Lightw. Technol.*, vol. 33, no. 2, pp. 443–458, 2015.
- [23] F. Iqbal and F. Kuipers, "Spatiotemporal risk-averse routing," in *IEEE INFOCOM Workshop on Cross-Layer Cyber Physical Systems Security (CPSS)*, 2016.
- [24] J. Tapolcai, L. Rónyai, B. Vass, and L. Gyimóthi, "List of shared risk link groups representing regional failures with limited size," in *IEEE INFOCOM*, Atlanta, USA, May 2017.
- [25] T. Gomes, J. Tapolcai, C. Esposito, D. Hutchison, F. Kuipers *et al.*, "A survey of strategies for communication networks to protect against large-scale natural disasters," in *Int. Workshop on Reliable Networks Design and Modeling (RNDM)*, 2016.
- [26] 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, no. 6, pp. 2421–2434, 2020.
- [27] 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.
- [28] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE network*, vol. 14, no. 6, pp. 16–23, 2000.
- [29] O. Crochat, J.-Y. Le Boudec, and O. Gerstel, "Protection interoperability

- for WDM optical networks,” *IEEE/ACM Trans. Netw.*, vol. 8, no. 3, pp. 384–395, 2000.
- [30] S. Verbrugge, D. Colle, P. Demeester, R. Huelsermann, and M. Jaeger, “General availability model for multilayer transport networks,” in *Design of Reliable Communication Networks (DRCN)*, Lacco Ameno, Italy, Oct. 16–19, 2005.
- [31] C. S. Ou and B. Mukherjee, *Survivable Optical WDM Networks*. Springer Science & Business Media, 2005.
- [32] A. Somani, *Survivability and traffic grooming in WDM optical networks*. Cambridge University Press, 2006.
- [33] S. Yang, S. Trajanovski, and F. Kuipers, “Availability-based path selection and network vulnerability assessment,” *Wiley Networks*, vol. 66, no. 4, pp. 306–319, 2015.
- [34] H.-W. Lee, E. Modiano, and K. Lee, “Diverse routing in networks with probabilistic failures,” *IEEE/ACM Trans. Netw.*, vol. 18, no. 6, pp. 1895–1907, 2010.
- [35] J. Liu, J. Zhang, Y. Zhao, C. Ma, H. Yang *et al.*, “Differentiated quality-of-protection provisioning with probabilistic SRLG in flexi-grid optical networks,” in *OSA Asia Communications and Photonics Conference*, 2013, pp. AF2G–8.
- [36] F. Iqbal, S. Trajanovski, and F. Kuipers, “Detection of spatially-close fiber segments in optical networks,” in *Design of Reliable Communication Networks (DRCN)*, 2016, pp. 95–102.
- [37] M. Rahnmay-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)*, 2011, pp. 1–6.
- [38] US National Seismic Hazard Maps. Available at <https://earthquake.usgs.gov/hazards/hazmaps/conterminous/>.
- [39] D. Eppstein, M. T. Goodrich, and D. Strash, “Linear-time algorithms for geometric graphs with sublinearly many edge crossings,” *SIAM Journal on Computing*, vol. 39, no. 8, pp. 3814–3829, 2010.
- [40] C. Bohler, C.-H. Liu, E. Papadopoulou, and M. Zavershynskiy, “A randomized divide and conquer algorithm for higher-order abstract voronoi diagrams,” in *Algorithms and Computation*, H.-K. Ahn and C.-S. Shin, Eds. Cham: Springer International Publishing, 2014, pp. 27–37.
- [41] E. Papadopoulou and M. Zavershynskiy, “The higher-order Voronoi diagram of line segments,” *Algorithmica*, vol. 74, no. 1, pp. 415–439, 2016.
- [42] Network library. Available at <https://github.com/jtapolcai/regional-srlg/tree/master/psrlg>. Accessed:2020.11.16.
- [43] A. Rovida, M. Locati, R. Camassi, B. Lolli, and P. Gasperini, “Cpti15, the 2015 version of the parametric catalogue of italian earthquakes,” *Istituto Nazionale di Geofisica e Vulcanologia*, 2016.
- [44] H. Kanamori, “The energy release in great earthquakes,” *Journal of Geophysical Research (1896-1977)*, vol. 82, no. 20, pp. 2981–2987, 1977.
- [45] A. Sieberg, “Erdebeben,” *Handbuch der Geophysic*, vol. 4, pp. 552–554, 1931.
- [46] C. Mueller, “Earthquake catalogs for the usgs national seismic hazard maps,” *Seismological Research Letters*, vol. 90, 10 2018.
- [47] Y. Y. Kagan, “Statistical distributions of earthquake numbers: consequence of branching process,” *Geophysical Journal International*, vol. 180, no. 3, pp. 1313–1328, 2010.
- [48] 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.
- [49] A. Frankel, “Simulating strong motions of large earthquakes using recordings of small earthquakes: the loma prieta mainshock as a test case,” *Bulletin of the Seismological Society of America*, vol. 85, no. 4, pp. 1144–1160, 1995.
- [50] A. Valentini, F. Visini, and B. Pace, “Integrating faults and past earthquakes into a probabilistic seismic hazard model for peninsular italy,” *Natural Hazards and Earth System Sciences*, vol. 17, pp. 2017–2039, 2017.
- [51] M. D. Petersen, A. M. Shumway, P. M. Powers, C. S. Mueller, M. P. Moschetti *et al.*, “The 2018 update of the us national seismic hazard model: Overview of model and implications,” *Earthquake Spectra*, vol. 36, no. 1, pp. 5–41, 2020.
- [52] Y. Y. Kagan, “Seismic moment distribution revisited: I. statistical results,” *Geophysical Journal International*, vol. 148, no. 3, pp. 520–541, 2002.
- [53] D. H. Weichert, “Estimation of the earthquake recurrence parameters for unequal observation periods for different magnitudes,” *Bulletin of the Seismological Society of America*, vol. 70, no. 4, pp. 1337–1346, 1980.
- [54] S. Hiemer, J. Woessner, R. Basili, L. Danciu, D. Giardini, and S. Wiemer, “A smoothed stochastic earthquake rate model considering seismicity and fault moment release for europe,” *Geophysical Journal International*, vol. 198, pp. 1159–1172, 06 2014.
- [55] J. Giardini, D. and. Woessner, L. Danciu, H. Crowley, F. Cotton, G. Grunthal *et al.* (2013) Seismic Hazard Harmonization in Europe (SHARE): Online Data. DOI: 10.12686/SED-00000001-SHARE.
- [56] P. C., D. Albarello, P. Gasperini, V. D’Amico, and B. Lolli, “The attenuation of seismic intensity in italy, part ii: modeling and validation,” *Bulletin of the Seismological Society of America*, vol. 98, no. 2, pp. 692–708, 2008.
- [57] W. Bakun, “MMI attenuation and historical earthquakes in the basin and range province of western North America,” *Bulletin of the Seismological Society of America*, vol. 96, no. 6, pp. 2206–2220, 2006.

APPENDIX

A. Geometric Transformation of the Network

In our geometric reasoning, we transform the links of the graph into line segments. We also need to ensure that no two segments share a common endpoint. In the network, the adjacent links terminate in a single node; thus, we need to perform a minor transformation as follows.

Let $S \subseteq E$ be a set of segments and $\epsilon > 0$ a small number. Suppose that we shorten some segments e of S , in a way that we delete ϵ long subsegment from both ends, in such a way that the deleted intervals do not overlap. Let S' denote the set of segments S after shortening.

Lemma 3: We have $f(S, p) \geq f(S', p)$ and $f(S, p) - f(S', p) \leq \epsilon K$ hold for every point p .

Proof: For the first inequality note that

$$\begin{aligned} f(S, p) &= \int_{s=0}^1 \prod_{e \in S} I_{R(p, s)}(e) ds \\ &\geq \int_{s=0}^1 \prod_{e' \in S'} I_{R(p, s)}(e') ds = f(S', p) \end{aligned} \quad (24)$$

because $I_{R(p, s)}(e) \geq I_{R(p, s)}(e')$ holds for every s , whenever $e \in S$.

We turn now to the second inequality. Let s be the smallest value such that $\prod_{e \in S} I_{R(p, s)}(e) = 1$ (if there is any), and set $s' = s + \epsilon K$. Let d and d' be the radii of $r(p, s)$ and $r(p, s')$, resp. By the Lipschitz property we have $\epsilon K = s' - s \leq K(d' - d)$ giving that $d' > d + \epsilon$. We know by the definition of s that $r(p, s)$ intersects every segment $e \in S$ in some point Q_e . But then $r(p, s')$ intersects e' . This holds, because the larger disk $r(p, s')$ clearly contains the disk of radius ϵ centered at Q_e , and the latter disk must intersect e' because we deleted disjoint subintervals of length at most ϵ from e to obtain e' . We have therefore $\prod_{e' \in S'} I_{R(p, s')}(e') = 1$, hence

$$\begin{aligned} f(p, S) - f(p, S') &= \int_{y=0}^1 \left(\prod_{e \in S} I_{R(p, y)}(e) - \prod_{e' \in S'} I_{R(p, y)}(e') \right) dy \\ &\leq \int_{y=s}^{s'} 1 dy = \epsilon K. \end{aligned} \quad (25)$$

We transform our set of segments into one, where no segment e has an endpoint A on any other segment. If we have such

a segment, then we carry out the transformation by deleting an ϵ long subsegment of e starting at A . Lemma 3 gives that if we set ϵ sufficiently small, then all the values $f(p, S)$ and $f(p, S')$ will be very close to each other, hence $\text{CFP}(S)$ and $\text{CFP}(S')$ will be very close to each other. Moreover, for any two segments $e_1, e_2 \in E$, we have that either $e_1 \cap e_2 = \emptyset$, or $e_1 \cap e_2$ is an interior point of both segments.

As a simple example illustrating the Lipschitz condition 2) from IV-A, suppose that $r(p, s)$ is a disk centered at p having radius sR_p , where R_p is the radius of $r(p, 1)$. Then for radii $d = sR_p$ and $d' = s'R_p$ we have $|s' - s| = \frac{1}{R_p}|d' - d|$. The Lipschitz condition then holds if there exists a $k > 0$ such that $R_p \geq k$ for every p .

Balázs Vass received his MSc degree in applied mathematics at Eötvös Loránd University (ELTE), Budapest in 2016. He started his PhD studies in informatics in the same year on the Budapest University of Technology and Economics (BME). His research interests include networking, survivability, combinatorial optimisation, and graph theory. He received the Best Paper Award of NaNa'16, and the Best-in-Session Presentation award on INFOCOM'18. He was an invited speaker of COST RECODIS Training School on Design of Disaster-resilient Communication Networks in 2019.

János Tapolcai received the MSc degree in technical informatics and the PhD degree in computer science from the Budapest University of Technology and Economics (BME), Budapest, in 2000 and 2005, respectively, and the D.Sc. degree in engineering science from the Hungarian Academy of Sciences (MTA) in 2013. He is currently a Full Professor with the High-Speed Networks Laboratory, Department of Telecommunications and Media Informatics, BME. He has authored over 150 scientific publications. He was a recipient of several Best Paper

Awards, including ICC'06, DRCN'11, HPSR'15, and NaNa'16. He is a winner of the MTA Lendület Program and the Google Faculty Award in 2012, Microsoft Azure Research Award in 2018. He is a TPC member of leading conferences, e.g., IEEE INFOCOM 2012-, and the general chair of ACM SIGCOMM 2018.

Zsolt Heszberger received his M.Sc. and Ph.D. degree in electrical engineering at the Budapest University of Technology and Economics (BME), Budapest, Hungary in 1997 and 2007, respectively. Currently he is an Associate Professor at the Dept. of Telecommunications and Media Informatics, BME. His main research interests are future internet technologies and complex networking. Currently, he is working on clean-slate design internet routing and network management algorithms.

and geometry of networks.

József Bíró received his M.Sc. and Ph.D. degree in electrical engineering at the Budapest University of Technology and Economics (BME), Budapest, Hungary in 1993 and 1998, respectively, and the D.Sc. degree in engineering science from the Hungarian Academy of Sciences (MTA) in 2009. Currently he is Full Professor at the High-Speed Networks Laboratory, Dept. of Telecommunications and Media Informatics and head of the Doctoral School of Electrical Engineering, BME. His main research interests are network science, stochastic modelling

David Hay is an Associate Professor and the chair of the Federmann Cybersecurity Research Center in the Rachel and Selim Benin School of Computer Science and Engineering, Hebrew University, Jerusalem, Israel. He received the B.A. (summa cum laude) and Ph.D. degrees in computer science from the Technion-Israel Institute of Technology, Haifa, Israel, in 2001 and 2007, respectively. In addition, he was with IBM Haifa Research Labs, Haifa, Israel; Cisco Systems, San Jose, CA, USA; the Electronic Department, Politecnico di Torino, Turin, Italy; and

the Electrical Engineering Department with Columbia University, New York, NY, USA. He has served as a technical program committee member of numerous networking conferences, and since 2018 serves as an editor of ACM/IEEE Transactions on Networking. His research interests are in computer networks—in particular, network algorithmics, packet classification, deep packet inspection, network survivability and resilience, software-defined networking, network-function virtualization, and various aspects of network security.

Fernando Kuipers is leading the Lab on Internet Science at the Delft University of Technology (TU Delft). His research revolves around understanding and improving the performance and reliability of Internet and communications infrastructures. In 2004, Fernando obtained his Ph.D. degree cum laude, the highest possible distinction at TU Delft, and was subsequently promoted to assistant professor. Since 2010, he is an associate professor at TU Delft. He also held visiting scholar positions at Technion (2009) and Columbia University (2016). Fernando is

the scientific director and co-founder of Do IoT, a TU Delft research initiative and fieldlab on 5G and IoT. He is also board member of the TU Delft Safety & Security institute. Previously, he was in the program advisory board of SURFnet (2017-2018), board member of KIVI Telecom (2017-2018), co-founder and board member of PowerWeb (2012-2015), and chairman of the board of examiners (2008-2016). Internationally, he currently serves as vice-chair of IFIP Working Group 6.2 on Network and Internetwork Architectures, board member of the IEEE Benelux chapter on communications and vehicular technology, and general co-chair of ACM SIGCOMM 2021.

Jorik Oostenbrink is a PhD candidate at the TU Delft, where he has also obtained his MSc degree in Computer Science in 2015. He is a member of the Lab on Internet Science. His research interests include network resilience, programmable networks, and machine learning.

Alessandro Valentini received the MSc (April 2015) in Geological Sciences and Technologies and the PhD (March 2019) in Earth Systems and Built Environments at the University of the Study G. D'Annunzio, Chieti, Italy. His research interests include probabilistic seismic hazard analysis, fault displacement hazard analysis, earthquake geology and speleoseismology. He co-authored more than 10 papers on these topics and he presented more than 20 abstracts at national and international conferences, meetings and workshops.

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