This is the author's accepted manuscript of the following book chapter: **Teresa Gomes, Lúcia Martins, Rita Girão-Silva, David Tipper, Alija Pašić, Balázs Vass, Luís Garrote, Urbano J Nunes, Martin Zachariasen, Jacek Rak. Enhancing Availability for Critical Services. In: Jacek Rak, David Hutchison (eds) Guide to Disaster-Resilient Communication Networks, chapter 22, pages 557-581. Computer Communications and Networks. Springer, Cham, 2020. DOI: 10.1007/978-3-030-44685-7_22**, which has been published in final form at https://link.springer.com/chapter/10.1007/978-3-030-44685-7_22.

Enhancing Availability for Critical Services*

Teresa Gomes[†] Lúcia Martins[†] Rita Girão-Silva[†] David Tipper[‡]
Alija Pašić[§] Balázs Vass[§] Luís Garrote[¶] Urbano Nunes[¶]
Martin Zachariasen[‖] Jacek Rak[#]

teresa@deec.uc.pt lucia@deec.uc.pt rita@deec.uc.pt dtipper@pitt.edu pasic@tmit.bme.hu
vb@tmit.bme.hu garrote@isr.uc.pt urbano@isr.uc.pt marz@itu.dk jrak@pg.edu.pl

January 2020

*This chapter is based on work from COST Action CA15127 ("Resilient communication services protecting end-user applications from disaster-based failures – RECODIS") supported by COST (European Cooperation in Science and Technology). This work is funded by ERDF Funds through the Centre's Regional Operational Program and by National Funds through the FCT – Fundação para a Ciência e a Tecnologia, I.P. under the project CENTRO-01-0145-FEDER-029312. This work was also partially supported by FCT under projects UID/EEA/50008/2019 and UIDB/00308/2020.

[†]University of Coimbra, Department of Electrical and Computer Engineering & INESC Coimbra, Rua Sílvio Lima, 3030-290 Coimbra, Portugal

[‡]University of Pittsburgh, Department of Informatics and Networked Systems, Pittsburgh, PA, USA

[§]MTA-BME Future Internet Research Group, High-Speed Networks Laboratory (HSN*Lab*) and Budapest University of Technology and Economics, Budapest, Hungary

[¶]Institute of Systems and Robotics, Coimbra, Portugal

[‖]IT University of Copenhagen, Denmark

[#]Gdansk University of Technology, Faculty of Electronics, Telecommunications and Informatics, G. Narutowicza 11/12, PL-80-233 Gdansk, Poland

Abstract

Traditional approaches to provide classes of resilient service take the physical network availability as an input and then deploy redundancy and restoration techniques at various layers, often without full knowledge of mappings between layers. This makes it hard (and often inefficient) to ensure the high availability required by critical services which are typically a small fraction of the total traffic. Here, the innovative technique of embedding a higher availability sub-structure, designated the *spine*, into the network at the physical layer, is explored. In the *spine*-based approach, it is considered that high availability must begin at the physical level, and then must be reinforced in upper layers. A recent disaster resilience framework, named FRAmework for DIsaster Resilience, which incorporates reliable network design (i.e., using the *spine*), disaster failure modelling and protection routing to improve the availability of critical services is discussed. Next, a proposal to select network links for availability upgrade to ensure high availability is presented. This is followed by a study assuming that if disaster-prone areas are known, they can be represented as obstacles which should be avoided when deploying the physical backbone of a communications network. Hence, a heuristic for a minimum cost Euclidean Steiner tree taking into account the presence of soft obstacles is presented.

1 The Spine Concept as an Approach to Increase Critical Services Resilience to Disasters

Society increasingly depends on communication networks for many aspects of daily life. For example, many critical infrastructures such as finance/banking, transportation, power grid SCADA systems, and emergency response depend on communication networks in order to function. The spine concept was proposed in [1] as an approach to efficiently provide classes of resilient service in backbone networks without modifying the topology. This was motivated in part by the need to constrain cost and that typically only a small fraction of the total traffic needs high reliability guarantees. To achieve clearly discriminated resilience classes, the concept proposed is to embed a higher availability set of links and nodes into the physical layer topology. This high availability sub-structure of the network is designated as the "spine". The spine serves as the underlying basis of resilience classes, as the availability differentiation at the physical layer can be leveraged with routing, restoration techniques, cross layer mapping and other methods to provide a wide range of availability guarantees. For example, one could route the most demanding quality of resilience class traffic on

the spine (or it could be used as backup if the routes in the spine have some lower key performance indicators).

Several methods can be used to create the spine at the physical layer. First, equipment that is more reliable (e.g., configured with redundant fans, hot standby line cards, etc.) can be utilised to construct the spine. Further, spare equipment can be deployed in parallel (e.g., hot standby fibre in physically diverse duct, redundant OXC, etc.) if needed. Moreover, the mean time to failure (MTTF) can be improved by applying additional strategies, like more robust and better insulated exterior cabinets, superior air conditioning, longer life backup batteries, underground cables instead of overhead lines, underground cables in ducts and at greater depth, communication towers resistant to intense winds, etc. Similarly, the mean time to repair (MTTR) of elements comprising the spine can be diminished using a number of techniques [2–4]. For example, one can adopt best practices and training procedures as specified by various government and trade organizations (e.g., NRIC, FCC, ATIS) and standards bodies (e.g., ITU, ETSI). Also, spare parts, equipment, software and test equipment can be distributed in a pre-calculated way to be readily available at spine locations. Further, the spine part of the network can also be monitored more closely by the network operations centre (NOC). Additionally, as shown in [2], the MTTR can be reduced by a careful deployment of the workforce to network operations, administration and management (OAM). The combination of techniques (i.e., hardware, equipment housing, workforce management, etc.) adopted to improve the reliability of the components making up the spine will depend on the cost versus benefit structure of the network owner.

The main outcome of adopting a spine is achieving *levels of availability differentiation at the physical level* which can be reinforced with restoration techniques, virtual network topology routing, cross layer mapping and other methods for a larger separation among resilience classes. Briefly summarizing, given a physical network topology graph $G = (N, E)$ where N is the set of nodes and E the set of links (undirected arcs), a high availability sub-structure $G_s = (N_s, E_s)$ is made part of the network, to increase the value of the relevant availability-based metric.

Here we illustrate the spine concept via a simple example. In Fig. 1 a full-mesh four node network is presented, which is used to illustrate the advantages of the spine approach. For clarity we assume nodes do not fail and only study the impact of adopting differentiated link availability. Further, we consider the case that the network uses path protection (i.e., disjoint working and backup paths) for each source-destination pair, to further improve the end-to-end availability. Specifically, we assume for each of the 12 source-destination

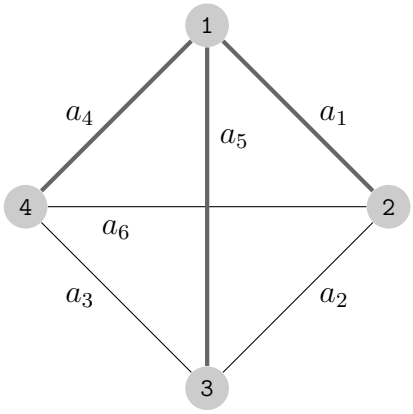

Figure 1: Full mesh network (thicker grey lines denote the spine)

pairs, the working path (WP) uses the direct single hop path and the backup path (BP) uses a disjoint two-hop path. The average of the end-to-end availability over all source-destination pairs will be denoted by A_S, and it can be obtained using series and parallel calculations. First, we study the case of homogeneous links, $a_e = a$, for all $e \in E$. Here, A_S is the result of the parallel combination of the single hop WP and the two-hop BP, the latter of which is a series system. Thus

$$A_S(a) = 1 - (1-a)(1-a^2) = -a^3 + a^2 + a \tag{1}$$

Now, consider the non-homogeneous edge availability case corresponding to the spine concept. In Fig. 1, the *spine* is the spanning tree made of edges `1`, `5` and `4`, marked with thicker grey lines. All edges on the spine have availability given by a_S (i.e., $a_1 = a_4 = a_5 = a_S$) and the remaining edges have availability equal to a_O (i.e., $a_2 = a_3 = a_6 = a_O$) The three edges on the spine result in six node pairs having a single hop WP on the spine and a BP with two hops, one of which is on the spine. Hence availability for these *s-d* pairs is $1 - (1 - a_S)(1 - a_S a_O)$. The other six node pairs have a two-hop WP on the spine and a BP, off the spine, with a single hop. Therefore the availability for these *s-d* pairs is $1 - (1 - a_S^2)(1 - a_O)$. Thus, the average end-to-end availability, as a function of a_S and a_O is given by: $A_S(a_S, a_0) = \frac{1}{12}\left(6\left(1 - (1 - a_S)(1 - a_S a_O)\right) + 6\left(1 - (1 - a_S^2)(1 - a_O)\right)\right)$. If we assume that $a_S = a + \Delta$ and $a_O = a - \Delta$, then A_S can be calculated as

$$A_S(a, \Delta) = -a^3 + (1 - \Delta)a^2 + (1 + \Delta)a + a\Delta^2 + \Delta^3 \tag{2}$$

Note that since $a_S = a + \Delta$ and $a_O = a - \Delta$, the network-wide sum of link availability and

Table 1: Effect of varying Δ on A_S and downtime

Case	A_S	Downtime (hours/year)
$a = .9,\ \Delta = 0$	.981	166.44
$a = .9,\ \Delta = 0.09$	.99712	25.23756
$a = .9,\ \Delta = 0.099$	.999701	2.61749

average link availability for the network with the spine are identical to the homogeneous network ($a_i = a$, $i = 1, \ldots, 6$). Hence if the the cost of increasing/decreasing availability is linear then both scenarios will present identical cost. We define δ as the difference in A_S between the spine and homogeneous scenarios, then $\delta = A_S(a, \Delta) - A_S(a)$, which can be shown to be $\delta = \Delta^3 + a\Delta^2 + a\Delta(1-a)$, and $\delta > 0$ if $\Delta > 0, a > 0$. Therefore, the existence of edges with non-identical availability leads to an average end-to-end availability which is greater than the end-to-end availability which would result from considering all edges with identical availability. Thus, the spine has the potential to improve the average end-to-end availability.

The impact of the variation of Δ on A_S and on the downtime per year is presented in Table 1 for the network of Fig. 1. Table 1 illustrates the potential improvement in A_S achieved by creating a spine. It is worth observing that with the spine, different s-d node pairs can have different availability levels. For example, if we take $\Delta = 0.09$ there are six single hop WPs on the spine with end-to-end availability of 0.998, whereas the six s-d pairs with two-hop WPs have end-to-end availability of 0.9962. Note that both single and two-hop WPs have an end-to-end availability greater than the uniform end-to-end availability achieved by the homogeneous case with $\Delta = 0$. It should be pointed out that although there are alternative solutions for the spine with maximal A_S (i.e., using edges `1`, `2`, and `6` results in the same A_S as in Fig. 1), the selection of the edges must be done carefully as other alternatives may result in a lower A_S (i.e., using edges `1`, `5`, and `6`).

In addition to improving A_S, the spine provides a wider range of availability options for deploying resilience classes. To illustrate this, consider two classes: class one only has a WP (i.e., no protection), whereas class two has path protection with a WP and BP that are disjoint. Table 2 presents the availability that a service provider can attain, in the network of Fig. 1, by routing the class one unprotected flows on paths with different number of hops. Next, we consider two groups: (a) the $s - d$ pairs with a direct one-hop WP route ON the spine {(1,2), (1,3), (1,4), (2,1), (3,1), (4,1)}; (b) the $s - d$ pairs with a one-hop WP route OFF the spine {(2,3), (2,4), (3,4), (3,2), (4,2),(4,3)}. The bottom two rows of Table 2

correspond to the spine-based network, where the routes can contain links that are ON the spine, OFF the spine or a mix of ON and OFF the spine links. It can be observed that in the spine case there is a greater availability range (i.e., max - min = 0.344 for spine vs. 0.171 for no spine case). The possible availability values for protected flows (i.e., with WP and BP) are shown in Table 3. The spine-based network, similarly to what was observed for the no protection case, provides a wider availability range (i.e., the range is 0.0633 for spine vs. 0.0171 for the no spine case). From the perspective of a service provider, the highest resilience class customers can be offered a superior level of resilience by combining routing and protection in the spine-based network, than what can be achieved in the no-spine network. Also, best effort lower paying customers will be offered a service without the increase in resilience which is the usual collateral effect of offering high availability resilience class.

Although availability is related to average performance under independent failures, the *spine* concept contributes to make networks more robust to disasters, by enhancing the availability of links and nodes subject to disaster-based challenges.

This chapter is organised as follows. In Sect. 2 a framework for disaster resilience, which consists of: a resilient network design model (using the *spine* concept to enhance network availability), a disaster failure model and a protection routing approach is proposed. This framework seeks to show the interdependence between these three elements, and how it can be explored to improve the availability of critical services. Next, in Sect. 3 a proposal on how to select edges for availability upgrade to ensure high availability is presented. These set of upgraded edges can be considered a *spine*. Finally, in Sect. 4, the knowledge about disaster-prone areas is taken into account in the design of the physical backbone of a communications network. Those critical areas are represented as obstacles which should be avoided, or that can be crossed at given cost (a soft obstacle). A heuristic for a minimum cost Euclidean Steiner tree representing the backbone network – the *spine* – taking into account the presence of soft obstacles is described and illustrative results are presented. Some concluding remarks can be found in Sect. 5.

2 Enhancing End-to-End Service Availability with the General Dedicated Protection and Spine

In this section, we discuss a novel framework for disaster resilience, called FRADIR (FRAmework for DIsaster Resilience) introduced in [5], which incorporates reliable network design,

Table 2: Flow availability values with no protection

Case	$s-d$ pairs	1-Hop WP	2-Hop WP	3-Hop WP
$a = .9$, $\Delta = 0$	All $s-d$ pairs the same	0.9	0.81	0.729
$a = .9$, $\Delta = 0.09$	6 $s-d$ pairs with 1-Hop WP ON spine	0.99	0.8019	0.64595
$-$	6 $s-d$ pairs with 1-Hop WP OFF spine	0.81	0.9801 or 0.6561	0.79388

Table 3: Flow availability values with protection

Case	$s-d$ pairs	1-Hop $\parallel$ 2-Hop	1-Hop $\parallel$ 3-Hop	2-Hop $\parallel$ 2-Hop
$a = .9$, $\Delta = 0$	All $s-d$ pairs the same	0.981	0.9729	0.9639
$a = .9$, $\Delta = 0.09$	6 $s-d$ pairs with 1-Hop WP ON spine	0.998	0.99645	0.960756
$-$	6 $s-d$ pairs with 1-Hop WP OFF spine	0.996219 or 0.934659	0.96084	0.993156

disaster failure modelling and protection routing so as to improve the availability of critical services. The method introduces a new probabilistic regional failure model, which takes into account not only the distance from the epicentre of the failure, but also the availability values of the network components. Based on the availability-aware disaster failure model, dedicated protection approaches are used to route the connection requests. The results show that with the interplay between protection routing, failure modelling and network update procedure the network performance can be significantly improved in terms of blocking probability and average resource consumption, which makes FRADIR an excellent competitor to provide disaster resilience in critical infrastructures in order to enhance the availability for critical services.

2.1 Motivation

There are three well-known research areas i.e., tools which can be utilised in order to enhance the reliability and availability of the network. These are the following:

1. Failure modelling: This research area deals with the question of how to model a failure event properly in order to find the appropriate failure mitigation technique. It is not a direct tool for enhancing availability, however it is a crucial task for that purpose. If we are able to model a failure event properly, we can locate the sensitive points in our network/system (see Chapter).

2. Network design: Network design investigates how to create or update a network in such a manner that we can satisfy certain criteria, e.g., a given level of availability or reliability can be enforced. It includes the resolution of problems such as where to add new links, or which network components should be upgraded in order to achieve higher availability. However, these topology upgrades are costly and time consuming. Furthermore, the majority of these approaches do not take into account how to route the connection requests themselves, they just provide the underlying network with given parameters (see Chapter).

3. Protection Approaches, i.e., Survivable Routing: This research area deals with the question of how to route connection requests through the network, in order to satisfy certain requirements, e.g., minimise the capacity consumption, or protect the connections against certain failure events (for example regional failures) (see Chapter).

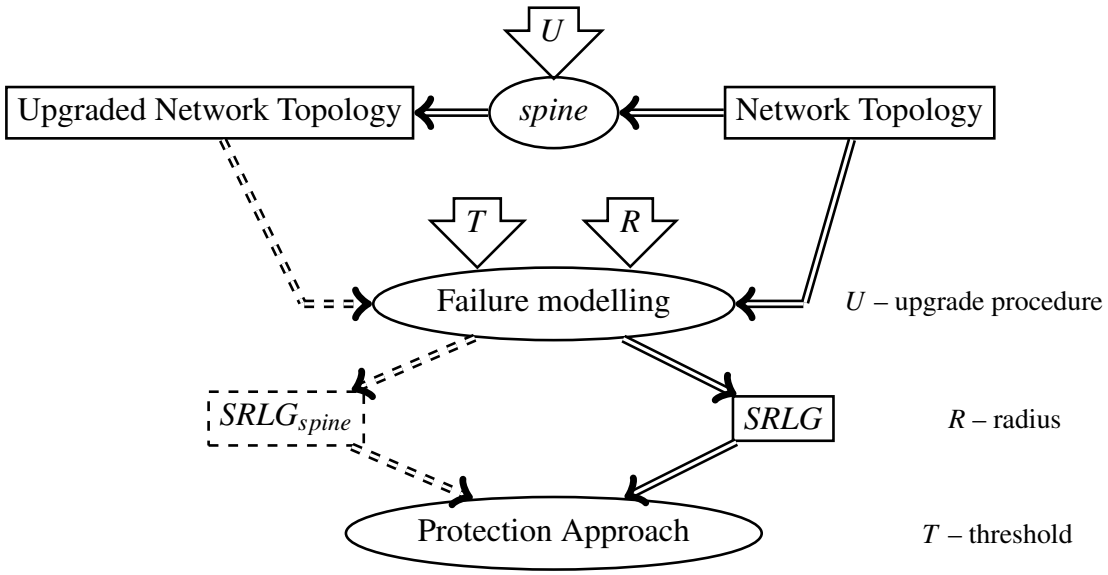

Figure 2: The concept of FRADIR: dashed lines represent the scenario with the *spine* topology upgrade; the full lines represents the scenario without the *spine* concept

In this section, we present the FRAmework for DIsaster Resilience (*FRADIR*), which includes network design, failure modelling and protection routing for the improvement of disaster-resilience of mission-critical applications.

2.2 FRADIR – Disaster-resilient Transport Networks

A truly disaster-resilient network cannot be created in an efficient manner if the focus is solely on the network design, the routing problem or the failure modelling. In fact, these problems have to be jointly addressed. In order to achieve this, the FRADIR integrates three well established methods: (*i*) for networks design, the *spine* (see Sect. 3 and Sect. 4); (*ii*) for failure modelling the so-called regional failure modelling [6]; (*iii*) for the protection itself the GDP (General Dedicated Protection) [7]. The basic concept of FRADIR is shown in Fig. 2.

The *spine* concept [1] is used to compute a high availability structure, in this case a spanning tree, in the topology. The link availability on the tree is considered to be increased by a certain amount, which is represented by the upgrade procedure U in Fig. 2. A description on the spine concept can be found in Sect. 1, and details of the spine design used by FRADIR can be found in [8].

In the next step, failure modelling is utilised, namely the concept of *Probabilistic SRLGs*

(PSRLGs) is used, where a probability of failing is associated to every SRLG. The list of SRLGs is generated, based on the distances of edges from the epicentre of possible disasters and the availability of the given component. Afterwards, the PSRLGs with a failing probability higher than a pre-defined threshold (T) are selected (details in Sect. 2.2A). To protect against failures in the SRLG list, 1+1 and generalised dedicated protection approaches are used to route the connections (details in Sect. 2.2B).

A. Unavailability-based Regional Failure Model

In this subsection, the novel method for modelling regional failures introduced in [5] to assess the benefits of the *spine* is presented. The input to the regional failure model consists of a graph G, a maximal radius of the failure $R \geq 0$ and a threshold $T \in [0, 1]$, while the output is a list of SRLGs containing all the SRLGs with a probability of failure above T. Note that a high T value yields a list of a few probable SRLGs (e.g., single link failures outside the spine), while a low threshold leads to the list of nearly all the edge sets that can be hit by a disk of radius $r \leq R$, which may include highly improbable scenarios.

This model is a modification of the one in [6], where failing probabilities are generated based on the distances of edges from the disaster epicentre. In this case, the failing probability of each PSRLG is modified, taking into account the unavailability values of its edges. The result is that the probability of PSRLGs containing spine edges is reduced, allowing to access the benefits of the *spine* concept. As a final step, the list $\mathcal{F}$ of SRLGs with a failing probability higher than a threshold T is considered, as these SRLGs are the most probable ones.

In our model, disasters d have an epicentre $\mathcal{P}$ (random variable) taking values $p \in \mathbb{R}^2$, with a shape overestimated by a circular disk of radius $\mathcal{R}$ (random variable) taking values $r \in [0, R]$, where R is the maximum range of disasters to be protected. We define $h(p)$ and $g(r)$ to be the density functions of disaster epicentre and disaster range, respectively. Let $U(e)$ be the unavailability of each link $e \in E$ (equal to 1 minus the availability value). The $U(e)$ values are scaled with a constant factor so that the average of the normalised unavailabilities of the edges, $\overline{U}(e)$, is 1.

Finally, we consider $I_{S,p,r}$ as indicator variables which are 1 if the disk with centre p and radius r hits all the edges of a set $S \subseteq E$, and 0 otherwise. Therefore, in our model, the probability of failing of a link set S is

$$P(S \text{ is hit}) = \prod_{e \in S} \overline{U}(e) \int_{p \in \mathbb{R}^2} \int_{r \in [0,R]} I_{S,p,r} g(r) \mathrm{d}r \; h(p) \mathrm{d}p. \tag{3}$$

The precision of the results is not affected by a sufficiently fine discretization, as some imprecision in the available network data is expected. With this in mind, the problem is discretized by defining a sufficiently fine resolution, say 1 km, and placing a grid of 1 km $\times$ 1 km squares over the plane to assume that the values of the inner integrals (i.e., $\int_{r\in[0,R]} I_{S,p,r} g(r)\mathrm{d}r$) are almost identical for every p inside each grid cell c. This way, the computationally hard integration problem becomes a simple summation. Due to the lack of space, detailing the discretization is omitted here. However, we mention that as the failure probability defined by Eq. (3) is almost identical to the one in [6], it can be done similarly as in the cited paper.

B. Dedicated Protection Approaches

It is a well known fact that in transport networks, even a short disruption may result in a huge amount of data loss. To avoid this, *instantaneous recovery* is required in today's transport networks, i.e., the recovery time should be held under 50 ms in order to ensure a seamless operation in the event of failures [9]. To fulfil these strict requirements, *dedicated protection* is implemented (see Chapters 3.0, 3.1).

FRADIR utilises the so-called General Dedicated Protection (GDP) introduced in [7]. GDP enables instantaneous failure recovery against arbitrary failure patterns listed in the SRLG list ($\mathcal{F}$). This SRLG list is generated in our case by the probabilistic failure model introduced in the previous section. The GDP is able to protect all the protectable[1] SRLGs by generalizing the rigid SRLG-disjoint path structure of 1+1 to an arbitrary directed acyclic graph between the source and destination nodes. In addition the GDP minimises the total bandwidth cost providing the optimal minimum cost survivable routing solution. Note that in FRADIR a non-bifurcated scenario (called GDP-R) was used, i.e., we considered that all the data is sent along all links of the connection as in 1+1. For ease of notation, we simply mention GDP.

The concept (i.e., framework) of FRADIR is very flexible and allows us to change any of its methods or components easily. This means that a new network design, other failure modelling methods or routing approaches could be easily incorporated into the framework.

[1]We call a failure $f \in \mathcal{F}$ *protectable* if the network topology remains $s-d$ connected after removing the links in f.

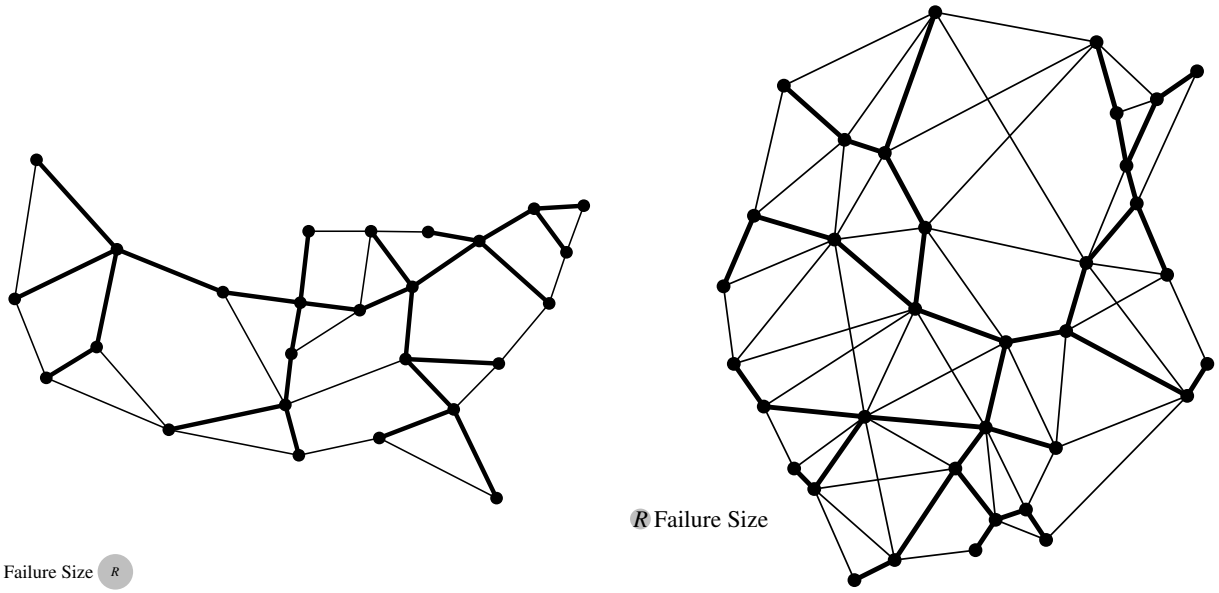

(a) US (26 nodes, 42 edges) [10] (b) India (35 nodes, 80 edges) [11]

Figure 3: The investigated networks: to visualize the size of the regional failure ($R = 50$) the radius is displayed as a gray circle; the edges in bold are in the spine [5]

2.3 The Effectiveness of FRADIR

Some experimental results of the FRADIR framework are in [5]. Two real-world topologies – see Fig. 3 – with different threshold values T and upgrade possibilities (corresponding to different values of U), leading to instances with different number of SRLGs, were investigated. Two performance measures were calculated: the average capacity allocated per connection and the blocking probability of the protection approaches (1+1 protection and GDP-R) with and without the upgraded availability values of the *spine* links. The 1+1 was calculated with the two-step approach. After calculating the working path, we try to calculate an SRLG-disjoint path. The connection is blocked if no SRLG-disjoint path pair is found.

Two different scenarios were devised: (i) all SRLGs generated by the probabilistic failure model are included in $\mathcal{F}$, i.e., the unprotectable SRLGs are also considered, which may lead to situations where the topology graph does not remain connected after a failure occurs; (ii) only the protectable SRLGs are included in an SRLG list $\mathcal{F}'_{s-d} \subseteq \mathcal{F}_{s-d}$ for every $s-d$ pair, so that the network remains $s-d$ connected if failure $f \in \mathcal{F}'_{s-d}$ occurs. In this second scenario, a disaster-resilient sub-graph for the connection may always be found by the GDP approach. Further details on the experimental setup are available in [5].

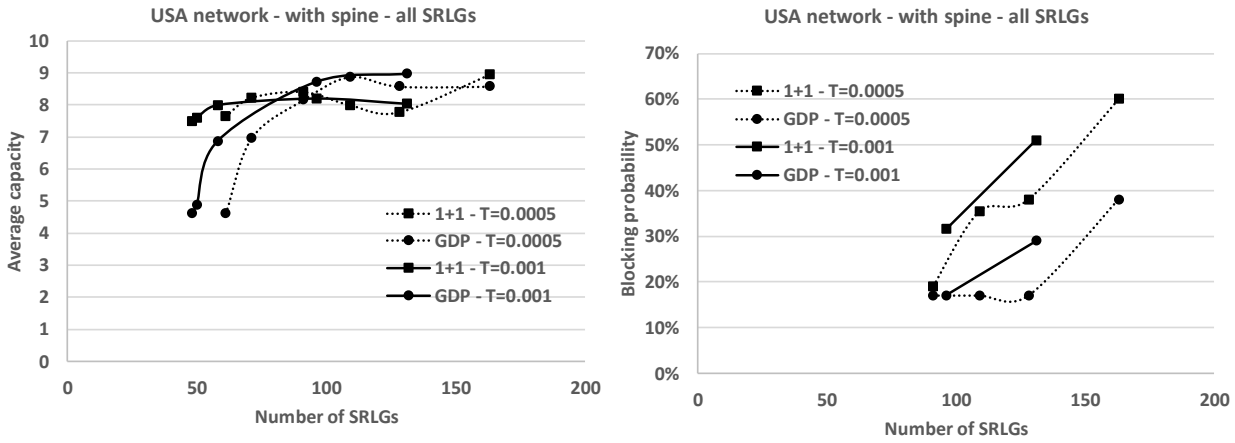

(a) Average capacity

(b) Blocking probability

Figure 4: Results for the USA network, considering $R = 50$ km in scenario (i)

The number of listed SRLGs depends on the radius of failure event R, the threshold value of SRLG failing probability T, and of the use of the upgrade procedure. Hence, a larger number of SRLGs corresponds to a more severe failure situation (larger R) and/or a lower T and/or a lower level of network upgrade. Instead of considering all the identified SRLGs, one could have considered only maximal SRLGs (that is, SRLGs which are not contained in any other SRLG).

The results for the USA network (26 nodes, 42 edges) considering the spine are displayed in Fig. 4 (first scenario) and in Fig. 5 (second scenario). Note that in Fig. 5(b), the blocking probabilities for the GDP are not displayed because they are always 0. Without the spine, the following results were obtained: for scenario (i), the average capacity varies between 8.62 and 12.92, and the blocking probability varies between 70.5% and 100%, for a total of 450 SRLGs (if $T = 0.0005$) or 293 SRLGs (if $T = 0.001$); for scenario (ii), the average capacity varies between 7.00 and 13.65, and the blocking probability varies between 0 and 99.5%, for a total of 435 SRLGs (if $T = 0.0005$) or 283 SRLGs (if $T = 0.001$).

The results for the average capacity allocation in the India network (35 nodes, 80 edges) considering the spine are displayed in Fig. 6 for both scenarios. The blocking probabilities are not displayed, because they are always 0, except for $T = 0.0005$ and higher number of SRLGs (they are under 10% in this case). Without the spine, the following results were obtained: for scenario (i), the average capacity varies between 7.57 and 8.16, and the blocking probability varies between 18.0% and 69.0%, for a total of 1126 SRLGs (if $T = 0.0005$) or 410 SRLGs (if $T = 0.001$); for scenario (ii), the average capacity varies

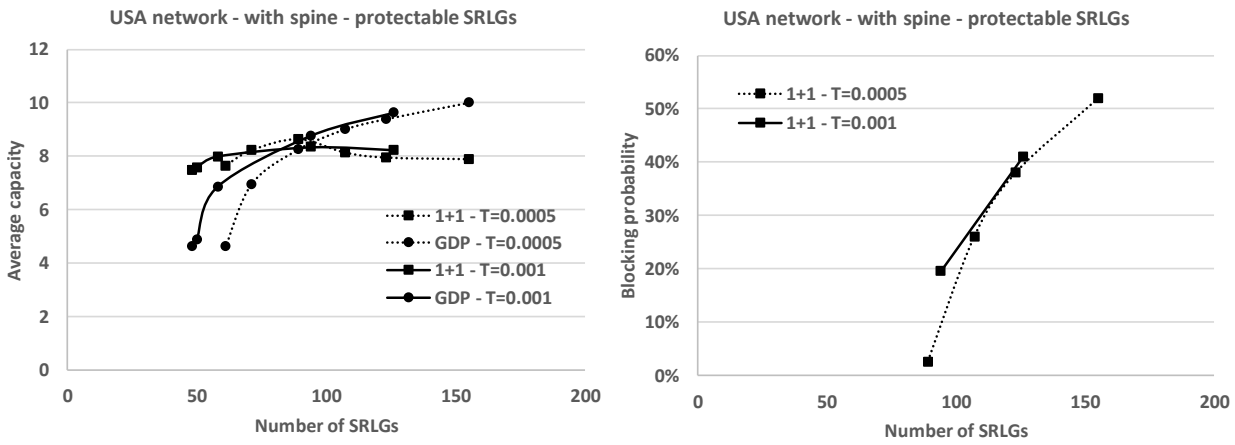

(a) Average capacity

(b) Blocking probability

Figure 5: Results for the USA network, considering $R = 50$ km in scenario (ii)

between 7.75 and 9.55, and the blocking probability varies between 0 and 57.5%, for a total of 1114 SRLGs (if $T = 0.0005$) or 406 SRLGs (if $T = 0.001$).

In terms of blocking probability, it has a very high value when the networks are not upgraded (i.e., without the *spine*). The exception is the case when only the protectable SRLGs are considered (scenario (ii)) and GDP is used, as the blocking probability is 0. When the spine is considered, the blocking probabilities are lower and the GDP clearly presents better results than the 1+1. In some cases, see Fig. 5(b), the blocking probability for 1+1 goes up to 52%, whereas for the GDP, it is is 0 (not shown in the figure). A final note to mention that for higher thresholds, the blocking probability is lower.

The use of the *spine* concept to update the network allows the FRADIR to reduce the number of link sets (all listed as SRLGs here) which fail with a larger probability than the considered threshold (T). An example is for the USA network in both scenarios, where the number of SRLGs decreases to about a third or lower (depending on the upgrade factor – see Figs. 4-5), just by using the spine. When the spine is used, the number of SRLGs is similar in both scenarios, i.e., whether all SRLGs or only the protectable SRLGs are considered.

The results show that the GDP is more efficient than the 1+1 in terms of average capacity allocation (per connection) in situations of equal blocking probability. This is due to the flexible structure of GDP, which manages to protect against failures without having to deploy long disjoint backup paths.

The results illustrate the benefits and the potential of FRADIR. The network per-

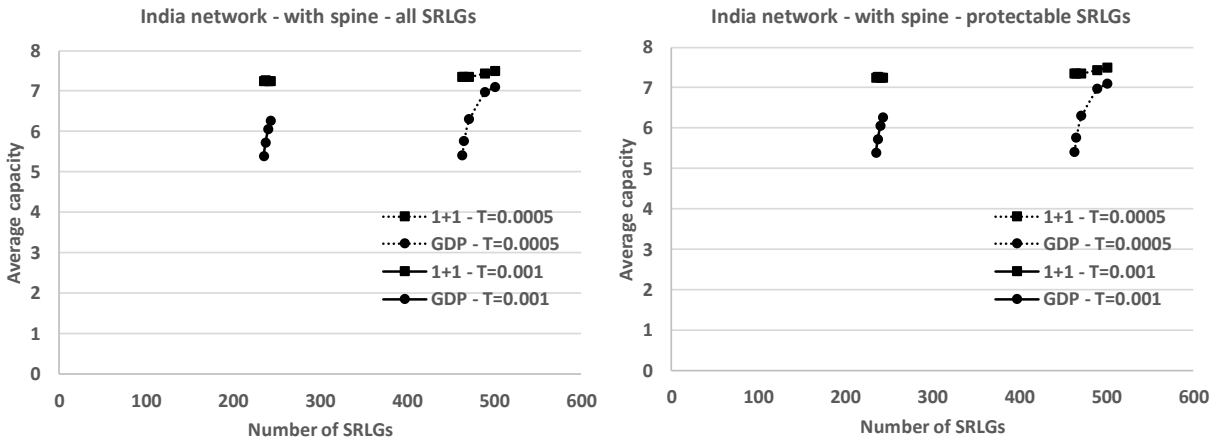

(a) Scenario (*i*)

(b) Scenario (*ii*)

Figure 6: Average capacity results for the India network, considering $R = 50$ km

formance is improved when different aspects such as reliable network design, appropriate modelling of disaster failures and adequate protection routing scheme are incorporated in an unified and coherent way. The impact of this framework may be even higher if the GDP is to consider the location of the more reliable links of the spine, in order to further minimise resource consumption.

Note that in [12] an improved framework for disaster resilience was proposed, i.e., an extension to FRADIR [5]. The major additions were that a cost function to select the high availability links on the spine [13] was introduced making the network design approach more realistic. The failure modelling, i.e., the identification of relevant SRLGs was refined with the help of a new regional failure model, where the availability of each link is translated into information regarding the distance of the link from the disaster epicentre area. In addition, a heuristic to select and upgrade a set of links to ensure no SRLG contains a cut-set was proposed. In other words, FRADIR-II ensures that the network remains connected in case of the most probable failure scenarios (according to the failure model).

3 Network Upgrade for Geodiverse Routing with Availability Constraints

The interdependence of critical services and communication networks drives the increasing need for the design of resilient networks. Service Level Agreements for some services (and

node pairs) may require a high level of availability. Geodiversity [14, 15] seeks to ensure the active and backup path are geographically separated, thus reducing the impact of geo-correlated failures. A pair of working and backup paths established between certain two end nodes is called D-geodiverse, if the minimal distance between any intermediate node or edge of the active path and any node or edge of the backup path is at least D. The formal definition of D-geodiversity used in the present work, can be found in [16] (and also in Chapter 3.2). Geodiverse routing is very useful for providing a pre-planned protection of communication paths against multiple correlated failures occurring, e.g., due to a natural disaster such as an earthquake or a volcano eruption with a specified size of a disaster region. A comprehensive survey on strategies for communication networks to protect against large-scale natural disasters can be found in [17], including references to works on geographically-diverse routing.

The objective is to improve the network resilience to disasters by selecting a subset of edges of total minimal length, for availability upgrade, to ensure that using 1+1 protection, a D-geodiverse path pair, for a given set of end-to-end connections, exists with a desired level of availability. Algorithm GAPGP described in [18, 19] (and also in Chapter 3.2) is instrumental to identify the connections that need to be improved. We will assume that, for a given network topology, the geographical distance between nodes and edges and between edges has been previously determined.

A larger D will require longer path(s), hence with lower availability, but on the other hand, it will tend to make the connection resistant to events with broader area impact. The maximal geodiversity value D_{sd} for a particular node pair (s, d) is determined by the location of the nodes and the geographical paths of the links – see Fig. 7. Therefore, when D is unfeasible for a node pair (s, d), the maximum geodiversity value, D_{sd} for that particular node pair is used instead. Hence the D-geodiverse problem always has a solution in bi-connected networks (i.e., no single node removal will make two nodes mutually unreachable). For a formal definition of D_{sd} see [16].

For simplicity, it is assumed that when an edge is upgraded, its new availability is equal to the value that would have been obtained placing a parallel edge, with the same availability as the original edge. Moreover, it is assumed that each edge can only be upgraded once.

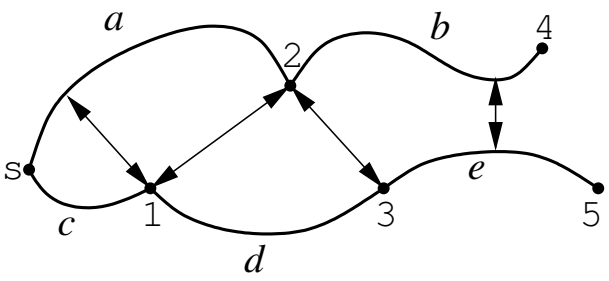

Figure 7: Distances between nodes and edges relevant to determine D_{sd} in part of a network where only `s` is marked – based on Fig. 3 of [16]

3.1 Additional Notation

Each edge $e \in E$ is characterised by an initial availability $a_e \in [0, 1]$ and a cost $c_e > 0$ to upgrade the link to an availability a_e^u. The variable x_e takes the value 1, if edge e has been upgraded, and 0 otherwise. The network is assumed to be bi-connected. Let K denote a set of node pairs of interest. The set of all node-disjoint path pairs between a given node pair s and d in G is represented by R_{sd} and each path pair $r \in R_{sd}$ is made of two sets of edges: S_{r1} and S_{r2}, the edge set of the first and second path of r, respectively. The availability of a path pair $r \in R_{sd}$ is denoted by Λ_r.

3.2 A Heuristic Approach

Given a set of node pairs K, a distance D and a desired level of availability Λ, the heuristic starts by calculating the most available path pair for all node pairs for which no D-geodiverse path pair with the desired availability Λ exists in the present network. Then, an edge belonging to one (or more) of those paths is selected for upgrade. Several criteria for selecting this edge can be considered and will result in different variants for the proposed heuristic. After each edge is upgraded, the procedure is repeated until all path pairs satisfy the desired availability and geodiversity constraints, or all edges have been upgraded.

The above described resolution approach is presented in [18], and revisited in Minimum Upgrade Cost with Availability and Geodiversity (MUCAG) – Alg. 1. This algorithm requires the Guaranteed Available Pair of Geodiverse Paths (GAPGP) algorithm proposed in [18] and reviewed in Chapter 3.2). For a given (s, d) pair, this algorithm calculates a path pair $r \in R_{sd}$ such that Λ_r is greater than or equal to the desired availability Λ, while respecting the D_{sd} geodiversity constraint. If no such path pair can be found, it returns the most available path pair among those that are D-geodiverse.

Algorithm 1 MUCAG

Require: G, K, Λ, $(a_e, a_e^u, c_e) : \forall e \in E$, $D_{sd} : \forall (s,d) \in K$
Ensure: $x_e : \forall e \in E$

for all $e \in E$ **do**
  $x_e \leftarrow 0$ ▷ Initially no edge is upgraded
end for
repeat
  $K' \leftarrow \emptyset$ ▷ Set of (s,d) pairs for which Λ was not reached – initially it is empty
  $\mathcal{R} \leftarrow \emptyset$ ▷ Set of the most available path pairs, for each $(s,d) \in K'$ – initially it is empty
  for all $(s,d) \in K$ **do**
    $(r, \Lambda_r) \longleftarrow \text{GAPGP}(s, d, D_{sd}, \Lambda, x_e : e \in E)$ ▷ If $\Lambda_r \geq \Lambda$ then (s,d) is solved
    if $\Lambda_r < \Lambda$ **then**
      $K' \leftarrow K' \cup \{(s,d)\}$ ▷ Inserts (s,d) pair in K' because Λ was not reached
      $\mathcal{R} \leftarrow \mathcal{R} \cup \{r\}$ ▷ Stores in $\mathcal{R}$ the most available path pair for (s,d)
    end if
  end for
  $K \leftarrow K'$ ▷ K is the set of unsatisfied node pairs
  if $|K| > 0$ **then** ▷ If the set of unsatisfied node pairs is not empty
    $(e', K') \leftarrow \text{selectEdge}(\mathcal{R}, K, x_e : e \in E)$ ▷ Selects edge e' for upgrade and updates K'
    $x_{e'} \leftarrow 1$ ▷ Signals that e' has been upgraded
    $K \leftarrow K \setminus K'$ ▷ K is updated due to the upgrade in e' but based on the routes in $\mathcal{R}$
  end if
until $K = \emptyset \vee \left(\sum_{e \in E} x_e = |E|\right)$ ▷ All node pairs have availability Λ or all edges were upgraded

In each iteration of the main cycle, two major tasks are performed:

1. the identification of the node pairs for which no D-geodiverse path pair has an availability greater than or equal to Λ (see lines 7-13);

2. the selection of the edge to upgrade e' (see lines 15-19). In function *selectEdge*, after e' is upgraded, the node pairs whose routes $\hat{r} \in \mathcal{R} : e' \in S_{\hat{r}1} \cup S_{\hat{r}2}$, have $\Lambda_{\hat{r}}$ greater than or equal to Λ, are placed in K', which allows for a further reduction of the size of K (line 18). Note $\mathcal{R}$ is defined in line 6 of the algorithm and constructed in line 11.

3.3 Selecting the Edge to Upgrade

Several strategies were considered [18,19] for iterative selection of an edge to upgrade. Let $E_0(\mathcal{R})$ be the set of all edges in the paths belonging to $\mathcal{R}$ with $x_e = 0$ (i.e., the set of edges in the path pairs in $\mathcal{R}$ that can still be upgraded). An additional set is defined: $E_P(\mathcal{R})$, the subset of edges of $E_0(\mathcal{R})$ with maximal number of occurrences in $r \in \mathcal{R}$.

The following three strategies are discussed in this section:

- **Min-Cost:** Select the edge in $E_0(\mathcal{R})$ with minimum upgrade cost.

- **Min-Cost–Max-Count:** Select the edge in $E_P(\mathcal{R})$ with minimum upgrade cost.

- **Max-On–Max-Count:** Select the edge in $E_P(\mathcal{R})$ that, if upgraded, maximizes the number of elements of $\mathcal{R}$ whose availability becomes at least Λ. If no such edge exists, select any edge in $E_P(\mathcal{R})$.

3.4 Computational Results

In [16], two networks, representative of typical telecommunication transport networks, were considered. Here we revisit the results obtained for only one of them, the Germany50 topology (geographical location of nodes available in [11]). Germany50, as its name indicates has 50 nodes, 88 edges (or links) and an average node degree of 3.52. It is assumed that links follow the shortest path over the terrestrial surface assuming that the Earth is a sphere, resulting in an average link length of 100.67 km, where the longest link has 252 km. The maximum geodistance that can be achieved, for any source-destination in Germany50, is 166 km – this value was obtained solving the Maximum Distance D of Geodiverse Paths (MDDGP) optimization problem [16]. The information concerning edge

lengths and geographical distance between edge-edge pairs and node-edge pairs is available at [20].

The edge availability values, a_e of $e \in E$, were calculated based on the edges length [9]:

$$a_e = 1 - \frac{MTTR}{MTBF_e} \tag{4}$$

with

$$MTBF_e[\text{hrs}] = \frac{CC \times 365 \times 24}{\ell_e} \tag{5}$$

where $MTBF$ and $MTTR$ are the mean time between failures and the mean time to repair in hours, respectively. CC is the cable cut metric and ℓ_e is the cable length of $e \in E$, both in km. As in [21], we consider $MTTR = 24$ h and $CC = 450$ km). As already stated, it was assumed an upgraded availability a_e^u for each edge $e \in E$ equivalent to the addition of a parallel edge of the same length, i.e. $a_e^u = a_e\,(2 - a_e)$.

Given the difficulty of determining the cost upgrade of a link [22, 23], here we simplify and consider the upgrade cost of each edge is given by its length. Nevertheless, results are analyzed not only considering the total length of the upgraded edges but also regarding the number of upgraded edges.

The minimum required availability was set to $\Lambda = 0.99999$. Table 4 displays the number of upgraded edges, the upgrade cost (the corresponding length in km for those edges) and the CPU time in s for the geodiversity values D of 40 km, 80 km, 120 km and 160 km. In the table, the solutions showing the minimum upgrade cost are highlighted in bold. The CPU times refer to a server with a Intel Xeon CPU X5660 @ 2.80GHz, with 6 cores and 48 GB of RAM.

Table 4 shows that Max-On–Max-Count is the best approach, because it presents the lowest cost (i.e., total length of the selected edges) and also the lowest number of edges selected for upgrade. Note that in [18, 19] other strategies are presented that led to lower values of the number of updated edges albeit with higher cost. However, the number of upgraded edges of Max-On–Max-Count exceeds, at most, 16% of the minimum values found with the other tested strategies. Therefore, one may conclude that Max-On–Max-Count is the best strategy for Germany50. The Min-Cost solutions have the largest number of upgraded edges and also the largest total length, hence it is clearly the strategy with worst performance.

A larger D requires a higher number of upgraded edges and also a higher cost and a higher CPU time, because the number of iterations of MUCAG (see Alg. 1) is determined

Table 4: Number of upgraded edges ($\#\mathcal{E}$), cost (in km) and CPU time (in s) for achieving $\Lambda = 0.99999$ between all node pairs in Germany50

Algorithm	$D = 40$ km			$D = 80$ km			$D = 120$ km			$D = 160$ km		
	$\#\mathcal{E}$	Cost	CPU	$\#\mathcal{E}$	Cost	CPU	$\#\mathcal{E}$	Cost	CPU	$\#\mathcal{E}$	Cost	CPU
Min-Cost	40	2561	77.9	52	3684	3731.5	52	3684	5261.0	55	4046	6730.5
Min-Cost–Max-Count	18	1507	13.5	25	2289	139.8	32	2843	447.0	37	3285	648.0
Max-On–Max-Count	16	**1469**	12.7	22	**2260**	137.9	28	**2839**	428.3	29	**2914**	573.0

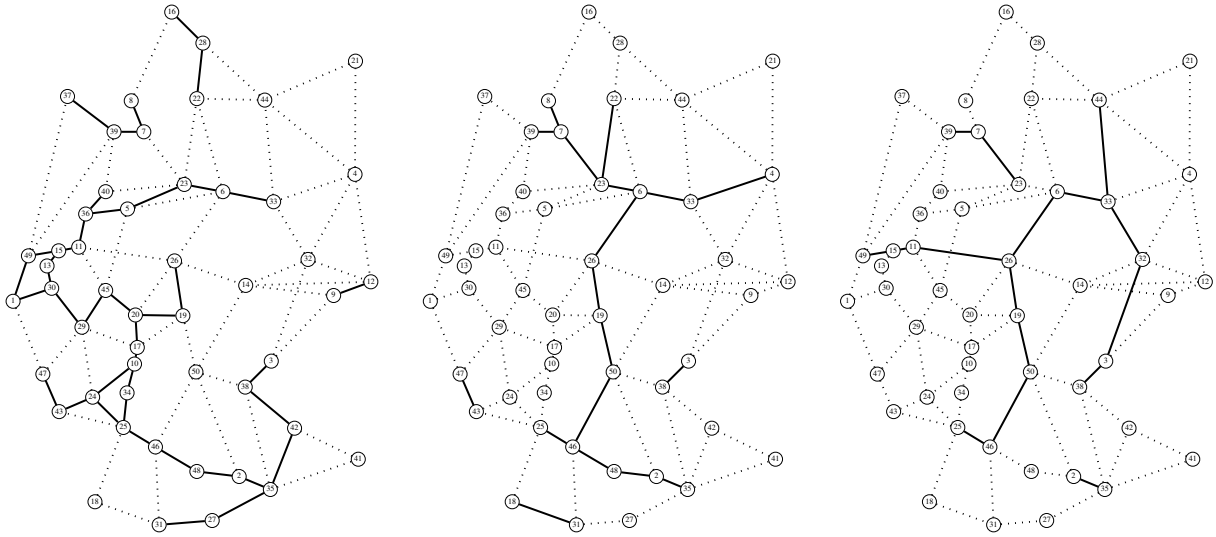

(a) Min–Cost (b) Min-Cost–Max-Count (c) Max-On–Max-Count

Figure 8: $D = 40$ [km]: upgraded edges for achieving availability $\Lambda = 0.99999$ between all node pairs in Germany50 are shown as solid lines

by the number of upgraded edges. Note that an analysis of the upgraded edges for different values of D in a certain method shows that the set of edges for a given D does not necessarily contain the set of selected edges for a smaller D.

In Fig. 8, the results for the worst strategy (Min-Cost), for the strategy with smaller cost (Max-On–Max-Count) and for the strategy Min-Cost–Max-Count are displayed. The Min-Cost solution (Fig. 8(a)) illustrates the handicap of this approach as it tends to select a larger number of shorter edges, which eventually results in a larger upgrade cost. The solutions of the other two strategies point out one main feature, also visible in Table 4. Strategies that favor the edges in $E_P(\mathcal{R})$ (i.e., the most frequent edges in the path pairs of $\mathcal{R}$), the Min-Cost–Max-Count (Fig. 8(b)) and the Max-On–Max-Count (Fig. 8(c)) lead to solutions with a lower upgrade cost, but higher number of upgraded edges when compared with other strategies, presented in [18, 19], that favour the minimum number of edges to be upgraded.

4 Exploring the Spine Concept in Disaster-prone Areas

The work presented next applies to the design of a spine from scratch or to the design of new branches of an existing spine. The purpose is to obtain a minimum-length tree connecting a given set of nodes and avoiding certain critical areas such as disaster-prone areas. The problem of obtaining a minimum-length tree in the Euclidean plane is the well-known Euclidean Steiner problem [24]. An extension of this problem is the one of obtaining the Euclidean Steiner tree avoiding solid obstacles [25]. The present work is based on the assumption that instead of simply avoiding obstacles the Euclidean tree can traverse them with an increased cost per unit of length. Such obstacles are called soft obstacles, in contrast to solid obstacles, which cannot be traversed. The consideration of soft obstacles was already done in the context of rectilinear Steiner trees [26] but to the best of our knowledge the first time they were considered in the context of the classical Euclidean Steiner trees was in [27]. In this work soft obstacles represent geographic areas prone to natural disasters such as floods, hurricanes, fires and so on but they can also represent geographic obstacles such as rivers, mountains or sea, among others. Therefore, the present work aims to obtain a spine which is a minimum-cost (or weighted-length) Euclidean Steiner tree, taking into account the presence of soft obstacles, that corresponds to the final layout of fibers or microwave links constituents of the spine.

The proposed strategy can also be applied in other contexts of communication networks design or even in a more general context of transport networks design.

4.1 Problem Definition

Given a set of terminal nodes T and a set of soft obstacles O, which are non-self-intersecting polygonal regions with an associated weight (or cost per unit of length) in the Euclidean plane, the goal is to find the weighted minimum-length Euclidean Steiner tree. Soft obstacles can be homogeneous and non-homogeneous. Homogeneous soft obstacles are the ones for which the weight is constant inside the obstacle. In the context of this work only homogeneous obstacles were considered. Each obstacle is defined by their extreme 2D Cartesian points and an associated weight c_o where $o \in O$. If two or more obstacles overlap, the weight is the highest among the overlapping polygons. The cost or weighted length of the Euclidean Steiner tree is then defined as the Euclidean distance multiplied by the underlying weight of each region represented.

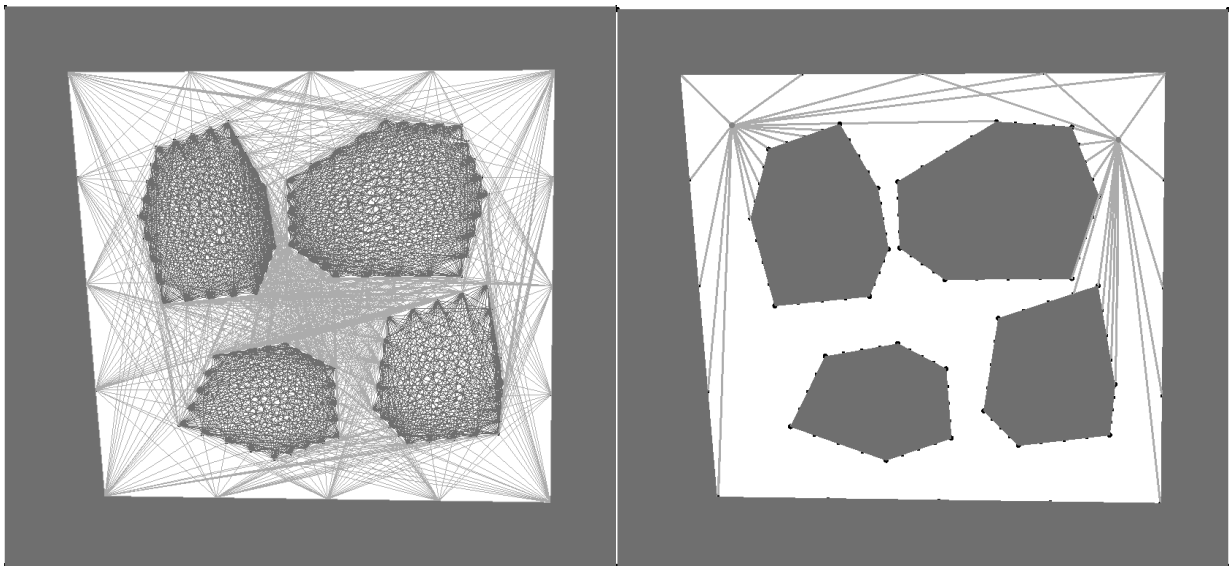

(a) Inter-obstacle connections (b) Node-obstacle connections

Figure 9: Interconnections involved in shortest path computation

In order to obtain the desired tree it is important to tackle two important problems taking into account soft obstacles: the shortest path and the shortest tree between three terminal nodes.

4.2 The Weighted Euclidean Shortest Path Problem

The weighted Euclidean shortest path is defined as the weighted minimum-length path between two points taking into account the weight of each soft obstacle. The computation of this path is inspired by the strategy proposed in [28] in which each obstacle's edge is uniformly discretized. Dijkstra's algorithm is then used in a graph were the edges inside obstacles, between obstacles and between the two nodes were considered. In scenarios where multiple obstacles exist, the computation of this graph can be heavy particularly for non-convex polygons although the edges between obstacles and inside obstacles can be precomputed for each scenario (see Fig. 9(a)). Computing all the connections between the obstacles and the two nodes (see Fig. 9(b)) can be achieved using line intersection algorithms such as Bresenham's [29] to discard connections that cross different costs regions (e.g., free space and inside each soft obstacle). For a small number of obstacles, this can be a fast approach but as the number of obstacles increases, the computational cost increases rapidly meaning that with multiple obstacles or obstacles with an high number of edges

this method will be computationally heavy. Even if we restrict to the line of sight of the line intersection algorithm, we still require multiple computations of intersections.

To speed up the the computation of the shortest path between each pair of nodes, two heuristics strategies were developed as presented in [27]. The first one is based on an adaptive grid for each obstacle to allow a dynamic adjustment of the cell size to better fit to the obstacle's shape. The second one is based on intersections between bounding circumferences for each obstacle and the circumference containing the two end nodes. This heuristic strategy allows to discard interconnections in the graph that are not in line of sight or are not important for the shortest path computation.

4.3 The Weighted Euclidean Shortest Tree Between Three Terminal Nodes

The minimum-length tree that connects three points in the Euclidean plane can be obtained solving the Fermat–Torricelli problem [24]. If all the interior angles of the triangle defined by the three points are less than 120° then the tree has and additional point – a Steiner point – that can be obtained exactly if no obstacles were involved. In the presence of soft obstacles, the problem of obtaining the weighted minimum-length Euclidean Steiner tree (or the minimum-cost tree) connecting tree points is rather difficult as is illustrated in Fig. 10. In the left top corner, the three points and the soft obstacle are represented. The rest of the figures represent the best solution found for each cost considered for the obstacle and also the variation of the cost of the tree regarding the position of the Steiner point. Darker areas represent the position of the Steiner point for lowest-cost trees. The white dot in the figures represent the position of the Steiner point in the case without an obstacle. Note that if the cost of the obstacle is too high then the lowest-cost tree must avoid the obstacle [27].

A tabu search meta-heuristic was developed to tackle the three points problem consisting in a diversification strategy, an intensive local search strategy and a tabu list. The diversification strategy spreads out the initial solutions, that are potential candidates to Steiner points, all over the entire region of interest. The intensive local search strategy focuses the search around a given point if there are points near it that, if taken as Steiner points, give rise to low-cost trees. Therefore, there are good chances to find even better solutions near that point. The tabu list avoids temporarily the search of solutions near some points that apparently are not good candidates to intensify the search around them. The local search and the diversification strategies are described in detail in [27] and the

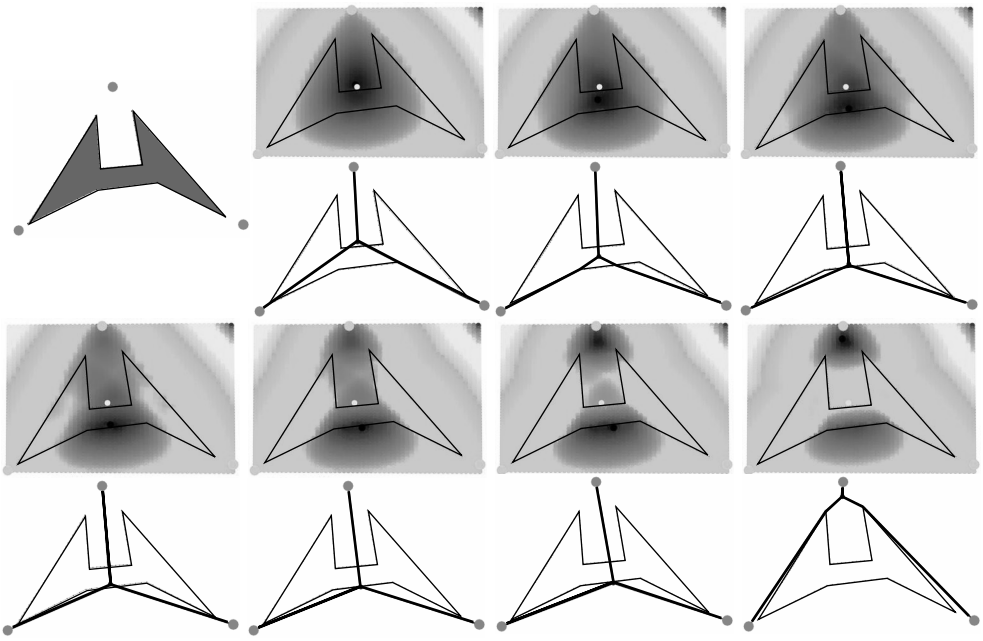

Figure 10: Connecting three terminal nodes with the costs of the obstacle equal to 1, 1.05, 1.1, 1.25, 1.5, 1.75 and 2 (increasing from left to right and from top to bottom)

tabu list was added latter in order to obtain the same solutions in less computation time.

4.4 Euclidean Steiner Tree Heuristic

The problem of computing a minimum-cost Euclidean Steiner tree is a well-known $\mathcal{NP}$-hard problem and the extension with soft obstacles presented herein clearly a generalization. To tackle this problem a heuristic was proposed in [27] based on an initial tree that is a minimum spanning tree (MST) obtained with Prim's algorithm. The inputs of this algorithm are the weighted distances between all pairs of terminal nodes obtained with the shortest path algorithm presented in Sect. 4.2. In general, the Euclidean Steiner heuristic then takes successively three nodes of the MST (advancing one node at a time) and tries to decrease the cost of the final tree by interconnecting these three points through a Steiner point using the strategy presented in [27] for the three points problem.

The results presented next were obtained with an improved version of the previous mentioned heuristic. The main improvements are the initial MST and the strategy used for the three points problem which is based in the tabu search meta-heuristic presented in the previous subsection. The initial MST (MST_{ini}) instead of simply using the shortest paths between each pair of terminal nodes, as in the previous Steiner heuristic, now tries to

use intermediary nodes of the shortest paths (that are not terminal nodes) to decrease the cost of the initial tree. In fact, some of these intermediary nodes might allow to connect more than two terminal nodes with a lower cost than the cost of the shortest paths between them.

4.5 Experimental Results

Since no dataset is available to validate the proposed heuristic considering soft obstacles, the dataset in [27] considering solid obstacles was used because the optimal solutions are known [25]. The results presented in Table 5 show that the strategy followed to obtain the initial tree (MST_{ini}) managed to improve the initial solution in 50% of the cases (compared with the MST). The maximum error percentage for the cost of the final tree presented in Table 5 is 0.7484% which is much lower than 10.17% presented in [27] and so it can be concluded that the new heuristic performs better than the previous one.

Designing a network's spine that connects a set of terminal nodes (e.g., nodes that can represent communication equipment) through a minimum-cost infrastructure, and considering costs associated to disaster-prone areas, requires real world data. In order to represent a near-realistic scenario, the validation of the proposed approach was carried out using polygons (see Fig. 11) extracted from one of the models in the 2013 European Seismic Hazard Model [31]. Additional obstacles with high cost were added to avoid unwanted areas, as is the case of the sea. The terminal nodes of the validation scenario were adapted from the COST266 network [30].

5 Conclusions

In this chapter, we revisited the *spine* idea of embedding a high-availability sub-structure, at the physical layer, complemented with other protection and restoration techniques at upper layers, in order to allow a network provider to offer a larger range of resilience classes in a cost-effective fashion. Here the focus was on providing high availability levels required by critical services in the context of disaster situations. The *spine* concept was shown to make networks more robust to disasters, by enhancing links resistance to disaster-based challenges, as was illustrated in three approaches presented here. FRADIR, a new disaster-resilience framework, showed the *spine* concept can be used to mitigate the impact of regional failures (e.g., disasters). A proposal on how to select edges for availability upgrade to ensure high availability under D-geodiverse constraints was also presented.

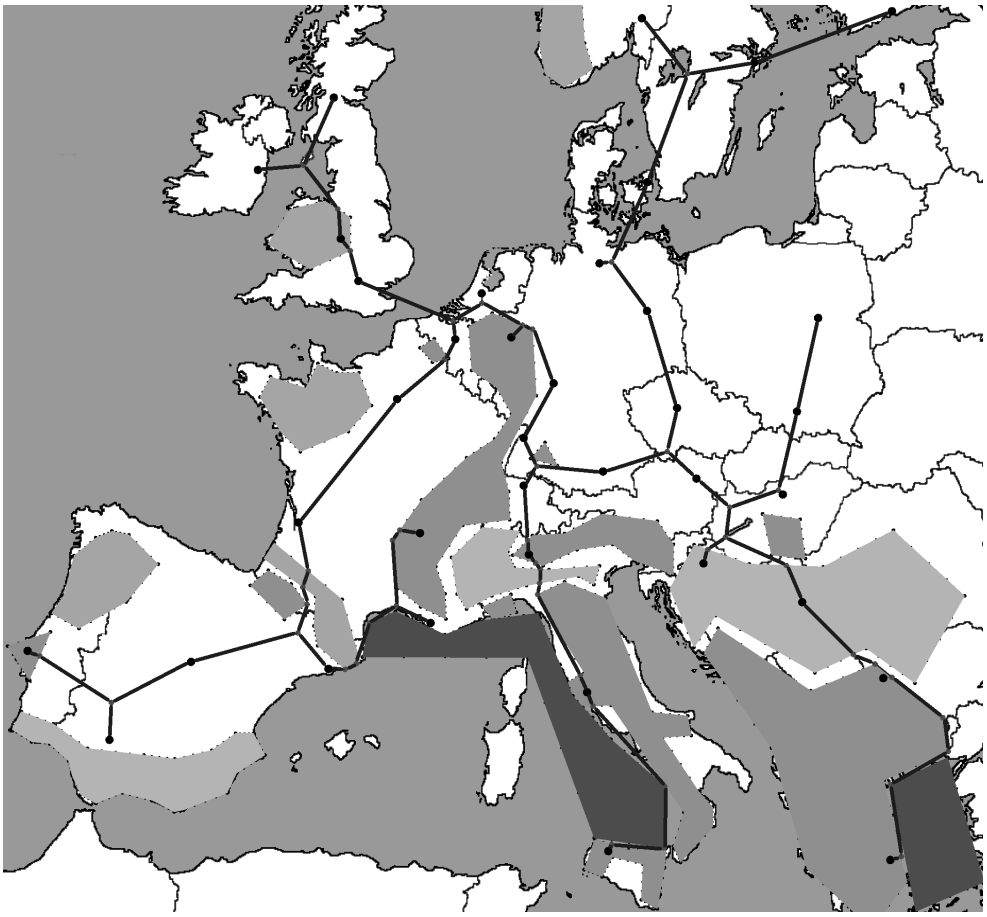

Figure 11: Europe COST 266 network nodes [30]: soft obstacles are derived from the seismic hazard model in [31]; the costs of the obstacles are 1.5, 3, 5 and 10 being the higher costs related to obstacles represented by darker regions

Table 5: Results for the proposed approach applied to the dataset used in [27]

Scenario	1	2	3	4	5	6	7	8	9
Exact	1.23162	5.47342	5.93503	15.3413	11.5136	15.9702	6.00962	15.7772	1.11558
Method	1.23203	5.51438	5.94391	15.3561	11.588	15.9745	6.00997	15.7772	1.11623
Error %	0.0333	0.7484	0.1496	0.0961	0.6462	0.0270	0.0057	0	0.0588
MST	1.3393	5.7883	6.21489	16.0679	12.8935	16.3397	6.94067	17.8215	1.27075
MST_{ini}	1.3393	5.77299	6.21489	15.8179	11.6403	16.3397	6.22076	15.9978	1.20912

Scenario	10	11	12	13	14	15	16	17	18
Exact	12.3587	15.0233	17.093	18.1115	2.18564	2.19569	2.29706	14.1726	17.0214
Method	12.3647	15.0233	17.093	18.1115	2.20129	2.20129	2.29706	14.2489	17.0223
Error %	0.0486	0	0	0	0.7162	0.2553	0	0.5384	0.0052
MST	13.7153	15.1859	17.3573	18.8219	2.4	2.4	2.4	16.0864	23.49
MST_{ini}	13.194	15.1859	17.3573	18.8219	2.4	2.4	2.4	16.0566	21.5887

Finally, a heuristic for a minimum-cost Euclidean Steiner tree representing the backbone network – the *spine* – taking into account the presence of soft obstacles (areas that should be avoided, or crossed at a given cost) was described. Illustrative results for these three approaches were also presented.

References

[1] D. Tipper. Resilient network design: challenges and future directions. *Telecommunication Systems*, 56(1):5–16, 2014.

[2] H.-Y. Chang and P.-C. Wang. Upgrading service availability of optical networks: A labor force perspective. *International Journal of Communication Systems*, page e3553, 2018. First published: 8 March 2018.

[3] I. Rados, T. Sunaric, and P. Turalija. Suggestions for availability improvement of optical cables. In *Circuits and Systems for Communications, 2002. Proceedings. ICCSC '02. 1st IEEE International Conference on*, pages 234–239, 2002.

[4] E. Marcus and H. Stern. *Blueprints for High Availability.* Wiley Publishing, 2003.

[5] A. Pašić, R. Girão-Silva, B. Vass, T. Gomes, and P. Babarczi. FRADIR: A novel framework for disaster resilience. In *10th International Workshop on Resilient Networks Design and Modeling (RNDM 2018)*, Aug. 2018.

[6] J. Tapolcai, B. Vass, Z. Heszberger, J. Biró, D. Hay, F. A. Kuipers, and L. Rónyai. A tractable stochastic model of correlated link failures caused by disasters. In *Proc. IEEE INFOCOM*, Honolulu, USA, April 2018.

[7] P. Babarczi, A Pasic, J. Tapolcai, F. Németh, and B. Ladóczki. Instantaneous recovery of unicast connections in transport networks: Routing versus coding. *Elsevier Computer Networks*, 82:68–80, 2015.

[8] R. Girão-Silva, L. Martins, T. Gomes, A. Alashaikh, and D. Tipper. Improving network availability—a design perspective. In X.-S. Yang, S. Sherratt, N. Dey, and A. Joshi, editors, *Third International Congress on Information and Communication Technology*, Advances in Intelligent Systems and Computing, pages 799–815, Singapore, 2019. Springer Singapore.

[9] J.-P. Vasseur, M. Pickavet, and P. Demeester. *Network Recovery – Protection and Restoration of Optical, SONET-SDH, IP, and MPLS*. Elsevier, 2004.

[10] US Network. `http://lendulet.tmit.bme.hu/~pasic/networks/`. Accessed: 2019-04-08.

[11] S. Orlowski, R. Wessäly, M. Pióro, and A. Tomaszewski. SNDlib 1.0–Survivable Network Design library. *Networks*, 55(3):276–286, 2010. http://sndlib.zib.de.

[12] A. Pašić, R. Girão-Silva, B. Vass, T. Gomes, F. Mogyorósi, P. Babarczi, and J. Tapolcai. FRADIR-II: An improved framework for disaster resilience. In *11th International Workshop on Resilient Networks Design and Modeling (RNDM 2019)*, Oct. 2019.

[13] A. Alashaikh, D. Tipper, and T. Gomes. Designing a high availability subnetwork to support availability differentiation. In *14th International Conference on the Design of Reliable Communication Networks (DRCN 2018)*, Paris, France, 19 Feb. 2018.

[14] J. Rohrer, A. Jabbar, and J. P. G. Sterbenz. Path diversification: a multipath resiliency mechanism. In *7th International Workshop on the Design of Reliable Communication Networks – DRCN 2009*, pages 343–351, Washington, DC, USA,, October 25-28 2009.

[15] Y. Cheng, M. T. Gardner, J. Li, R. May, D. Medhi, and J. P. G. Sterbenz. Analysing geopath diversity and improving routing performance in optical networks. *Computer Networks*, 82:50–67, 2015.

[16] A. de Sousa, D. Santos, and P. Monteiro. Determination of the minimum cost pair of D-geodiverse paths. In *The 2017 International Conference on Design of Reliable Communication Networks (DRCN 2017)*, Munich, March 8-10 2017.

[17] T. Gomes, J. Tapolcai, C. Esposito, D. Hutchison, F. Kuipers, J. Rak, A. de Sousa, A. Iossifides, R. Travanca, J. André, L. Jorge, L. Martins, P. O. Ugalde, A. Pašić, D. Pezaros, S. Jouet, S. Secci, and M. Tornatore. A survey of strategies for communication networks to protect against large-scale natural disasters. In *8th International Workshop on Resilient Networks Design and Modeling (RNDM)*, pages 11–22, Sept 2016.

[18] A. de Sousa, T. Gomes, R. Girão-Silva, and L. Martins. Minimizing the network availability upgrade cost with geodiversity guarantees. In *2017 9th International Workshop on Resilient Networks Design and Modeling (RNDM)*, pages 1–8, Sept. 4-6 2017.

[19] A. de Sousa, T. Gomes, R. Girão-Silva, and L. Martins. Minimization of the network availability upgrade cost with geodiverse routing for disaster resilience. *Optical Switching and Networking*, 31:127–143, Jan. 2019.

[20] `http://www.av.it.pt/asou/geodiverse.htm`, March 2017.

[21] A. Alashaikh, T. Gomes, and D. Tipper. The spine concept for improving network availability. *Computer Networks*, 82(0):4–19, 2015.

[22] T. Nagae and H. Wakabayashi. Differences in network reliability improvement by several importance indices. *Transportation Research Procedia*, 10:155–165, 2015. 18th Euro Working Group on Transportation, EWGT 2015, 14-16 July 2015, Delft, The Netherlands.

[23] U. Franke. Optimal it service availability: Shorter outages, or fewer? *IEEE Transactions on Network and Service Management*, 9(1):22–33, March 2012.

[24] M. Brazil and M. Zachariasen. *Optimal Interconnection Trees in the Plane: Theory, Algorithms and Applications*. Algorithms and Combinatorics. Springer, 2015.

[25] M. Zachariasen and P. Winter. *Obstacle-Avoiding Euclidean Steiner Trees in the Plane: An Exact Algorithm*, pages 286–299. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999.

[26] M. Zachariasen. A catalog of hanan grid problems. *Networks*, 38(2):76–83, 2001.

[27] L. Garrote, L. Martins, U. Nunes, and M. Zachariasen. Weighted euclidean steiner trees for disaster-aware network design. In *DRCN 2019 - Design of Reliable Communication Networks; 15th International Conference*, Coimbra, Portugal, March 2019.

[28] M. Lanthier, A. Maheshwari, and J.-R. Sack. Approximating weighted shortest paths on polyhedral surfaces. In *Proceedings of the Thirteenth Annual Symposium on Computational Geometry*, SCG '97, pages 485–486, New York, NY, USA, 1997. ACM.

[29] J. E. Bresenham. Algorithm for computer control of a digital plotter. *IBM Systems Journal*, 4(1):25–30, 1965.

[30] S. Orlowski, R. Wessäly, M. Pióro, and A. Tomaszewski. SNDlib 1.0 – Survivable Network Design library. *Networks*, 55(3):276–286, 2010. http://sndlib.zib.de.

[31] J. Woessner, D. Laurentiu, D. Giardini, H. Crowley, F. Cotton, G. Grünthal, G. Valensise, R. Arvidsson, R. Basili, M. B. Demircioglu, S. Hiemer, C. Meletti, R. W. Musson, A. N. Rovida, K. Sesetyan, and M. Stucchi. The 2013 european seismic hazard model: key components and results. *Bulletin of Earthquake Engineering*, 13(12):3553–3596, Dec 2015.