Minimizing the network availability upgrade cost with geodiversity guarantees

Amaro de Sousa
Instituto de Telecomunicações, 3810-193 Aveiro, Portugal
Departamento de Eletrónica, Telecomunicações e Informática,
Universidade de Aveiro, 3810-193 Aveiro, Portugal
Email: asou@ua.pt

Teresa Gomes, Rita Girão-Silva and Lúcia Martins
Department of Electrical and Computer Engineering,
University of Coimbra, 3030-290 Coimbra, Portugal
INESC Coimbra, 3030-290 Coimbra, Portugal
Email: {teresa,rita,lucia}@deec.uc.pt

***Abstract*—As telecommunication networks are a critical infrastructure of our society, they must evolve to provide high end-to-end availability and high resilience to large-scale disasters. Path protection mechanisms can improve end-to-end availability but, in general, might not be enough to reach the availability required by critical services. Moreover, adding geodiversity to the routing paths (i.e., selecting path pairs with higher geographical distance between them) enhances the network disaster resilience but also makes it more challenging to reach a high end-to-end availability as the resulting paths tend to be longer. So, for a network where each link is characterized by its current availability and by the cost of upgrading its availability to a new value, this paper proposes some strategies aiming to determine a set of links to be upgraded at a minimum cost ensuring a desired level of availability and geodiversity. The problem is defined as an integer non-linear programming model, a solving algorithm based on different greedy strategies is proposed and the relative performance of the different strategies is evaluated on a set of problem instances.**

***Index Terms*—availability; path geodiversity; resilience; disaster.**

I. Introduction

It is well known that telecommunication networks are currently one of the critical infrastructures upon which our society depends and services are expected to be always-on. Moreover, disaster-based failures are becoming more frequent in time and wider in scope, degrading drastically the services supported by telecommunication networks [1] (a survey of strategies to protect networks against large-scale natural disasters is presented in [2]). Thus, it is imperative that telecommunication networks become both highly available and resilient to disasters.

Networks need to guarantee that all node pairs of interest (i.e., node pairs involved in critical services such as emergency calls, smart grid communications, financial transactions, etc.) have a high end-to-end availability [3], [4], [5]. Moreover, when a disaster-based failure occurs, it is important not only to quickly recover the network in the disaster area but also to minimize the disaster impact for node pairs outside the disaster area. Here, we explore the idea of path geodiversity, i.e., how to take into consideration the geographical diversity of the network topology when making routing decisions. Path geodiversity has been used to improve services availability [6] and disaster resilience [7].

Consider a network topology such that the geographical distance between any two network elements (nodes or links) is known. Like in [8], [9], we consider a path geodiversity strategy where the routing between two network nodes is based on two paths geographically separated by at least a minimum distance D, so that, if a disaster with a geographical coverage whose diameter is lower than D affecting one intermediate element of one path cannot affect the other path. Here, we follow [9] which considers intermediate elements both as nodes and links, while [8] considers only intermediate nodes. Note that the larger the value D is, the more resilient the network is to disasters but the more difficult it becomes to reach high end-to-end availability since the resulting paths tend to be longer.

In this work, we consider jointly the end-to-end availability and the disaster resilience of a given network topology. We address the network availability upgrade problem where the availability of some links must be upgraded to provide the required end-to-end availability and disaster resilience for a set of node pairs of interest. In [10], the authors also select a set of links to be upgraded (shielded, so that they become invulnerable). However only end-to-end connectivity against geographic based attacks is ensured, without considering availability requirements.

The upgraded network must guarantee that each node pair of interest is provided with at least one path pair fulfilling two requirements. Concerning end-to-end availability, at least one path pair with a minimum target availability of Λ must be provided to each node pair. As pointed out in [9], the maximum geodiversity value that can be provided by a path pair to a given node pair is constrained by the geographical locations of the nodes and the geographical paths of the links. Hence, concerning disaster resilience, a target geodiversity value D with a soft requirement is used: at least one path pair with a minimum geodiversity of D must be provided to each node pair whose maximum geodiversity value is at least D; for node pairs such that the network topology cannot provide D, the maximum possible geodiversity value is used.

The paper is organized as follows. In section II, the network availability upgrade problem is described. In section III, a solving algorithm based on different greedy strategies is proposed. Section IV presents the computational results comparing the efficiency of the different greedy strategies. Finally, section V presents the main conclusions, along with some further work.

II. Problem Definition

Consider two given parameters: a minimum availability parameter Λ and a minimum geodiversity parameter D. Consider a given biconnected network defined by an undirected graph $G = (N, E)$ where N is the set of nodes and E is the set of edges representing node pairs connected by a direct link. Each edge $e \in E$ is characterized by its current availability a_e, its upgraded availability a_e^u and the cost c_e required to upgrade its availability from a_e to a_e^u. The aim is to determine a set of edges to be upgraded at a minimum cost. For a given set of node pairs of interest K, the upgraded availability solution must guarantee the existence of at least one pair of paths for each node pair $(s,t) \in K$ with (i) a minimum availability of Λ and (ii) a minimum geodiversity of D, if the topology of G allows it, or the maximum possible geodiversity value if it is lower than D.

For each node pair $(s,t) \in K$, consider the set of all pairs of paths available in G defined by R_{st} where each path pair $r \in R_{st}$ is defined by two sets of edges: S_{r1} (the edge set of the first path of r) and S_{r2} (the edge set of the second path of r). Then, consider the following binary variables:

x_e is equal to 1 if edge $e \in E$ is upgraded; 0 otherwise.
y_{str} is equal to 1 if node pair $(s,t) \in K$ may be provided with path pair $r \in R_{st}$; 0 otherwise.

Using these variables, the availability of a path pair $r \in R_{st}$, represented by Λ_r, is given by:

$$\begin{aligned}\Lambda_r \;=\; 1 - &\left(1 \;-\; \prod_{e \in S_{r1}} \left((1-x_e)\,a_e + x_e a_e^u\right)\right) \times \\ &\left(1 \;-\; \prod_{e \in S_{r2}} \left((1-x_e)\,a_e + x_e a_e^u\right)\right) \qquad (1)\end{aligned}$$

Following [9], the geodiversity of a path pair is the minimum distance between any intermediate node or edge of one path and any node or edge of the other path. In [9], it is shown that the geodiversity of a path pair can be modelled based only on geographical distances between edges provided that these distances are defined as follows. For each $r \in R_{st}$ between s and t, the geographical distance between one edge $e_i \in S_{r1}$ and one edge $e_j \in S_{r2}$, represented by $\delta(e_i, e_j)$, is defined as (i) the minimum distance between any point in the geographical path of e_i and any point in the geographical path of e_j if they do not share s or t or (ii) the minimum distance between one edge and the non-common end node of the other edge if they share either s or t. Fig. 1 shows part of a network with the source node s, five other nodes (2 to 6) and five edges (a to e) illustrating the geographical distances between edges in the different cases. Examples of case (i) are the distances $\delta(a,e)$ and $\delta(b,e)$. Note that the zero distances $\delta(a,b)$, $\delta(c,d)$ and $\delta(d,e)$ also illustrate this case, because each pair of edges shares a node, which is neither s nor t. As for case (ii), the distance $\delta(a,c)$ between edges a and c, since they share the source node s, is the minimum between the distance between node 2 (the non-common end node of a) and edge c and the distance between node 3 (the non-common end node of c) and edge a. In this work, it will be considered that links follow the shortest path over a sphere that represents Earth.

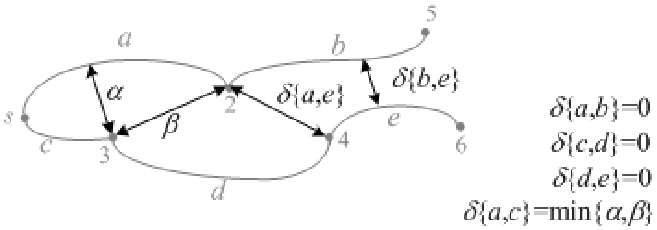

Fig. 1: Geographical distances between edges (adapted from [9])

Then, the geodiversity value of $r \in R_{st}$, represented by D_r, is given by:

$$D_r = \min_{e_i \in S_{r1}, e_j \in S_{r2}} \delta(e_i, e_j) \qquad (2)$$

Note that the geodiversity value of a path pair only depends on the geographical path of each edge. Moreover, for each node pair $(s,t) \in K$, there is a maximum geographical distance, which we represent by D_{st}^{Max}, above which a pair of paths is infeasible (these values can be computed in advance using [9]). Therefore, we consider the geodiversity requirement using $D_{st} = \min(D, D_{st}^{Max})$, i.e., we require D if $D_{st}^{Max} \geq D$ or we require D_{st}^{Max}, if $D_{st}^{Max} < D$.

The network upgrade problem is, then, defined by the following integer non-linear programming model:

$$\text{Minimize} \quad \sum_{e \in E} c_e x_e \qquad (3)$$

Subject to:

$$\begin{aligned} \sum_{r \in R_{st}} y_{str} &\geq 1 && (s,t) \in K && (4)\\ \Lambda_r &\geq \Lambda\, y_{str} && (s,t) \in K, r \in R_{st} && (5)\\ D_r &\geq D_{st}\, y_{str} && (s,t) \in K, r \in R_{st} && (6)\\ x_e &\in \{0,1\} && e \in E && (7)\\ y_{str} &\in \{0,1\} && (s,t) \in K, r \in R_{st} && (8)\end{aligned}$$

The objective function (3) is the minimization of the total cost of all upgraded edges. Constraints (4) guarantee that each $(s,t) \in K$ is provided with at least one path pair (i.e., the paths $r \in R_{st}$ such that the variable y_{str} is set to 1) and each of these path pairs has an availability value not lower than Λ, guaranteed by constraints (5), and a geodiversity value not lower than D_{st}, guaranteed by constraints (6). Finally, constraints (7)-(8) are the variable domain constraints.

Note that, in constraints (5), Λ_r is given by expression (1). These constraints relate variables x_e with variables y_{str} in a non-linear way which turns the proposed formulation in an integer non-linear programming model. In this case, standard solution techniques are not valid and appropriate exact methods (i.e., able to compute the optimal solutions) must be investigated.

III. Solving Algorithm

The solving algorithm proposed here, named Minimum Upgrade Cost with Availability and Geodiversity (MUCAG), uses an iterative approach based on a greedy strategy. Starting with the network configuration without any upgraded edge, the algorithm selects iteratively one edge to be upgraded until the resulting network configuration provides at least a pair of paths for each node pair $(s,t) \in K$ with a minimum availability of Λ and a minimum geodiversity of D_{st}.

We designate a pair of paths with a minimum geodiversity D_{st} as a pair of geodiverse paths. As already explained, a pair of geodiverse paths is always feasible for all $(s,t) \in K$, as $D_{st} = \min(D, D_{st}^{Max})$. Nevertheless, the existence of a pair of geodiverse paths with a minimum availability Λ depends on the set of upgraded edges. So, a nuclear task of MUCAG is, for a given network configuration and a given $(s,t) \in K$, to compute a pair of geodiverse paths r whose availability Λ_r is at least Λ. This task is implemented through an algorithm, named Guaranteed Available Pair of Geodiverse Paths (GAPGP), which is an adaptation of the algorithm in [11] for calculating the most reliable pair of link disjoint paths. For reasons that will be explained in the detailed description of MUCAG, algorithm GAPGP also computes the pair of geodiverse paths with the highest availability value if such value is lower than the required Λ.

In the following three subsections, we describe separately first the GAPGP algorithm, then, the MUCAG algorithm and, finally, the different greedy strategies tested in practice.

A. GAPGP algorithm description

GAPGP is specified in Algorithm 1. For a given node pair $(s,t) \in K$, minimum availability value Λ, minimum geodiversity value D_{st} and network configuration (defined by the values assigned to the binary variables x_e), GAPGP computes a pair of geodiverse paths $r \in R_{st}$ with an availability value Λ_r which is either not lower than Λ, if such path pair exists, or is maximal if $\max_{r \in R_{st}} \Lambda_r \leq \Lambda$.

GAPGP starts by computing an edge cost c'_e for all $e \in E$ (lines 1–3) such that enumerating the k shortest paths using these costs corresponds to the enumeration of the k paths with the highest availability.

Then, in the main while cycle (lines 6–27), the algorithm iteratively generates a new first path p with function *next-shortest-path* (line 7) by non-increasing order of availability value (*next-shortest-path* corresponds to the iterative use of Yen's [12] k-shortest path algorithm or of the loopless version [13] of the MPS algorithm [14]). For each first path p, a second path q is computed by function *path-geo-distance* (line 16) as the path with the highest availability and a geodiversity of D_{st} with p (function *path-geo-distance* runs a shortest path algorithm in an auxiliary graph given by G without the edges of p and the edges with a distance from any edge of p below D_{st}). If the second path q exists ($q \neq \emptyset$ in line 17), the availability of path pair $r' = (p,q)$ is evaluated to check if the current best solution r must be updated (lines 18–21).

The algorithm stops (line 6) when either the availability of the current best path pair r is at least Λ or if variable *opt* becomes true (which means that the path pair with the highest availability has been reached). Variable *opt* becomes true in one of two cases. The first case (line 25) is when function *next-shortest-path* (line 7) returns no path (i.e. $p = \emptyset$ in line 8), which means that all possible paths have already been enumerated. The second case (line 12) is when the availability of the current best path pair cannot be further improved (line 11) (condition also used in [11]).

To understand the condition of line 11, let $Av(p)$ represent the availability of a path p, i.e., $Av(p) = e^{-c'(p)}$, where $c'(p) = \sum_{e \in p} c'_e$. Consequently, the availability of a path pair (p,q) is $Av(p) + (1 - Av(p))Av(q)$. Let the current best path pair be $r = (p_w, q_w)$ with availability Λ_r, where w is the order of generation of the first element of the pair (path p obtained in line 7) and q_w is the second path (path q obtained in line 16). Finally, let the new first path generated by *next-shortest-path* be p_i (with $i > w$). If $Av(p_i) + (1 - Av(p_i))Av(p_i) \leq \Lambda_r$ (line 11), then $r = (p_w, q_w)$ is optimal. The verification of this statement is straightforward. Following [11], notice first that $Av(p_i) \leq Av(p_w)$. Let q_i be the path with the highest availability which guarantees a geodiversity of D_{st} with p_i. If $Av(q_i) > Av(p_i)$, this path pair would have already been obtained when $p = q_i$. On the other hand, if $Av(q_i) \leq Av(p_i)$, this path pair has an availability which is at most $Av(p_i) + (1 - Av(p_i)))Av(p_i)$ and, by the above condition, lower than Λ_r. So, any path pair obtained from this point onwards has an availability not better than Λ_r.

B. MUCAG algorithm description

MUCAG is specified in Algorithm 2. At the beginning (lines 1-3), the algorithm sets all variables x_e to 0, representing the network with no upgraded edges.Then, the algorithm runs a while cycle (lines 4-19). At the beginning of each cycle, set K has the node pairs (s,t) for which there is still no pair of geodiverse paths with the required availability Λ. The while cycle runs until set K is empty (line 4).

In the first step of each cycle (line 5), the auxiliary sets K' and $\mathcal{R}$ are first initialized empty (note that at the end of each cycle, $\mathcal{R}$ includes all the best pairs of geodiverse paths, which still do not attain Λ). Then, for each $(s,t) \in K$ (lines 6-12), algorithm GAPGP (line 7) computes a pair of geodiverse paths r with an availability Λ_r (recall the description in the previous subsection) that, if below the required value Λ (line 8), makes node pair (s,t) to be added to K' (line 9) and path pair r to be added to $\mathcal{R}$ (line 10). Note that when $\Lambda_r < \Lambda$, r is the path pair with the highest availability of the current network configuration for node pair (s,t). So, the edges involved in such path pairs are likely to be the most promising ones to be upgraded to reach the required network configuration in subsequent cycles.

Algorithm 1 GAPGP

```
Require: $G$, $s$, $t$, $\Lambda$, $D_{st}$, $(a_e, a_e^u, x_e) : \forall e \in E$, $\delta(e_i, e_j) : \forall e_i, e_j \in E$
Ensure: $(r, \Lambda_r)$
for all $e \in E$ do
    $c'_e = -\log(a_e)(1 - x_e) - \log(a_e^u)x_e$ ▷ New edge cost
end for
$r \leftarrow (\emptyset, \emptyset)$; $\Lambda_r \leftarrow 0$
$opt \leftarrow false$ ▷ $r$ is not yet the most available path pair
while $\Lambda > \Lambda_r \wedge \neg opt$ do
    $p \leftarrow$ next-shortest-path$(s, t, G, c')$
    if $p \neq \emptyset$ then
        if $\Lambda_r \neq 0$ then
            $g \leftarrow e^{-c'(p)}$
            if $g(2 - g) \leq \Lambda_r$ then ▷ $Av(p) \triangleq Av(q)$
                $opt \leftarrow true$ ▷ $r$ is the optimal solution
            end if
        end if
        if $\neg opt$ then
            $q \leftarrow$ path-geo-distance$(\delta, D_{st}, p, G, c')$
            if $q \neq \emptyset$ then
                $r' \leftarrow (p, q)$
                if $\Lambda_{r'} > \Lambda_r$ then
                    $r \leftarrow r'$ ▷ Updates solution
                end if
            end if
        end if
    else
        $opt \leftarrow true$ ▷ No more improvement possible
    end if
end while
```

Algorithm 2 MUCAG

```
Require: $G$, $K$, $\Lambda$, $(a_e, a_e^u, c_e) : \forall e \in E$, $D_{st} : \forall (s, t) \in K$
Ensure: $x_e : \forall e \in E$
 1: for all $e \in E$ do
 2:     $x_e \leftarrow 0$
 3: end for
 4: while $K \neq \emptyset$ do
 5:     $K' \leftarrow \emptyset, \mathcal{R} \leftarrow \emptyset$
 6:     for all $(s, t) \in K$ do
 7:         $(r, \Lambda_r) \longleftarrow$ GAPGP$(s, t, D_{st}, \Lambda, x_e : e \in E)$
 8:         if $\Lambda_r < \Lambda$ then
 9:             $K' \leftarrow K' \cup \{(s, t)\}$
10:             $\mathcal{R} \leftarrow \mathcal{R} \cup \{r\}$
11:         end if
12:     end for
13:     $K \leftarrow K'$
14:     if $|K| > 0$ then
15:         $(e, K') \leftarrow$ selectEdge$(\mathcal{R}, K, x_e : e \in E)$
16:         $x_e \leftarrow 1$
17:         $K \leftarrow K \setminus K'$
18:     end if
19: end while
```

Then, K is set with K' (line 13), i.e., the new set K has only the node pairs (s, t) for which there is still no pair of geodiverse paths compliant with the availability requirement. Finally, if set K is not empty (lines 14-18): (i) function *selectEdge* (line 15) selects one edge e (among the non-upgraded ones belonging to path pairs in $\mathcal{R}$) and calculates the end nodes K' of the path pairs whose availability becomes at least Λ with the selected edge (this set can be an empty set), (ii) the selected edge e is upgraded (line 16) and (iii) K' is removed from K (line 17).

Note that if all the edges in set $\mathcal{R}$ (see line 15) have already been upgraded, the algorithm ends without achieving the desired end-to-end availability for the demands currently in set K. For simplicity, this situation is not considered in the algorithm description.

C. Variants for the greedy algorithm

In line 15 of Algorithm 2, the edge to be upgraded next may be selected according to different greedy strategies. Consider first that $E(\mathcal{R})$ represents the set of edges, among the ones still not upgraded, that belong to at least one of the path pairs in $\mathcal{R}$. The edge to be upgraded next is one edge of $E(\mathcal{R})$. The different strategies described next use one (or a combination) of the following subsets of $E(\mathcal{R})$:

$E_M(\mathcal{R})$ is defined as follows: first, a counter is associated to each edge in $E(\mathcal{R})$ with the number of times it is in the path pairs of $\mathcal{R}$; then, $E_M(\mathcal{R})$ is composed by the edges with maximum counter value (i.e. the edges whose upgrade improves the availability of more path pairs).

$E_O(\mathcal{R})$ is defined as follows: first, a counter is associated to each edge in $E(\mathcal{R})$ with the number of path pairs of $\mathcal{R}$ whose availability becomes at least Λ if the edge is upgraded; then, $E_O(\mathcal{R})$ is composed by the edges with maximal counter value or is empty if the maximal counter value is 0 (i.e. when not empty, it contains the edges whose upgrade results in the highest number of path pairs fulfilling the required availability).

In the following, we describe different greedy strategies, creating different variants for the greedy algorithm:

- **Min-Cost:** Select the minimum upgrade cost edge in $E(\mathcal{R})$.
- **Min-Cost–Max-Count:** Select the minimum upgrade cost edge in $E_M(\mathcal{R})$.
- **Min-Cost–Max-On:** Select the minimum upgrade cost edge in $E_O(\mathcal{R})$ if $E_O(\mathcal{R})$ is not empty, or in $E_M(\mathcal{R})$ if $E_O(\mathcal{R})$ is empty.
- **Max-On–Max-Count:** Select the edge in $E_M(\mathcal{R})$ that, if upgraded, maximizes the number of path pairs whose availability becomes at least Λ. If no such edge exists, select any edge in $E_M(\mathcal{R})$.
- **Max-Count–Max-On:** Select among the edges in

$E_O(\mathcal{R})$ the one whose upgrade improves the availability of more path pairs, or any edge in $E_M(\mathcal{R})$ if $E_O(\mathcal{R})$ is empty.

IV. Computational Results

The computational results presented in this section consider two network topologies (also used in [9]) representative of typical telecommunication transport networks: the Germany50 topology (geographical location of nodes available at [15]) and the CORONET CONUS topology (geographical location of nodes available at http://www.monarchna.com/topology.html), referred to as Coronet in this section. Since there is no available information on the geographical path of each link, we have considered that links follow the shortest path over the terrestrial surface assuming that Earth is a sphere, as already mentioned in section II. For both topologies, the information concerning edge lengths and geographical distance between edge-edge pairs and node-edge pairs is available at http://www.av.it.pt/asou/geodiverse.htm.

In both topologies, we have considered the worst case of set K composed by all node pairs. Then, for each $(s,t) \in K$, the maximum geodiversity value D_{st}^{Max} was obtained solving the Maximum Distance D of Geodiverse Paths (MDDGP) optimization problem proposed in [9]. Table I shows the topology characteristics of both networks.

Concerning the edge availability values, we have considered the current availability a_e of $e \in E$ based on its length [16]:

$$a_e = 1 - \frac{MTTR}{MTBF_e} \tag{9}$$

with

$$MTBF_e[\text{hrs}] = \frac{CC \times 365 \times 24}{\text{cable length}_e} \tag{10}$$

where $MTBF$ and $MTTR$ are the mean time between failures and mean time to repair in hours, respectively (as in in [4], we consider $MTTR = 24$ and $CC = 450$). CC is the cable cut metric in km (cable lengths also in km). Moreover, we have 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., $a_e^u = a_e\,(2 - a_e)$.

Concerning the upgrade cost values, note first that in the general case, the upgrade cost c_e of each edge $e \in E$ should be composed by a fixed cost and a cost per unit length of the edge. Since the expression to calculate the actual cost is uncertain (it depends on many different factors), here, we analyze the results of the different greedy strategies for the two extreme cases: (1) the upgrade cost is the number of upgraded edges (equivalent to consider a cost of 1 per upgraded edge) and (2) the upgrade cost of each edge is given by its length (equivalent to consider a cost of 1 per length unit of each edge).

TABLE I: Network characteristics ($|N|$, $|E|$, d – average node degree, max L – maximum edge length, avg L – average edge length, $\max_{st} D_{st}^{Max}$) – All lengths are in km.

Network	$\lvert N\rvert$	$\lvert E\rvert$	d	$\max L$	avg L	$\max_{st} D_{st}^{Max}$
Germany50	50	88	3.52	252	100.67	166
Coronet	75	99	2.64	1017	329.72	707

In the following two subsections, we present and discuss separately the computational results obtained to each network topology.

A. Results for the Germany50 network

Due to its geographical coverage and the considered edge availability values, Germany50 already provides four nines (0.9999) availability between all node pairs even in the case of requiring a geodiversity value D_{st}^{Max} for all node pairs $(s,t) \in K$. So, we have considered a minimum availability $\Lambda = 0.99999$. Table II presents the number of upgraded edges, the upgrade cost (length based) and the CPU time for the geodiversity values D of 40 km, 80 km, 120 km and 160 km (solutions providing the minimum number of upgraded edges and/or the minimum upgrade cost highlighted in bold).

Table II shows that, in terms of number of upgraded edges, the Min-Cost–Max-On has the best results, on average, but both Max-On–Max-Count and Max-Count–Max-On are only slightly worse. On the other hand, in terms of upgrade cost, Max-On–Max-Count is clearly the best strategy. So, the main conclusion is that Max-On–Max-Count is the best compromise strategy for Germany50, as it finds solutions with the lowest costs and a number of upgraded edges at most 16% above the minimum values found by any strategy.

Moreover, the simplest Min-Cost strategy is, by far, the worst strategy both in terms of number of upgraded edges and upgrade cost, showing that considering the number of path pairs whose availability is improved by the selected edge (as exploited by the other strategies) leads to more efficient algorithms. The Min-Cost strategy was considered because it is strongly associated with the objective function in equation (3). As expected, it resulted in a large number of short edges being selected for upgrade.

Finally, two expected observations are that, in the overall: (1) higher geodiversity values D impose upgrade solutions with both higher number of upgraded edges and cost (with only a few exceptions); (2) the CPU times are also higher since the number of iterations of MUCAG (see Algorithm 2) grows with the number of upgraded edges.

Fig. 2 presents the solutions found by the best compromise Max-On–Max-Count strategy for the four considered values of D. Besides illustrating that higher values of D impose a higher number of required upgraded edges, it also shows that the set of upgraded edges required by a given D is not a subset of the set of upgraded edges required by a geodiversity value higher than D. So, in practice, the right choice of parameter D is of paramount importance since if later on the required value D becomes larger, the previous upgraded edges might not be the best choices.

Fig. 3 presents the solutions found by the other four strategies for D = 80 km (the solution of Max-On–Max-Count is in Fig. 2b). The Min-Cost solution (Fig. 3a) illustrates the inefficiency of this strategy since it gives preference to many shorter upgraded edges, some of them with a minor

TABLE II: No. of upgraded edges, cost and CPU time for achieving $\Lambda = 0.99999$ between all node pairs in Germany50

	Min-Cost			Min-Cost–Max-Count			Min-Cost–Max-On			Max-On–Max-Count			Max-Count–Max-On		
D	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)
40	40	2561	77.9	18	1507	13.5	15	1842	12.1	16	**1469**	12.7	**14**	1688	11.1
80	52	3684	3731.5	25	2289	139.8	**20**	2664	153.5	22	**2260**	137.9	21	2563	155.2
120	52	3684	5261.0	32	2843	447.0	**25**	3355	324.5	28	**2839**	428.3	27	3389	358.8
160	55	4046	6730.5	37	3285	648.0	26	3405	353.4	29	**2914**	573.0	**25**	3190	358.3

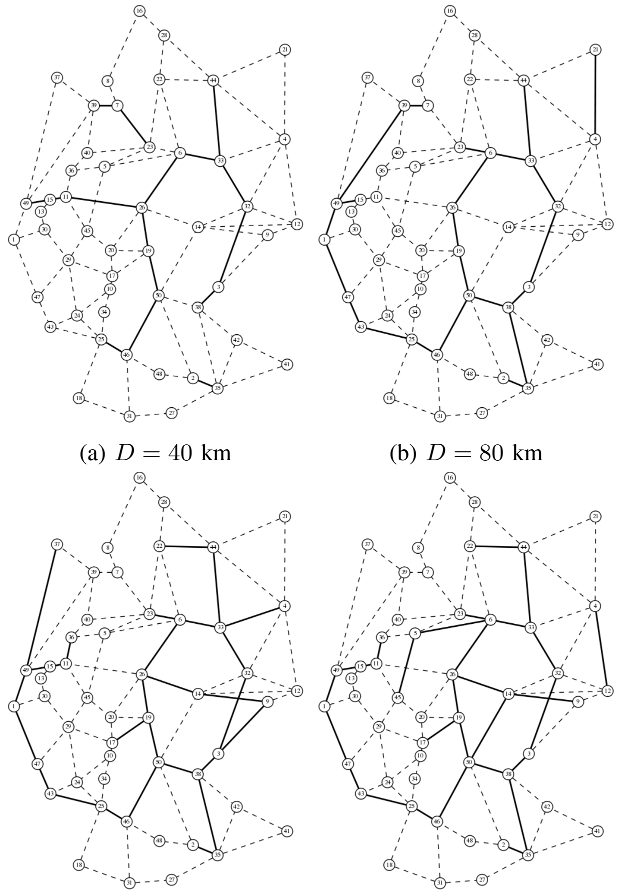
(a) $D = 40$ km
(b) $D = 80$ km
(c) $D = 120$ km
(d) $D = 160$ km

Fig. 2: Max-On–Max-Count: upgraded edges for achieving availability $\Lambda = 0.99999$ between all node pairs in Germany50

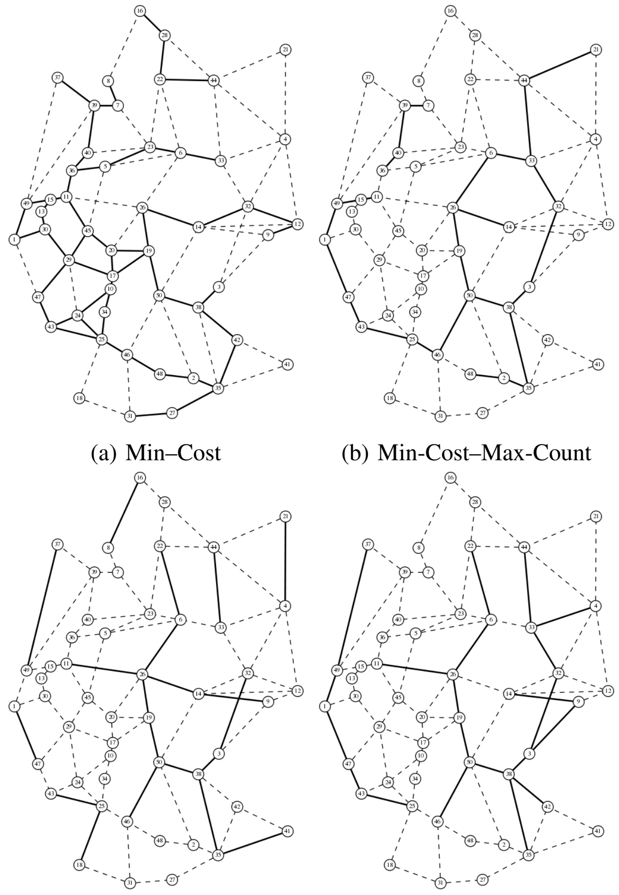
(a) Min–Cost
(b) Min-Cost–Max-Count
(c) Min-Cost–Max-On
(d) Max-Count–Max-On

Fig. 3: Upgraded edges for achieving availability $\Lambda = 0.99999$ between all node pairs in Germany50, when $D = 80$ km

contribution to the availability improvement of the different path pairs, which ultimately leads to a higher upgrade cost.

The solutions of the other four strategies highlight two main characteristics also observable in the results of Table II. First, strategies that consider the edges in $E_M(\mathcal{R})$ (i.e., the most frequent edges in the path pairs of $\mathcal{R}$), namely, the Max-On–Max-Count (Fig. 2b) and the Min-Cost–Max-Count (Fig. 3b), find solutions with lower upgrade costs. Second, strategies that consider the edges in $E_O(\mathcal{R})$ (i.e., edges whose upgrade makes the availability of more path pairs of $\mathcal{R}$ to become at least Λ), namely, the Min-Cost–Max-On (Fig. 3c) and the Max-Count–Max-On (Fig. 3d), find solutions with a lower number of short edges, leading to solutions with a lower number of upgraded edges.

B. Results for the Coronet network

The Coronet network has a much wider geographical coverage and cannot provide four nines (0.9999) availability to many node pairs without upgraded edges. So, in this case, we have considered two minimum availability values. Table III (for $\Lambda = 0.9999$) and Table IV (for $\Lambda = 0.99999$) present the number of upgraded edges, the upgrade cost and the CPU time for the geodiversity values D of 100 km, 200 km, 400 km and 600 km (as before, solutions providing the minimum number of upgraded edges and/or the minimum upgrade cost highlighted in bold).

Table III shows that for $\Lambda = 0.9999$ the Max-On–Max-Count and Max-Count–Max-On strategies are the best compromise between the number of upgraded edges and the

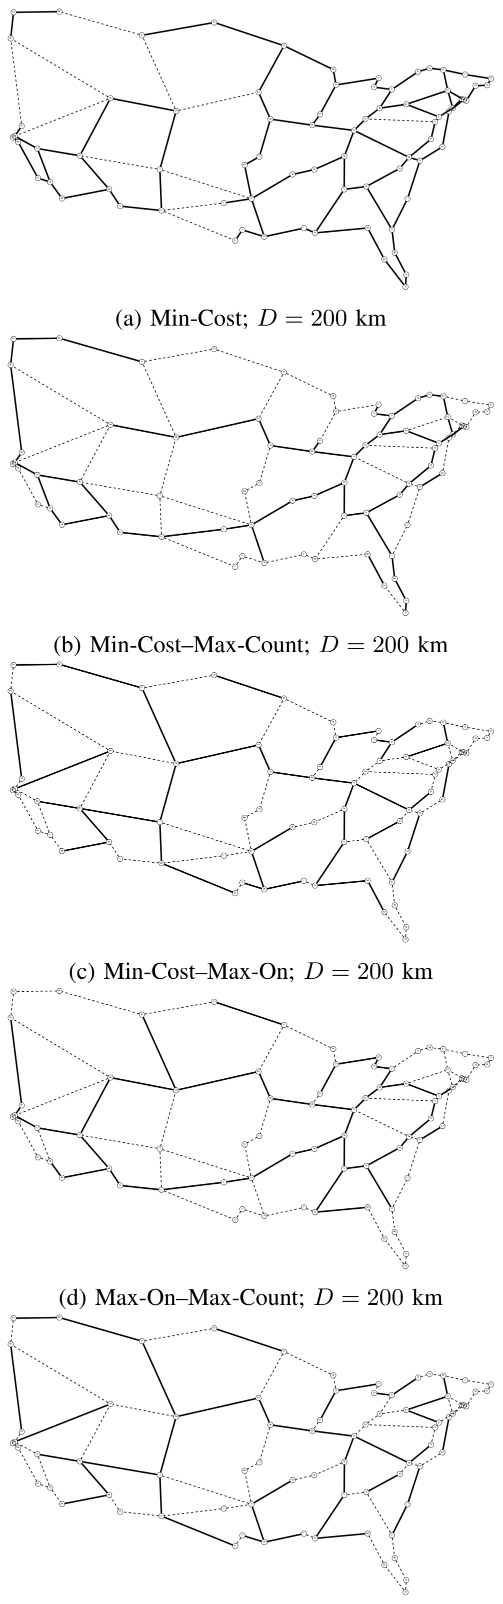

(a) Min-Cost; $D = 200$ km

(b) Min-Cost–Max-Count; $D = 200$ km

(c) Min-Cost–Max-On; $D = 200$ km

(d) Max-On–Max-Count; $D = 200$ km

(e) Max-Count–Max-On; $D = 200$ km

Fig. 4: Upgraded edges for achieving availability $\Lambda = 0.9999$ between all node pairs in Coronet, when $D = 200$ km

upgrade cost: the Max-On–Max-Count obtains the lowest (or very close to the lowest) upgrade costs with a number of upgraded edges at most 16% higher than the lowest values and the Max-Count–Max-On obtains the lowest number of upgraded edges with an upgrade cost at most 18% higher than the lowest values. On the other hand, Table IV show that for $\Lambda = 0.99999$, with the exception of Min-Cost, which is much worse, none of the other strategies represents a better compromise between the number of upgraded edges and the upgrade cost.

The results of both tables reinforce the observations already made for Germany50, namely, (1) Min-Cost is always the worst strategy, (2) the Max-Cost–Max-Count and Max-On–Max-Count strategies (that give preference to the most frequent edges in the path pairs of $\mathcal{R}$) find solutions with lower upgrade costs and (3) the Min-Cost–Max-On and the Max-Count–Max-On strategies (that give preference to edges whose upgrade makes the availability of more path pairs to become at least Λ) find solutions with a lower number of upgraded edges. Fig. 4 presents the solutions found by the five strategies for $\Lambda = 0.9999$ and $D = 200$ km. In this figure, it is possible to check that the Min-Cost strategy selects a much higher number of upgraded edges. Due to the larger dimension of this network (when compared to the dimension of the Germany50 network), the differences between the lengths of the selected edges are not so clear. However, it is possible to observe that the Min-Cost–Max-On (Fig. 4c) and the Max-Count–Max-On (Fig. 4e) strategies tend to have less short length edges.

V. Conclusions and Further Work

Telecommunication networks must provide high end-to-end availability and high resilience to large-scale disasters. Path protection improves end-to-end availability but might be not enough to reach the availability required by critical services. Moreover, adding path geodiversity to enhance disaster-resilience of networks makes the provision of high end-to-end availability even more challenging.

Here, we have addressed the problem of selecting a set of edges to be upgraded at a minimum cost ensuring a required level of availability and geodiversity. We have proposed a solving algorithm which uses an iterative approach based on a greedy strategy: starting with the network configuration without upgraded edges, the algorithm selects iteratively one edge to be upgraded until the resulting network configuration fulfils the required availability and geodiversity levels. Different edge selection strategies were proposed and tested on a set of problem instances. The computational tests showed that a simple strategy of edge selection only based on edge cost is very inefficient, while strategies taking into account the improvement impact of the selected edge on the end-to-end availability of node pairs lead to more efficient algorithms.

Regarding future work, an exact resolution approach and/or lower bounds for the cost of the solutions will be pursued for small network instances, which will allow to gain some insight into the quality of the approximate solutions and also to allow some tuning of the heuristics.

TABLE III: No. of upgraded edges, cost and CPU time for achieving $\Lambda = 0.9999$ between all node pairs in Coronet

	Min-Cost			Min-Cost–Max-Count			Min-Cost–Max-On			Max-On–Max-Count			Max-Count–Max-On		
D	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)
100	84	21758	14491.2	45	**13439**	425.3	34	14593	297.3	38	13598	385.9	**33**	13494	472.2
200	88	23877	38019.7	53	15493	616.6	39	17237	600.1	44	**15130**	580.6	**38**	17031	603.3
400	91	26084	44959.8	56	18194	1126.6	45	19610	1002.2	46	**17860**	1036.8	**44**	19376	955.7
600	92	26718	46632.4	56	18482	1149.4	47	20584	1024.5	47	**17865**	1082.5	**46**	21066	984.7

TABLE IV: No. of upgraded edges, cost and CPU time for achieving $\Lambda = 0.99999$ between all node pairs in Coronet

	Min-Cost			Min-Cost–Max-Count			Min-Cost–Max-On			Max-On–Max-Count			Max-Count–Max-On		
D	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)	Edges	Cost	CPU(s)
100	91	26409	72949.7	82	24534	1482.1	**66**	26499	1775.3	68	**23083**	1139.6	**66**	26499	1812.9
200	94	28182	166467.9	86	25966	2125.3	69	26939	13086.8	72	**24495**	1649.9	**67**	26949	4101.8
400	97	30679	189654.7	88	**27617**	6223.4	**72**	28274	16829.8	79	28045	4218.5	**72**	28696	15630.1
600	98	31625	197249.8	89	**28350**	7070.0	**71**	28786	13718.6	79	28542	4688.7	72	29247	17911.4

ACKNOWLEDGMENT

A. de Sousa was supported by Fundação para a Ciência e a Tecnologia (FCT) under project UID/EEA/50008/2013. T. Gomes, R. Girão-Silva and L. Martins were supported by FCT under project UID/MULTI/00308/2013 of INESC Coimbra. This article is based upon 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).

REFERENCES

[1] J. Rak, D. Hutchison, E. Calle, T. Gomes, M. Gunkel, P. Smith, J. T. S. Verbrugge, and L. Wosinska, "Recodis: Resilient communication services protecting end-user applications from disaster-based failures," in *18th International Conference on Transparent Optical Networks (ICTON)*, July 10-14 2016. Invited paper.

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

[3] D. Tipper, "Resilient network design: challenges and future directions," *Telecommunication Systems*, vol. 56, no. 1, pp. 5–16, 2014.

[4] A. Alashaikh, T. Gomes, and D. Tipper, "The spine concept for improving network availability," *Computer Networks*, vol. 82, no. 0, pp. 4 – 19, 2015.

[5] A. Alashaikh, D. Tipper, and T. Gomes, "Supporting differentiated resilience classes in multilayer networks," in *12th International Conference on the Design of Reliable Communication Networks (DRCN)*, pp. 31–38, March 2016.

[6] 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*, (Washington, DC, USA,), pp. 343–351, October 25-28 2009.

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

[8] S. Trajanovski, F. A. Kuipers, A. Ilić, J. Crowcroft, and P. V. Mieghem, "Finding critical regions and region-disjoint paths in a network," *IEEE/ACM Transactions on Networking*, vol. 23, pp. 908–921, June 2015.

[9] 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.

[10] J. Zhang, E. Modiano, and D. Hay, "Enhancing network robustness via shielding," in *Design of Reliable Communication Networks (DRCN), 2015 11th International Conference on the*, pp. 17–24, March 2015.

[11] T. Gomes and J. Craveirinha, "Efficient calculation of the most reliable pair of link disjoint paths in telecommunication networks," *European Journal of Operational Research*, vol. 182, no. 3, pp. 1055–1064, 2007.

[12] J. Y. Yen, "Finding the k shortest loopless paths in a network," *Management Science*, vol. 17, pp. 712–716, July 1971.

[13] E. Martins, M. Pascoal, and J. Santos, "An algorithm for ranking loopless paths," Tech. Rep. 99/007, CISUC, 1999.

[14] E. Martins, M. Pascoal, and J. Santos, "Deviation algorithms for ranking shortest paths," *International Journal of Foundations of Computer Science*, vol. 10, no. 3, pp. 247–263, 1999.

[15] S. Orlowski, R. Wessäly, M. Pióro, and A. Tomaszewski, "SNDlib 1.0–Survivable Network Design library," *Networks*, vol. 55, no. 3, pp. 276–286, 2010. http://sndlib.zib.de.

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