Earthquake-Induced Link Failure Scenarios for Traffic Engineering Resiliency

Bennet Janzen, Leon Richardt, Nils Aschenbruck
Osnabrück University, Institute of Computer Science
Osnabrück, Germany
{bennet.janzen, richardt, aschenbruck}@uos.de

***Abstract*—The resiliency of Traffic Engineering (TE) algorithms is a critical challenge. Examining the performance against correlated, multi-link failures requires realistic failure scenarios with associated probabilities. However, no publicly available dataset provides such scenarios for the network topologies commonly used in TE research. We address this gap by modeling the impact of seismic activity on network infrastructure and generating the Earthquake-induced Link Failure (ELF) dataset of Probabilistic Shared Risk Link Groups (PSRLGs) for 30 backbone topologies from the REPETITA dataset. Using global earthquake occurrence rates and ground-motion propagation models, we compute for each network all earthquake-induced link failure scenarios with their annual occurrence probabilities. To demonstrate the ELF dataset's value, we evaluate several state-of-the-art TE algorithms. Our results show conventional TE approaches exhibit high annual overload probabilities. A strategic TE approach cannot find a single pre-computed configuration that withstands all failure scenarios due to their large number and diversity. Some failure scenarios even cause overload under optimal routing. These findings provide evidence for combining strategic TE with tactical measure in order to withstand regional failures. We make the ELF dataset publicly available to encourage further research into regional failure resiliency.**

***Index Terms*—Traffic Engineering, Segment Routing, Resiliency, PSRLG, Regional Failures**

I. Introduction

Ensuring continuous network service under failures is a critical objective for operators. Communication networks are part of the critical infrastructure of modern society, and disruptions in the backbone networks of Internet Service Providers (ISPs) can trigger cascading impacts that extend far beyond the network itself. Consequently, network resiliency, the ability to maintain acceptable performance in the presence of failures, is among the primary design goals for ISP backbone networks.

Recent research addresses this challenge through Traffic Engineering (TE) algorithms that are designed to be resilient against failures. In the context of Segment Routing (SR), strategic approaches like PCA2SR [1] proactively compute configurations that remain performant under predefined failure scenarios. However, these approaches are typically evaluated against single-link failures and operator-defined Shared Risk Link Groups (SRLGs), such as fiber bundles sharing a conduit. A potentially more severe class of failures stems from regional failures where all network infrastructure within a geographic area is affected simultaneously. Regional failures can be caused by natural disasters like earthquakes, hurricanes, or floods. Such events can cause multi-link outages that could heavily influence the performance of the network.

ISBN 978-3-903176-82-9

Evaluating TE algorithms against these regional failure scenarios does not only require the knowledge of the affected links but also the associated occurrence probabilities of the scenarios. While seismic risk models for networks have been proposed in the literature [2], [3], no publicly available dataset exists that provides Probabilistic SRLGs (PSRLGs) for the network topologies commonly used in TE research. The widely used Internet Topology Zoo [4] and REPETITA datasets [5] provide topologies and traffic matrices for TE evaluation. They do not contain any information about SRLGs or correlated failure scenarios and their probabilities. The TE research community currently lacks the data needed to evaluate the resiliency of TE algorithms against regional failures.

In this paper, we address this gap. Our primary contribution is the Earthquake-induced Link Failure (ELF) dataset [6] containing earthquake-induced PSRLGs for 30 backbone topologies from the REPETITA dataset. By modeling the impact of seismic activity using forecasted earthquake rates and ground-motion propagation models, we compute, for each network, the full set of link failure scenarios that earthquakes can cause, together with their annual occurrence probabilities. To the best of our knowledge, ELF is the first publicly available PSRLG dataset of this scope for network topologies that are commonly used in TE research.

To demonstrate the utility of this dataset, we use it to evaluate the resiliency of several state-of-the-art TE algorithms. This evaluation yields three key insights. First, conventional TE approaches like optimizing IGP weights or 2-Segment Routing (2SR) are highly vulnerable to the failure scenarios in the dataset, with many scenarios causing network overload. Considering the earthquake occurrence frequencies, we find that the aggregate annual probability of network overutilization is remarkably high for some networks when using such conventional approaches. Second, even the strategic resilient TE approach PCA2SR [1] provided with the complete set of seismic failure scenarios as input leaves a substantial gap to the performance achievable when SR policies are recomputed for each individual failure scenario. This demonstrates that the number of failure scenarios is too large and diverse to be handled by a single pre-computed SR configuration. Strategic TE should be combined with a tactical TE that reacts to

specific failure scenarios. Third, for some failure scenarios even an optimal reconfiguration yields overutilization. This means that no tactical TE can solve these failure scenarios.

The remainder of this paper is structured as follows. The relevant background is introduced in Section II. In Section III, we describe the modeling approach and the generation of the ELF dataset. Section IV demonstrates how the dataset can be leveraged to create deeper insights into network resiliency under regional failures. Finally, the paper is concluded in Section V.

II. Background and Related Work

In this section, we provide background and related work for failure models and SRLGs, seismic risk modeling for networks, network datasets, and TE.

A. Failure Models and Shared Risk Link Groups

Network resiliency can be evaluated under failure models that define which elements can fail and how failures are correlated. The simplest model considers only independent single-link or single-node failures. SRLGs generalize this by defining sets of links that share a common risk. SRLGs are typically defined by operators based on physical infrastructure knowledge, such as fiber bundles sharing a conduit. A PSRLG additionally associates a failure probability with each link set.

Regional failures represent a special case of PSRLGs, as they are often caused by large-scale catastrophic events. They are defined as geographic events that simultaneously destroy all network infrastructure within an affected area. Typical reasons for regional failures are earthquakes, floods, hurricanes or tsunamis. They can affect a substantial number of links simultaneously, yet occur with comparatively low probability. Despite the practical relevance of regional failures, no publicly available PSRLG dataset exists for the network topologies commonly used in TE research.

Tapolcai et al. [7] propose a general stochastic model of regional failures in networks. This model calculates the joint failure probability of a set of links, i.e. the PSRLGs of the network. Their model can consider regional failures of arbitrary size and shape.

B. Seismic Risk Modeling for Networks

Valentini et al. [2] propose a model for computing link failure probabilities caused by seismic activity. They use historic earthquake catalogs to compute earthquake occurrence rates and model the impact of earthquakes using ground-motion propagation functions.

Vass et al. [3] extend this by computing PSRLGs, i.e., joint failure probabilities for sets of links. They compute their model on a few networks in Italy, the US and Europe. The PSRLGs can be used to compute the availability of routing paths. However, their probabilities are normalized, representing the conditional probability that a specific link set fails given that a regional link failure event occurs, rather than annual failure probabilities.

C. Network Datasets

The REPETITA dataset [5] is widely used for evaluating TE algorithms. It provides network topologies from the Internet Topology Zoo [4], the Rocketfuel project [8], and the DEFO project [9], together with synthetically generated traffic matrices using a random gravity model [10]. These datasets enable standardized comparison of TE approaches across a diverse set of network topologies. However, REPETITA contains no information about failure scenarios beyond the topology itself. The ELF dataset fills this gap by adding earthquake-induced PSRLGs for REPETITA networks.

D. Traffic Engineering

The Interior Gateway Protocols (IGPs) Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS) rely on simple shortest path routing. If the IGP operates with simple unary link weights, this results often in high Maximum Link Utilizations (MLUs). By adjusting link weights, shortest paths and consequently the MLU can be influenced. Fortz and Thorup [11] demonstrate that optimizing IGP link weights using a local-search heuristic can significantly reduce the link utilizations. IGP link weight optimization is a widely adopted and simple form of TE. For theoretically optimal TE, the multicommodity flow (MCF) linear program can be formulated. Given a network topology and a set of traffic demands, it computes a routing configuration that minimizes the MLU. Due to real-world hardware limitations, the solution of this problem only provides a lower bound on the achievable MLU. Segment Routing (SR) is a modern source-routing paradigm in which forwarding paths are encoded as sequences of segments [12]–[14]. These sequences of segments can be installed as SR policies at a headend node to steer traffic, e.g., to reduce link utilizations by distributing traffic more evenly across available paths. It has been shown that near-optimal TE results can often be achieved using only two node segments [15], [16], an approach commonly referred to as 2SR.

Resilient TE can be divided into strategic and tactical TE. Strategic TE approaches aim to compute a configuration that remains robust over an extended period of time despite failures and traffic variations [15]–[18]. Schüller et al. [1] propose Post-Convergence Aware 2SR (PCA2SR), which finds a single SR configuration optimized across a set of predefined failure scenarios. PCA2SR computes the post-convergence routing for each scenario and solves a linear program to find one set of SR policies that minimizes the sum of the overutilization across all scenarios. In contrast, tactical TE approaches [19]–[21] react to failures after they occur to restore network utilization to an acceptable level. Such approaches become necessary when unexpected or large-scale failures arise that cannot be adequately handled by the precomputed strategic TE configuration. The performance of strategic TE approaches depend on the set of provided failure scenarios. Only if the scenario set accurately reflects realistic failure events, a strategic TE can compute a solution that can be considered resilient.

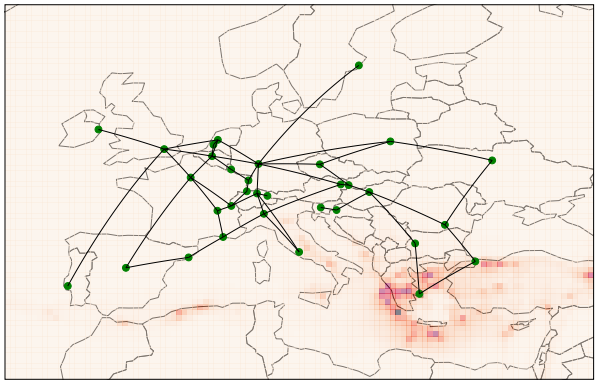

Fig. 1. Network topology of Bics (E4) overlaid on a map with earthquake occurrence rates.

III. Modeling and Dataset Generation

In this section, we describe the construction of the ELF dataset, including the selection of network data, the earthquake model and the computation of failure probabilities.

A. Network Data

We build the ELF dataset using the network topologies from the REPETITA dataset [5]. The REPETITA authors preprocess each topology from the Internet Topology Zoo to form a single connected component and assign uniform link capacities where no data is available. For each network, five traffic matrices are generated using a random gravity model and scaled so that an optimal TE solution achieves an MLU of 90 %.

To compute regional failure induced PSRLGs, geographic node coordinates are needed. Of the REPETITA topologies, 85 include this information. We exclude networks with fewer than 20 nodes, as they are not representative of ISP backbone networks. After that we exclude two networks that have a spanning tree topology, which are not interesting for TE evaluations, yielding a final set of 30 networks (Table I). Since no geographic link data is available, we model each link as the geodesic path between its endpoints, consistent with [2], [3].

The 30 networks span three geographic regions: eight European (E1–E8), 15 North American (A1–A15) and seven worldwide (W1–W7). For the performance evaluation, we use all 30 networks. Due to the lack of space, we select 6 representative network for a more in-depth analysis: Bics (E4), EliBackbone (A1), AttMpls(A6), Tw (A15), Packetexchange (W2) and Ntt (W7). The Bics (E4) topology is shown as an example in Figure 1.

B. Earthquake Modeling

To model the impact of earthquakes on network infrastructure, we combine earthquake source data with a ground-motion propagation model.

TABLE I
Networks in the ELF dataset with identifier, number of nodes $|V|$, edges $|E|$, and PSRLGs.

Network	ID	$\|V\|$	$\|E\|$	|PSRLG|
Fccn	E1	23	25	641
Renater2001	E2	24	27	1031
Geant2001	E3	27	38	1641
Bics	E4	33	48	2376
Arnes	E5	34	46	1681
Geant2009	E6	34	52	2029
Geant2010	E7	37	56	2561
Uninett2010	E8	74	101	7494
EliBackbone	A1	20	30	490
Atmnet	A2	21	22	198
Psinet	A3	24	25	372
Arpanet19723	A4	25	28	302
Agis	A5	25	30	505
AttMpls	A6	25	56	1693
Bbnplanet	A7	27	28	318
Integra	A8	27	36	729
Darkstrand	A9	28	31	362
Arpanet19728	A10	29	32	454
Digex	A11	31	35	855
CrlNetworkServices	A12	33	38	712
Xspedius	A13	34	49	2718
Bellcanada	A14	48	64	1940
Tw	A15	71	115	8159
Quest	W1	20	31	112
Packetexchange	W2	21	27	220
Bandcon	W3	21	28	232
Xeex	W4	24	34	620
HurricaneElectric	W5	24	37	500
DeutscheTelekom	W6	30	55	1931
Ntt	W7	32	63	1272

1) Earthquake Sources: We adopt a scenario-based approach grounded in Probabilistic Seismic Hazard Analysis (PSHA) [22] while preserving spatial correlations between neighboring locations. In contrast to Valentini et al. [2] we do not use a seismicity model to calculate earthquake rates per location but use an earthquake source dataset. We use the GEAR1 dataset [23], which provides annual earthquake occurrence rates for the entire globe on a 0.1-degree grid per 0.1-magnitude bin.

2) Ground-Motion Propagation: Given an earthquake with epicenter and magnitude, we model circular propagation of seismic waves using the following intensity attenuation relation from Bakun et al. [24]:

$$I = 0.44 + 1.7 \cdot M_w - 0.0048 \cdot D - 2.73 \log_{10} D \quad (1)$$

where M_w is the moment magnitude and $D = \sqrt{R^2 + h^2}$ the hypocentral distance with epicentral distance R and hypocentral depth $h = 3.91\,\mathrm{km}$. This equation computes the level of ground shaking for every location for a certain earthquake source. The propagation model was designed for the USA. It is also used in [3]. While using a single propagation function globally is a simplification, we observed comparable results with models calibrated for specific regions like [25] for Italy.

3) Destruction Radius: Network infrastructure is assumed to fail at locations experiencing a Modified Mercalli Intensity (MMI) of VI or higher [26]. An MMI level of VI

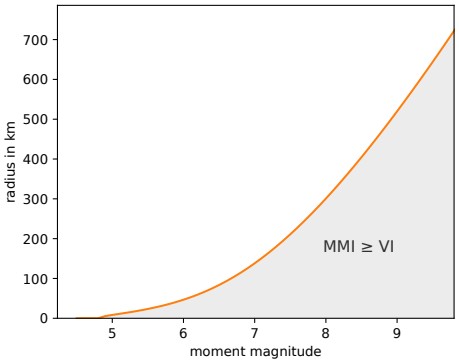

Fig. 2. Destruction radius per magnitude for MMI level VI.

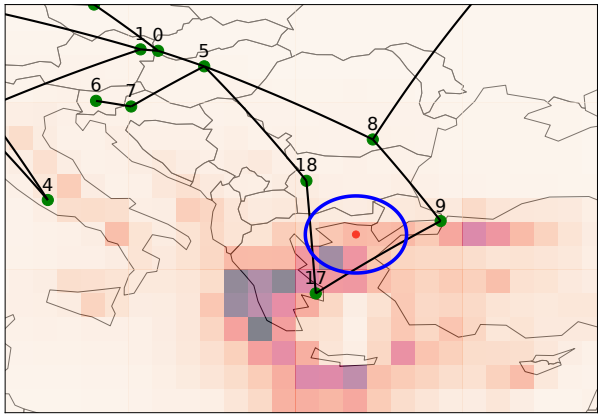

Fig. 3. A magnitude 7.2 earthquake near northern Greece affects links 9–17 and 17–18 in the Bics network.

is described as strong shaking and covers a peak ground acceleration of 0.1 to 0.2 g. This ground shaking level is high enough to consequence damage at buildings an also network infrastructure. Since the propagation model assumes circular wave propagation, we can precompute for each magnitude M_w the destruction radius $D(M_w)$. This depicts the maximum distance from the epicenter at which the MMI still exceeds level VI (Figure 2). Following prior approaches [2], [3], we add the grid-cell diameter to the destruction radius to compensate for placing earthquakes anywhere in the cells. An example destruction area on the Bics network is shown in Figure 3.

C. PSRLG Computation

So far, we know how to model earthquakes and their impact on the links of networks. Next we combine the probabilities of all possible earthquake scenarios to correlated annual failure probabilities of sets of links due to these earthquakes. The sets of links can be seen as SRLGs and together with the probability we build PSRLGs. We base our notations on [2].

Let $c_{x,y}$ denote a grid cell with $x \in \mathcal{G}_x$, $y \in \mathcal{G}_y$. We consider magnitudes $M_w \in \mathcal{M} = \{6.1, 6.2, \ldots, 9.0\}$, as they are included in the GEAR1 dataset. The event E_{x,y,M_w} represents an earthquake with magnitude in $(M_w - 0.1, M_w]$ and epicenter in cell $c_{x,y}$, occurring with annual rate r_{x,y,M_w}. Each event destroys all links within the disk of radius $D(M_w)$ centered at $c_{x,y}$. We denote the resulting set of failing links as F_{x,y,M_w}.

For a link set S, define the indicator

$$I_{x,y,M_w}(S) = \begin{cases} 1 & \text{if } S = F_{x,y,M_w}, \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

Assuming independent Poisson arrivals, the total annual rate at which exactly the link set S fails is

$$R_{\text{failure}}(S) = \sum_{x \in \mathcal{G}_x} \sum_{y \in \mathcal{G}_y} \sum_{M_w \in \mathcal{M}} r_{x,y,M_w} \cdot I_{x,y,M_w}(S). \tag{3}$$

The annual probability that at least one event causes exactly the failure of set S is

$$P_{\text{failure}}(S) = 1 - e^{-R_{\text{failure}}(S)}. \tag{4}$$

A PSRLG is now given by the pair $(S, P_{\text{failure}}(S))$. We compute PSRLGs by iterating over all grid cells covering the network's extent and all magnitudes in $\mathcal{M}$. In [2], [3] joint probabilities for the failure of a set S are computed and normalized to obtain the conditional probabilities that exactly the set S fails in the next earthquake event. In contrast, ELF provides additional annual failure probabilities that are useful for the evaluation of TE resiliency.

D. Dataset Characteristics

The ELF dataset comprises over 40 000 distinct link failure scenarios across the 30 networks. Table I lists the number per network, ranging from 112 (Quest) to over 8 000 (Tw). This variation reflects differences in network size, geographic extent and seismic exposure of the respective region.

Figure 4 shows the empirical cumulative distribution of annual failure probabilities for selected networks. The vast majority of individual PSRLGs have very low probabilities (well below 10^{-4}). The probabilities correspond to a very specific earthquake event occurring within a 0.1-degree grid cell and with a magnitude specified to a precision of 0.1. As we show in Section IV-C, the high number of failure scenarios lead to non-negligible aggregate probabilities when they are combined.

Figure 5 shows the distribution of PSRLG sizes, i.e., the number of simultaneously failing links. While most PSRLGs contain one or two links, some involve the simultaneous failure of five or more links, representing severe regional failures. Most PSRLGs contain only a small number of links but the distribution of the number of PSRLGs differs between the individual networks. This is due to the size of the networks, but also due to their density and their geographical exposure to earthquakes.

IV. Evaluating Traffic Engineering Resiliency

To demonstrate the value of the ELF dataset, we use it to evaluate the resiliency of several state-of-the-art TE algorithms under the earthquake-induced failure scenarios. First, we present the evaluation setup and the evaluated TE

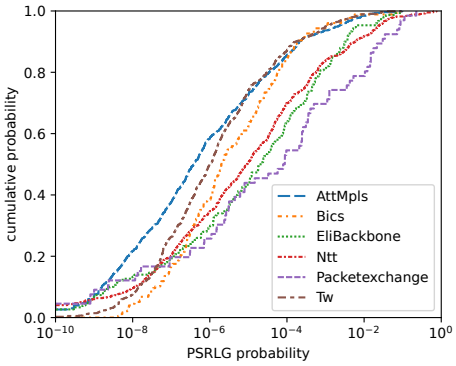

Fig. 4. ECDF of annual PSRLG failure probabilities.

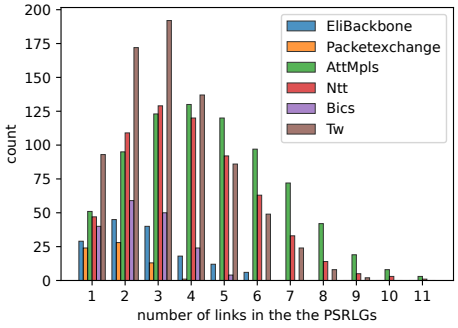

Fig. 5. Number of links per PSRLG for selected networks.

algorithms. After that we evaluate the MLU in the failure scenarios. Next, we incorporate the failure probabilities and derive the annual probability of exceeding certain MLU thresholds. Lastly, we deal with the expected failure MLU and overutilization probabilities.

A. Setup and Algorithms

In related work, PSRLGs are used to compute path failure probabilities [3]. We extend this and compute the probability that the network is overloaded due to failure scenarios derived by the PSRLGs. In a real network the MLU threshold 1 should not be exceeded, as this would indicate a network overutilization which causes packet loss and other unwanted behavior of the network. To calculate the probability that MLU thresholds are exceeded in a network, we consider all failure scenarios that arise from the PSRLGs. That means, for each PSRLG we consider the failure of all links it contains. To assess the network's utilization under the failure scenarios we evaluate the post-convergence utilizations as introduced in [1]. That means that the network configuration stays the same, but the IGP reconverges. SR policies stay configured but the shortest paths are recomputed on the updated topology without the failed links. Using this approach, we evaluate the performance of the following TE algorithms:

SPR is a simple ECMP shortest path routing with unary link weights, as implemented by link-state IGPs such as OSPF and IS-IS. Under failures, IGP reconvergence leads to new shortest paths using the unary link weights.

HeurSPR [11] is a TE approach introduced by Fortz and Thorup that influences the traffic paths by tuning the IGP link weights. It tries to lower the link utilizations using local search heuristics. Under failures, IGP reconvergence leads to new shortest paths using these pre-configured link weights.

2SR [15] is a basic SR-based TE algorithm that computes SR policies with at most two node segments. We use an LP formulation that minimizes the MLU for the intact network with unary link metrics. Under failures, the pre-configured SR policies remain in place. On each segment, the traffic is routed along the new shortest paths after IGP reconvergence.

PCA2SR [1] is a strategic resilient TE algorithm that computes resilient SR policies. Like 2SR, it is based on an LP and produces policies with at most two node segments. Its optimization objective is to lower the overutilization under a pre-defined set of failure scenarios. In our evaluation, we provide PCA2SR with the full set of earthquake-induced failure scenarios from the ELF dataset. PCA2SR has two parameters. The first parameter is the normal-case MLU Θ. We set it to be 0.8 as suggested by the authors. The second parameter is the failure-case MLU Φ. We set it to 1, as we want to provide flexibility for the algorithm to prevent network overutilization. Under failures, the pre-configured SR policies remain in place. On each segment, the traffic is routed along the new shortest paths after IGP reconvergence.

To assess whether failure scenarios can be mitigated by optimal tactical TE, we also evaluate two algorithms optimized for each failure scenario:

FO-2SR reoptimizes SR policies for the defect network, while still using at most two node segments. This solution shows to what extent the MLU can be reduced by tactically deploying modified 2SR policies. It therefore provides a lower bound on the MLU achievable by tactical 2SR-based TE algorithms.

FO-MCF reoptimizes an optimal MCF solution for each failure scenario. This yields an optimal routing for each failure scenario and thus a lower bound on the MLU in these scenarios.

These 6 algorithms are evaluated for the 30 networks of our dataset. We use all PSRLGs as failure scenarios in which the network is not segmented. We do not consider node failures. We use the traffic matrices derived from the REPETITA dataset [5] and scale them to 50 %, yielding an MLU of 0.45 under optimal TE on the intact network, which is reasonable for an ISP backbone network. The computations are done on a computer with an AMD EPYC 7452 CPU with 256GB of RAM and 64-bit Ubuntu 24.04.4. LPs are solved using CPLEX [27].

Since the probabilities of the failure scenarios are known through the PSRLGs, we can make further statements about the network. In particular, we can compute the probability that certain MLU thresholds are exceeded. Computing the probability that the MLU exceeds 1 corresponds to the annual probability of network overutilization.

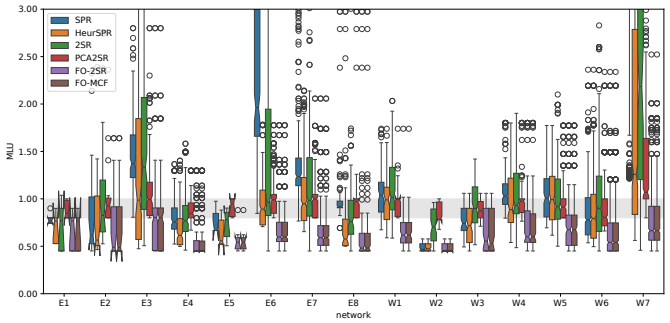

(a) European and Worldwide networks

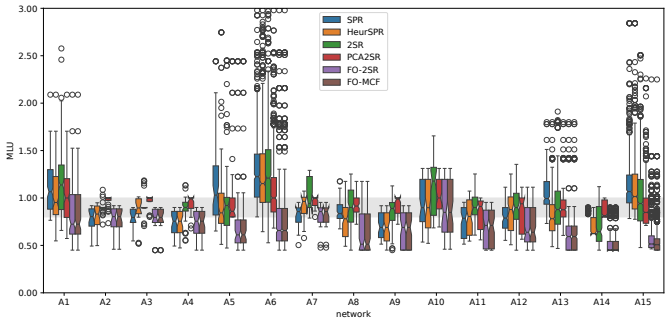

(b) North American networks

Fig. 6. MLU across all PSRLG failure scenarios by network and TE algorithm.

B. MLU in the failure scenarios

In this section, we are only interested in the link utilizations of the TE algorithms under the derived failure scenarios. Figure 6 shows the MLU for all failure scenarios, grouped by network and algorithm.

In Figure 6, the MLU is depicted across all failure scenarios by network and TE algorithm. The range between 0.8 and 1 is highlighted, as this interval is of relevance for network operators. The MLU should not exceed 0.8 in order to tolerate short-term traffic fluctuations or single-link failures. An MLU greater than 1 should never occur, as this corresponds to network overutilization, where traffic can no longer be fully routed. In a real network, this would cause packet loss, degraded quality of service, and potential violations of service level agreements. Since the MLU in the intact network under optimal routing is 0.45, this value constitutes the lower bound for all networks. It can be observed that in most networks the majority of failure scenarios leads to an increase in MLU, but for many cases the MLU remains below the threshold of 1. However, for some networks, a substantial fraction of scenarios results in significantly higher MLU values, with the MLU exceeding 1.

We now examine selected networks in greater detail, beginning with the network Tw (A15) as one of the largest and thus most representative of real-world backbone networks. For the non-FO algorithms, the MLU is between 0.8 and 1.2 in the majority of scenarios. Among these approaches, the highest MLU values are observed for SPR, followed by HeurSPR, then 2SR, and finally PCA2SR. Nevertheless, even for PCA2SR, the MLU exceeds 1 in a considerable number of scenarios. This shows that the failure scenarios from our dataset can have a great impact on the performance of the networks and lead to network overutilization, that even a strategic TE approach like PCA2SR could not prevent, even when all failure scenarios are known by the algorithm. This indicates that the number and diversity of failure scenarios is too large to find a pre-configured solution that solves all scenarios. For the FO-algorithms, the MLU values are significantly lower across the considered scenarios. In particular, FO-2SR yields slightly higher MLU values than FO-MCF. This can be explained by the greater flexibility of FO-MCF, in contrast to FO-2SR which is constraint to a maximum of two segments per demand.

Next, we consider the network Elibackbone (A1). For SPR, 2SR, and PCA2SR, nearly half of the scenarios reach the critical MLU threshold of 1. However, in this network 2SR exhibits the highest MLU values. This can be attributed to the fact that 2SR is optimized solely for the intact network and does not explicitly account for failure scenarios. Even the FO-algorithms result in MLU values exceeding 1 in more than 25 % of the scenarios. This indicates that even optimal tactical TE cannot fully mitigate the impact of these failure scenarios.

In the NTT (W7) network, we observe even higher MLU values, especially for the non-resilient TE algorithms. This results from heterogeneous link capacities. Most networks in the REPETITA dataset have uniform link capacities where this effect cannot be observed. The failures of high-capacity links lead to very high link utilization of alternative small links. Additionally the network has a high earthquake exposure due to its topology spanning mainly East Asia and the Pacific.

Overall, the results show that in many networks, the non-FO-algorithms lead to substantial increases in MLU compared to the failure-free case. In many scenarios, the MLU threshold of 1 is exceeded. While the FO-algorithms are able to mitigate many of these scenarios and limit the increase in MLU, there remain a few networks in which certain scenarios cannot be sufficiently compensated using tactical TE.

It should be emphasized that, in the preceding analysis, all scenarios have been evaluated solely on the basis of their resulting MLU values, irrespective of their associated occurrence probabilities. Consequently, no distinction has been made between highly probable yet minor scenarios and rare but severe ones. The following section addresses this limitation by incorporating the failure probabilities provided by the ELF dataset.

C. Annual Probability of Exceeding MLU Thresholds

A key advantage of the PSRLG dataset over pure scenario enumeration is that it provides annual failure probabilities. This enables computing the annual probability that the MLU exceeds a given threshold t due to any earthquake-induced failure scenario. Since the failure events follow independent

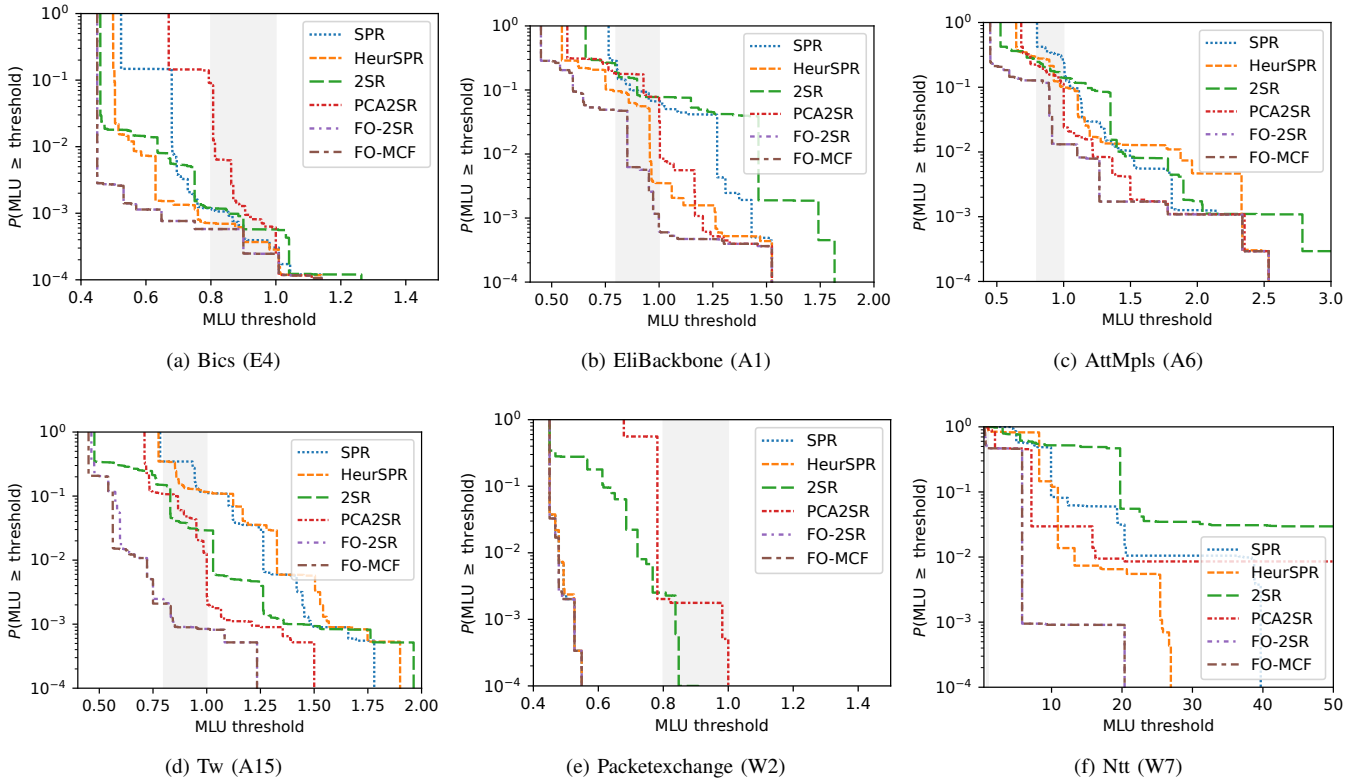

(a) Bics (E4) (b) EliBackbone (A1) (c) AttMpls (A6)

(d) Tw (A15) (e) Packetexchange (W2) (f) Ntt (W7)

Fig. 7. Annual probability of exceeding MLU thresholds under different TE algorithms.

Poisson processes, this probability is

$$P_{\text{PSRLG}}(\text{MLU} > t) = 1 - \prod_{\substack{S \in \text{PSRLG}, \\ \text{MLU}(S) > t}} \big(1 - P_{\text{failure}}(S)\big), \quad (5)$$

where $\text{MLU}(S)$ denotes the post-convergence MLU under a given TE algorithm when the links in S fail.

Figure 7 plots this probability as a function of the threshold for selected networks and algorithms. For example, if we consider the value at the MLU threshold of 1, the y-axis represents the annual probability that, when using the respective algorithm, the MLU exceeds 1. All plots use a logarithmic scale for the probabilities. The scaling of the x-axis differs between networks, as they are affected to varying degrees by earthquake-induced failure scenarios. Again the range between 0.8 and 1 is highlighted, as this interval is of relevance for network operators.

We begin again with the network Tw (A15, Figure 7d). The results observed here generalize to many networks in the dataset. In the intact network, 2SR and the FO-algorithms achieve an MLU of approximately 0.45, PCA2SR operates at about 0.75, and SPR as well as HeurSPR just below 0.8. In the operationally relevant range between MLU thresholds 0.8 and 1, the algorithms differ significantly. FO-MCF and FO-2SR achieve the lowest exceedance probabilities at roughly 0.1 %, followed by PCA2SR at approximately 0.2 %, 2SR at about 3 %, and SPR and HeurSPR at around 10 %. FO-MCF serves as a lower bound on the achievable MLU in the failure scenarios. FO-2SR deviates only marginally from FO-MCF, indicating that tactical 2SR-based TE is sufficient for near-optimal rerouting. PCA2SR is configured to have an MLU below 0.8 for the normal case and to minimize overutilization above the threshold 1. This explains the vertical drop at MLU $= 1$, as its policies are optimized to keep the MLU below this threshold. 2SR achieves a lower MLU for the normal case, but its exceedance probability is higher for MLU values greater than 1. HeurSPR and SPR yield comparable results. For higher MLUs, the FO-algorithms have the lowest probabilities, followed by PCA2SR, then SPR and HeurSPR. 2SR performs worst. This suggests that the policies worsen the routing performance under failures.

The performance of HeurSPR relative to the other algorithms varies across networks. In EliBackbone (A1, Figure 7b), it outperforms even PCA2SR in the relevant part of the network. It has a lower probability to exceed the MLU threshold 1 than PCA2SR, whereas in Tw it is partially worse than SPR. This can be explained by the additional routing flexibility that the IGP link weights provide over the segment policies. This suggests that the interplay between IGP link weight optimization and failure resiliency needs further evaluation.

The network Packetexchange (W2, Figure 7e) shows interesting behaviour. The SPR and HeurSPR solutions are as resilient as the FO-algorithms whereas 2SR and PCA2SR are more likely to exceed the relevant MLU threshold 0.8. The configured SR policies worsen the TE resiliency. This indicates

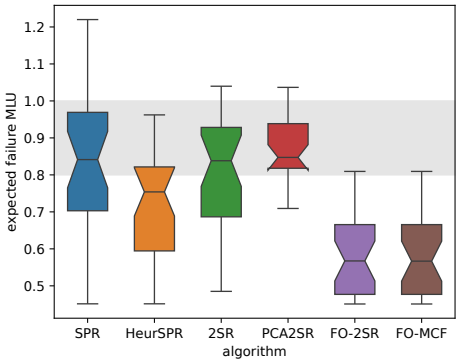

Fig. 8. Expected failure MLU by algorithm, over all networks.

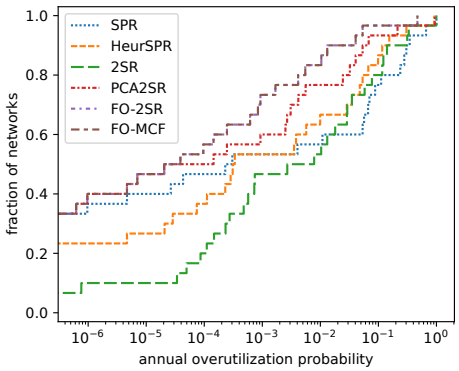

Fig. 9. ECDF of annual overutilization probabilities for all networks.

that there is potential to further enhance strategic TE, for example by minimizing the MLU in the failure scenarios to withstand additional effects like traffic fluctuations.

The Ntt network (W7, Figure 7f) exhibits exceptionally high MLU values across all algorithms. As explained earlier, this effect is caused by heterogeneous link capacities within the network. Since most topologies in the REPETITA dataset use unary link capacities, this phenomenon does not appear there. When high-capacity links fail, traffic is possibly redirected to lower-capacity links, resulting in severe overutilization on those smaller links. Furthermore, the network is particularly exposed to earthquakes because its topology primarily covers East Asia and the Pacific region. Finally, the gap between the FO-algorithms and PCA2SR is not present in every network. In AttMpls (A6, Figure 7c), strategic TE already mitigates most failure scenarios, leaving little room for tactical improvement. These effects show that the resiliency of TE algorithms depends on the network topology and the failure scenarios.

D. Expected failure MLU and overutilization probability

In Section IV-B, the MLU in failure scenarios is evaluated without probability weighting. For instance, the boxplots include highly unlikely but severe events alongside very likely yet minor events, treating them equally. To compare the TE approaches in a probability-weighted manner, we compute the expected MLU under the failure scenarios. More precisely, for each network we determine the conditional expected MLU given that any failure occurs. Figure 8 illustrates the expected failure MLU for all 30 networks, grouped by algorithm. We observe that the FO-algorithms yield the lowest expected MLU values. In most networks, these values remain below the threshold of 0.8, indicating that sufficient headroom is preserved to accommodate additional single-link failures or traffic fluctuations. SPR and 2SR exhibit median expected failure MLU values of approximately 0.85. However, there are networks in which the expected failure MLU exceeds 1, implying overutilization. HeurSPR achieves lower expected failure MLU values compared to SPR and 2SR. This suggests that the additional optimization of IGP link weights represents a possible direction for future work on resilient TE. PCA2SR also shows a median expected failure MLU of approximately 0.85 across the networks. However, in contrast to the non-resilient 2SR approach, the variance is lower, with most networks exhibiting values between 0.8 and 1. This behavior can be explained by the structure of the underlying linear program, which enforces a normal-case MLU of 0.8 and minimizes only the aggregate overutilization across all failure scenarios. Consequently, the expected failure MLU values tend to cluster close to the MLU threshold 1.

Next we evaluate the annual probability of overutilization, i.e., the probability that the MLU exceeds 1 in any given year. Figure 9 presents the annual overutilization probability for all networks in an ECDF plot. The curve shows, for each probability value on the x-axis, the cumulative fraction of networks that experience an overutilization with at most that annual probability. In this plot, FO-MCF is the upper bound. For 80 % of the networks, the annual probability of overutilization is below 0.5 %. In approximately 35 % of the networks, no overutilization is expected, as none of the considered failure scenarios leads to an MLU greater than 1. FO-2SR exhibits almost identical annual overutilization probabilities to those of FO-MCF. Thus, in a tactical TE setting, 2SR-based approaches are sufficient to achieve resiliency levels comparable to the optimal multicommodity flow solution. PCA2SR yields higher overutilization probabilities compared to the FO-algorithms. For 80 % of the networks, the annual probability of overutilization is below 3 %. Higher probabilities of overutilization are observed for the conventional TE algorithms. For SPR, HeurSPR, and 2SR, 80 % of the networks exhibit annual overutilization probabilities between 6 % and 10 %. Among these approaches, no algorithm consistently achieves lower overutilization probabilities than the others.

In summary, we observe that the probability of overutilization can be reduced compared to conventional TE by employing a strategic TE approach that explicitly incorporates the ELF dataset of failure scenarios to precompute a resilient configuration. However, in many networks, this probability can be further reduced through the application of tactical TE, which dynamically adapts the routing in response to occurring failure events.

V. Conclusion and Future Work

In this paper, we propose the ELF dataset that captures earthquake-induced correlated link failures for 30 backbone topologies from the REPETITA dataset.

The ELF dataset comprises earthquake-induced link failure scenarios with annual occurrence probabilities and is, to the best of our knowledge, the first publicly available PSRLG dataset for network topologies of the REPETITA dataset, enabling the evaluation of resilient TE algorithms against realistic, correlated failure scenarios. The ELF dataset is publicly available to enable reproducible research on TE resiliency [6].

To demonstrate the value of this dataset, we evaluated several state-of-the-art TE algorithms. Our results show that conventional approaches are highly vulnerable to regional failure scenarios and that even a strategic TE approach provided with all failure scenarios cannot prevent network overutilization for all scenarios due to the high number and diversity of the failure scenarios. This indicates that strategic TE should be combined with a tactical TE that reacts to specific failure scenarios. However, we also observe severe failure scenarios that cannot even be solved by an optimal tactical TE.

The dataset enables several directions for future work. Probability-aware strategic TE algorithms could leverage the failure probabilities of the ELF dataset to find SR configurations that minimize the expected overutilization probability rather than the worst-case MLU, focusing optimization effort on the most likely scenarios. Our work also indicates that optimizing IGP link weights can provide routing flexibility to improve strategic TE. Moreover, our PSRLG dataset can be extended with additional disaster types like floods, or hurricanes, in order to enable a more comprehensive resiliency assessment.

References

[1] T. Schüller, N. Aschenbruck, M. Chimani, and M. Horneffer, "Failure resiliency with only a few tunnels–enabling segment routing for traffic engineering," *IEEE/ACM Transactions on Networking*, pp. 262–274, 2020.

[2] A. Valentini, B. Vass, J. Oostenbrink, L. Csák, F. Kuipers, B. Pace, D. Hay, and J. Tapolcai, "Network resiliency against earthquakes," in *Proc. of the IEEE International Workshop on Resilient Networks Design and Modeling (RNDM)*, 2019, pp. 1–7.

[3] B. Vass, J. Tapolcai, Z. Heszberger, J. Bíró, D. Hay, F. A. Kuipers, J. Oostenbrink, A. Valentini, and L. Rónyai, "Probabilistic shared risk link groups modeling correlated resource failures caused by disasters," pp. 2672–2687, 2021.

[4] S. Knight, H. X. Nguyen, N. Falkner, R. A. Bowden, and M. Roughan, "The internet topology zoo," *IEEE J. Sel. Areas Commun.*, vol. 29, no. 9, pp. 1765–1775, 2011.

[5] S. Gay, P. Schaus, and S. Vissicchio, "Repetita: Repeatable experiments for performance evaluation of traffic-engineering algorithms," *ArXiv e-prints*, 2017.

[6] B. Janzen, L. Richardt, and N. Aschenbruck. ELF Dataset: Earthquake-induced Link Failure Scenarios for Traffic Engineering Resiliency. OsnaData, V1, Accessed: 2026-04-29. [Online]. Available: https://doi.org/10.26249/FK2/Z8NISK

[7] J. Tapolcai, B. Vass, Z. Heszberger, J. Bíró, D. Hay, F. A. Kuipers, and L. Rónyai, "A tractable stochastic model of correlated link failures caused by disasters," in *Proc. of the IEEE Conference on Computer Communications (INFOCOM)*, 2018, pp. 2105–2113.

[8] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring isp topologies with rocketfuel," *IEEE/ACM Transactions on Networking*, pp. 2–16, 2004.

[9] R. Hartert, S. Vissicchio, P. Schaus, O. Bonaventure, C. Filsfils, T. Telkamp, and P. Francois, "A declarative and expressive approach to control forwarding paths in carrier-grade networks," *ACM SIGCOMM computer communication review*, pp. 15–28, 2015.

[10] M. Roughan, "Simplifying the synthesis of internet traffic matrices," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, pp. 93–96, 2005.

[11] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," in *Proc. of the IEEE Conference on Computer Communications (INFOCOM)*, 2000, pp. 519–528.

[12] C. Filsfils, P. Camarillo, J. Leddy, D. Voyer, S. Matsushima, and Z. Li, "Segment routing over IPv6 (SRv6) network programming," RFC 8986, 2021.

[13] C. Filsfils, K. Talaulikar, D. Voyer, A. Bogdanov, and P. Mattes, "Segment Routing Policy Architecture," RFC 9256, 2022.

[14] P. L. Ventre, S. Salsano, M. Polverini, A. Cianfrani, A. Abdelsalam, C. Filsfils, P. Camarillo, and F. Clad, "Segment routing: A comprehensive survey of research activities, standardization efforts, and implementation results," *IEEE Communications Surveys & Tutorials*, pp. 182–221, 2020.

[15] R. Bhatia, F. Hao, M. S. Kodialam, and T. V. Lakshman, "Optimized network traffic engineering using segment routing," in *Proc. of the IEEE Conference on Computer Communications (INFOCOM)*, 2015, pp. 657–665.

[16] T. Schüller, N. Aschenbruck, M. Chimani, M. Horneffer, and S. Schnitter, "Traffic engineering using segment routing and considering requirements of a carrier ip network," *IEEE/ACM Transactions on Networking*, pp. 1851–1864, 2018.

[17] A. Brundiers, T. Schüller, and N. Aschenbruck, "Live long and prosper–on the potential of segment routing midpoint optimization to improve network robustness," in *Proc. of the IEEE Conference on Local Computer Networks (LCN)*, 2024, pp. 1–9.

[18] J. Bogle, N. Bhatia, M. Ghobadi, I. Menache, N. S. Bjørner, A. Valadarsky, and M. Schapira, "TEAVAR: striking the right utilization-availability balance in WAN traffic engineering," in *Proc. of the ACM SIGCOMM*. ACM, 2019, pp. 29–43.

[19] S. Gay, R. Hartert, and S. Vissicchio, "Expect the unexpected: Sub-second optimization for segment routing," in *Proc. of the IEEE Conference on Computer Communications (INFOCOM)*, 2017, pp. 1–9.

[20] A. Brundiers, T. Schüller, and N. Aschenbruck, "Tactical traffic engineering with segment routing midpoint optimization," in *Proc. of the IFIP Networking*, 2023, pp. 1–9.

[21] A. Brundiers, T. Schüller, and N. Aschenbruck, "Fast reoptimization with only a few changes: Enhancing tactical traffic engineering with segment routing midpoint optimization," *IEEE J. Sel. Areas Commun.*, vol. 43, no. 2, pp. 495–509, 2025.

[22] P. C. Thenhaus, K. W. Campbell, W. Chen, and C. Scawthorn, "Seismic hazard analysis," *Earthquake engineering handbook*, pp. 1–50, 2003.

[23] P. Bird, D. D. Jackson, Y. Y. Kagan, C. Kreemer, and R. Stein, "Gear1: A global earthquake activity rate model constructed from geodetic strain rates and smoothed seismicity," *Bulletin of the Seismological Society of America*, pp. 2538–2554, 2015.

[24] W. H. 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.

[25] C. Pasolini, 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.

[26] J. Mayer, V. Sahakian, E. Hooft, D. Toomey, and R. Durairajan, "On the resilience of internet infrastructures in pacific northwest to earthquakes," in *Proc. of the International Conference: Passive and Active Measurement*, 2021, pp. 247–265.

[27] IBM. ILOG CPLEX Optimization Studio. Accessed: 2026-02-19. [Online]. Available: https://www.ibm.com/docs/en/icos