Resilient SDN Intercontroller Network Design under Availability Requirements and Geodiversity Constraints

Dorabella Santos
INESC Coimbra
DEEC, Rua Sílvio Lima, Polo 2,
3030-290 Coimbra, Portugal
dorabella.santos@gmail.com

Teresa Gomes, Lúcia Martins
University of Coimbra, INESC Coimbra,
Department of Electrical and Computer Engineering,
Rua Sílvio Lima, Polo 2,
3030-290 Coimbra, Portugal
{teresa,lucia}@deec.uc.pt

João P. Vidal
INESC Coimbra
DEEC, Rua Sílvio Lima, Polo 2,
3030-290 Coimbra, Portugal
joaosantosvidal_1994@hotmail.com

***Abstract*—In Software-Defined Networking (SDN), the data plane is defined by the switches which are connected to a set of controllers defining the control plane. The intercontroller connections have lacked the same focus in the literature as the switch-controller connections, despite the set of logically centralized controllers being the brain of the SDN network. In this paper, we address a bi-objective joint optimization problem involving the controller placement problem, geodiversity and intercontroller availability guarantees. The availability guarantees imply a subgraph whose links can be upgraded at a given cost. The geodiversity constraints extend robustness to disaster-based failures. To solve this NP-complete non-linear problem, we employ a heuristic strategy based on the decomposition into manageable subproblems.**

***Index Terms*—SDN, controller placement, availability, heuristics, integer linear programming, bi-objective**

I. Introduction

Software-Defined Networking (SDN) has gained importance as a key enabler of more programmable and flexible networks. The paradigm of decoupling the data and control planes, where the data plane is composed of basic forwarding switches and the control plane is composed by a set of logically centralized controllers, allows for more manageable network configuration and optimization. However, several resiliency, security and scalability issues arise.

Inherent to SDN is the well-known controller placement problem (CPP) [1], which determines how many controllers should be placed in the network and where. The CPP is known to be NP-hard and has been studied in many contexts taking into account delay, availability, resiliency and scalability.

In this work, we consider the CPP focusing mainly on intercontroller availability and resiliency [2]. Resiliency can be achieved by path protection for the more frequent link failures and can be extended to disaster-based failures by using geodiverse routing.

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 project grant UIDB/00308/2020.

Other works have addressed the CPP for resiliency purposes. In [3], the CPP is studied for intercontroller resiliency against single link failures using path protection. The authors consider a Steiner tree for the intercontroller routing, but do not consider any availability guarantees nor intercontroller delay constraints. In [4], the CPP is studied guaranteeing "five-nines" availability for switch-controller connections. The authors show that some switches need at least two backup controllers to achieve the desired availability. They do not consider any other constraints such as delay constraints. In [5], the capacitated CPP is addressed considering resiliency against link failures and controller loss. Each switch can connect to more than one controller, to prevent switch-controller connectivity loss. The problem is formulated as an integer linear programming (ILP) model aiming to minimize controller deployment cost, while ensuring connectivity via controller redundancy and protection routing. However, delay constraints are not considered. In [6], the capacitated CPP is addressed considering switch-controller delay constraints and guaranteeing path protection to maximize availability. Intercontroller delay constraints are not considered, and availability guarantees are not imposed. None of these works consider geodiversity.

In [7], geodiversity is considered between the switches and controllers. To increase availability, several paths are generated guaranteeing geographical separation between them. The controller locations are known and delay constraints are not considered. In [8], the CPP is addressed where each switch is connected to a primary controller and a backup controller with geodiverse paths. Neither delay constraints nor availability guarantees were considered in this work.

Our approach considers delay constraints, availability guarantees and geodiversity constraints simultaneously. We employ a resolution strategy inspired in our previous work [9]. In this framework, to ensure control plane performance within desired delay constraints, we use the well-known strategy of imposing maximum switch-controller and intercontroller delay bounds [9]–[11]. We further assume intercontroller availability guarantees which can be accomplished to some extent by path protection. However, it is not always possible to achieve desired availability target values with path protection alone [4], [12]. Therefore, we adopt the spine concept [13], where a set of links is selected for upgrade to have improved availability, that can be achieved by reducing the mean time to repair (e.g. by prioritizing the repair of these links in detriment to others) [14] or alternatively by increasing the mean time between failures (by making the link more robust). Furthermore, we consider geodiversity constraints to extend resiliency to disaster-based failures [8].

A very closely related problem is our previous work [11], where we studied a complementary problem. Therein, we focused on availability and resiliency via geodiverse paths between the switch-controller connections. he main differences between our previous problem and our current problem are as follows. Firstly, in the previous problem, since every switch must connect to at least one controller, the problem becomes more manageable since we can use anycast routing to model the connections, even without prior knowledge of the controller loacations. This further allows us to assume that the subgraph for improved availability is a spanning tree. In our current problem, which focuses on the intercontroller connections this cannot be assumed, making the problem more difficult to model. Secondly, in the previous problem, increasing the number of controllers generally leads to shorter switch-controller paths which tends to a lower upgrade cost. However, increasing the number of controllers may lead to longer intercontroller paths in some cases, potentially increasing the upgrade cost in our current problem. We therefore expect a smaller range of trade-off solutions for the intercontroller connections. Thirdly, to tackle the non-linear availability constraints in our previous problem, we considered a pair of availability target values for the primary and backup paths respectively, which were calculated to ensure the final desired path pair availability guarantee. This strategy is restrictive but allowed us to solve the joint optimization problem as an ILP model.

Due to the complexity of our current problem which is NP-complete, we consider a bi-objective heuristic approach inspired by [9]. In this approach, the optimization problem is divided into its constituent subproblems: (i) controller placement; (ii) determination of a subgraph supporting the primary paths; (iii) selection of the links in the subgraph to be upgraded in order to achieve the path pair availability guarantee. The approach aims to minimize two objectives: (i) the number of controllers (so as to minimize intercontroller communication overhead); (ii) and the link upgrade cost.

The main contribution of this paper is the bi-objective heuristic approach inspired by [9] and adapted to include geodiversity. We compare the solutions for a range of geodiversity distances and also compare them with the solutions obtained without geodiversity imposition.

This paper is organized as follows. In Section II, we describe the main joint optimization problem and explain why we consider a decomposition approach to solve it. In Section III, we present the CPP subproblem and formulate it as an ILP model. In Section IV, we review the definition of D-geodiversity and present the ILP model that guarantees geodiverse routing between a pair of controllers. In Section V, we revisit the link availability upgrade problem. In Section VI, we describe the heuristic resolution approach. In Section VII, the computational results are discussed and finally in Section VIII, we draw some conclusions.

II. Main Optimization Problem

The main joint optimization problem addressed involves several smaller optimization subproblems. Inherent to SDN, we have the CPP aiming to minimize the number of controllers while guaranteeing delay constraints, which is NP-hard [1]. We assume switch-controller and intercontroller maximum delay bounds: (i) the delay between each switch and its primary controller cannot exceed a given value $D_{sc} > 0$; (ii) the delay between any two controllers cannot exceed a given value $D_{cc} > 0$. Given that switches communicate with controllers much more frequently than controllers among themselves, we assume $D_{sc} < D_{cc}$. Another optimization subproblem involves the determination of a subgraph $\mathcal{S}$ whose links can be upgraded to have improved availability at a given cost. This subgraph supports the intercontroller primary paths and ensures that the intercontroller availability (achieved by each pair of primary and backup paths) reaches the minimum target λ. This problem aiming to minimize the upgrade cost is NP-complete [13]. Hence, we have two objectives: (i) minimizing the number of controllers and (ii) minimizing the upgrade cost. We have a third optimization subproblem which is guaranteeing that each pair of primary and backup paths is D-geodiverse in order to extend resiliency to disaster-based failures.

The NP-complete bi-objective joint optimization model can be generally defined as:

$\min$ *number of controllers*

$\min$ *upgrade cost*

s.t.

switch-controller assignment

switch-controller and intercontroller delay constraints

geodiverse intercontroller protection routing

intercontroller availability guarantees

The switch-controller assignment constraints are non-linear since the controller locations are not known *a priori*, but can be linearized using decision variables that resort to McCormick envelopes. In turn, the availability constraints are also non-linear and to the best of our knowledge not linearizable, since it involves a pair of paths. Note that minimizing the number of controllers does not guarantee that the upgrade cost is minimum. Therefore, we obtain a set of solutions that represent the trade-off between the two objectives, instead of a global optimum solution.

III. Controller Placement Problem

Consider that the SDN data plane is represented by a graph $G = (N, E)$ where N is the set of nodes representing the switches and E is the set of edges representing the links. Each link is represented by its end nodes $\{i, j\} \in E$ and has an associated length ℓ_{ij} and an associated delay d_{ij}. Consider A as the set of arcs or directed links, where each link $\{i, j\} \in E$ is represented by a pair of arcs (i, j) directed from i to j and (j, i) from j to i. The delay between two nodes i and j is proportional to the length of the path connecting them. The controllers are placed in the network, co-located with the switches (in-band control plane).

The CPP with delay constraints can be formulated as an ILP model. Consider the following decision variables:

y_i binary variable that is 1 if there is controller located in node i, and 0 otherwise

The ILP model is given by:

$$\min C = \sum_{i \in N} y_i \tag{1}$$

s.t.

$$\sum_{j \in N: d_{ij} \leq D_{sc}} y_j \geq 1 \quad i \in N \tag{2}$$

$$y_i + y_j \leq 1 \quad i, j \in N : d_{ij} > D_{cc} \tag{3}$$

$$y_i \in \{0, 1\} \quad i \in N \tag{4}$$

The objective function (1) aims to minimize the number of controllers C, which is one of the objectives of the joint optimization problem. Constraints (2) guarantee that for any node $i \in N$, there is a controller placed with a delay of at most D_{sc} from it. Constraints (3) guarantee that the delay between any two controllers is at most D_{cc}. Finally, constraints (4) are the variable domain constraints. Note that D_{sc} directly influences the minimum number of controllers. In turn, D_{cc} indirectly bounds the maximum number of controllers that can be placed in the network.

IV. D-geodiversity

In this work, we consider path protection via a pair of primary and backup paths between each pair of controllers. To extend protection to disaster-based failures we also consider geodiversity. It is possible to circumvent a regionally delimited failure with a coverage diameter of at most $D > 0$, by ensuring that the primary and backup paths between two controllers are geographically separated at least D, as long as the controllers themselves are not in the affected region. In this way, if one of the paths transverses the affected region, it is guaranteed that the other path does not. We adopt the definition used in [8], that generalizes the notion in [15] to consider link geodiversity including for the links incident to the controllers.

Denote the geographical distance between two links $e_1, e_2 \in E$ as $\delta_e(e_1, e_2)$ and define it as the infimum geographical distance between any point of e_1 and any point of e_2, i.e.,

$$\delta_e(e_1, e_2) = \inf_{\substack{w_1 \in e_1 \\ w_2 \in e_2}} \delta(w_1, w_2) \tag{5}$$

Therefore, adjacent links have zero geographical distance between them, since they share an end node. The definition is easily adapted to define the distance between a node $n \in N$ and a link $e \in E$, denoted as $\delta_{ne}(n, e)$ and defined as the infimum distance between node n and any point of edge e, i.e.,

$$\delta_{ne}(n, e) = \inf_{w \in e} \delta(n, w) \tag{6}$$

Assume the controller placements are known and denote the controller set as $\mathcal{C}$, with $|\mathcal{C}| = C$. Consider two controllers $c_1, c_2 \in \mathcal{C}$. The primary path $p^{c_1 c_2}$ and the backup path $b^{c_1 c_2}$ connecting both controllers are said to be D-geodiverse if any link of $p^{c_1 c_2}$ that is not incident to c_1 nor c_2 is geographically separated from any link of $b^{c_1 c_2}$ at least D, and vice-versa. This is illustrated in Fig. 1(a). The links incident to the controllers are considered to be special cases. Consider the link in $p^{c_1 c_2}$ incident to c_1 denoted as $e_1 = \{c_1, i\}$ and the link in $b^{c_1 c_2}$ incident to c_1 denoted as $e_2 = \{c_1, j\}$. The D-geodiversity is guaranteed here by considering that the geographical distances between node i and link e_2, and between node j and link e_1 are at least D. In other words, the adjacent links to the controllers are D-geodiverse if each link is geographically separated at least D of the opposite end node of the other link, i.e., if $\delta'(e_1, e_2)$ defined as $\delta'(e_1, e_2) = \min\{\delta_{ne}(i, e_2), \delta_{ne}(j, e_1)\}$ is at least D. This is illustrated in Fig. 1(b).

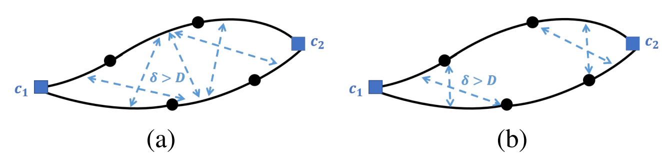

Fig. 1. D-geodiverse paths between controllers c_1 and c_2. (a) D-geodiversity guaranteed between the non-adjacent links of both paths. (b) D-geodiversity guaranteed between the adjacent links incident to the controllers.

To define the geodiverse routing constraints between a pair of controllers c_1 and c_2, consider the following parameters:

ζ_i parameter that is 1 if $i = c_1$, -1 if $i = c_2$, and 0 otherwise

and the following decision variables:

$x_{ij}^{c_1 c_2}$ binary variable that is 1 if arc $(i, j) \in A$ belongs to the *primary* path between c_1 and c_2

$u_{ij}^{c_1 c_2}$ binary variable that is 1 if arc $(i, j) \in A$ belongs to the *backup* path between c_1 and c_2

Moreover, consider set $\mathcal{P}_{c_1 c_2}$ as the set of incompatible link pairs w.r.t. D-geodiversity between c_1 and c_2. Then $\mathcal{P}_{c_1 c_2}$ includes the pair of links $e_1, e_2 \in E$ not incident to c_1 nor c_2 such that $\delta_e(e_1, e_2) < D$. It also includes the pairs of adjacent links $e_1, e_2 \in E$ such that the common end node is either c_1 or c_2 and $\delta'(e_1, e_2) < D$.

Since the path availability depends inversely on its length (as will be explained in Section V), we are interested in finding the pair of paths between c_1 and c_2 such that they are D-geodiverse and such that the sum of their delays is minimum. So the following ILP model is solved for each pair of controllers $c_1, c_2 \in \mathcal{C}$:

$$\min \sum_{\{i,j\}\in E} d_{ij}(x_{ij}^{c_1c_2} + x_{ji}^{c_1c_2} + u_{ij}^{c_1c_2} + u_{ji}^{c_1c_2}) \quad (7)$$

s.t.

$$\sum_{j\in V(i)} (x_{ij}^{c_1c_2} - x_{ji}^{c_1c_2}) = \zeta_i \qquad i \in N \quad (8)$$

$$\sum_{j\in V(i)} (u_{ij}^{c_1c_2} - u_{ji}^{c_1c_2}) = \zeta_i \qquad i \in N \quad (9)$$

$$\sum_{\{i,j\}\in E} d_{ij}(x_{ij}^{c_1c_2} + x_{ji}^{c_1c_2}) \le D_{cc} \quad (10)$$

$$x_{ij}^{c_1c_2} + x_{ji}^{c_1c_2} + u_{ij}^{c_1c_2} + u_{ji}^{c_1c_2} \le 1 \qquad \{i,j\} \in E \quad (11)$$

$$x_{ij}^{c_1c_2} + x_{ji}^{c_1c_2} + u_{lk}^{c_1c_2} + u_{kl}^{c_1c_2} \le 1 \qquad \{\{i,j\},\{l,k\}\} \in \mathcal{P}_{c_1c_2} \quad (12)$$

$$x_{ij}^{c_1c_2}, u_{ij}^{c_1c_2} \in \{0,1\} \qquad (i,j) \in A \quad (13)$$

The objective function (7) of this subproblem minimizes the total delay of the primary and backup paths. Constraints (8) and (9) guarantee a primary path and backup path between c_1 and c_2, respectively. Constraint (10) guarantees that the primary path satisfies the intercontroller delay constraint; this constraint is relaxed for the backup path since it must be D-geodiverse with the primary one and so if necessary may have a longer delay (but the objective function avoids it being longer than necessary). Constraints (11) ensure that the paths are link disjoint. These constraints are needed to complete (12) that together ensure that the primary and backup paths are D-geodiverse. Note that considering $D = 0$ makes these constraints inactive, reducing the model to simply guarantee link-disjoint path protection. In this case, there is still, in general, some geographical separation between the paths. Finally, constraints (13) are the variable domain constraints.

If D is too large, it may not be possible to guarantee the desired geodiversity between certain pairs of controllers. For each pair of controllers $c_1, c_2 \in \mathcal{C}$, there is a maximum geographical distance $D_{\max}^{c_1c_2}$ beyond which we cannot guarantee a greater geodiversity of the path pair. These distances can be determined in advance by solving an optimization problem [16]. In those cases, we consider the geodiversity distance to be $D'_{c_1c_2} = \min\{D, D_{\max}^{c_1c_2}\}$.

V. Link Availability Upgrade

For the sake of self-containment, in this section we revisit the link availability upgrade model already presented in [9].

The intercontroller availability is determined by the availability of the primary and backup path pairs. We impose that each pair of paths must have an availability of at least λ. In general, the target availability cannot be achieved without link upgrade. Consider a subgraph $\mathcal{S}$ to support the primary paths and whose links can be upgraded to have improved availability at a given cost.

Consider that each link $\{i,j\} \in E$ has a default availability α_{ij}^0 that depends on its length ℓ_{ij} (see details in [11]). Assume that each link of $\mathcal{S}$ can be upgraded incrementally up to κ levels, where in each level the link unavailability is decreased by a factor of $\varepsilon \in (0,1)$. Therefore, the link upgrade availability in level $k = 1, ..., \kappa$ is given by $\alpha_{ij}^k = \alpha_{ij}^{k-1} + \varepsilon(1 - \alpha_{ij}^{k-1})$, at a cost given by [13]:

$$c_{ij}^k = -\ell_{ij} \cdot \ln\left(\frac{1-\alpha_{ij}^k}{1-\alpha_{ij}^0}\right) \qquad k = 1, .., \kappa \quad (14)$$

The cost function increases exponentially as the link is upgraded to the next level.

The availability of a path is given by the product of the availability of its elements (links and nodes). To simplify the model, we consider only the link availabilities, i.e., the node availabilities are considered to be 1. Although this is not totally realistic, it is known that the node availabilities are typically much higher than the link availabilities [17].

Denote the primary and backup path availabilities between controllers c_1 and c_2 as $\mathcal{A}_p^{c_1c2}$ and $\mathcal{A}_b^{c_1c_2}$, respectively. Then the availability guarantee for that path pair is expressed as

$$1 - (1 - \mathcal{A}_p^{c_1c_2})(1 - \mathcal{A}_b^{c_1c_2}) \ge \lambda \quad (15)$$

Since $\mathcal{S}$ supports the primary paths and its links are subject to upgrade, the link availability improvement will be reflected in the primary paths. Note that, however, some backup paths can use links of $\mathcal{S}$ as long as the D-geodiverse constraints are satisfied. Consider that the set of primary paths obtained by the ILP model (7)-(13), between each pair of controllers, forms the subgraph $\mathcal{S}$. In other words, subgraph $\mathcal{S}$ is set of links defined by $\mathcal{S} = \{\{i,j\} \in E : x_{ij}^{c_1c_2} = 1,\, c_1, c_2 \in \mathcal{C}\}$.

In the case where geodiversity is relaxed, i.e. $D = 0$, we simply have link-disjoint path protection. This case was studied in [9] where we considered $\mathcal{S}$ as the set of links of the shortest paths or as a Steiner tree. We also considered a downgrade level for link availability that is not considered here.

To ensure the availability guarantees, consider the default availability of the backup path $b^{c_1c_2}$ which is given by (see [11] for details):

$$\mathcal{A}_b^{c_1c_2} = \prod_{\{i,j\}\in b^{c_1c_2}} \alpha_{ij}^0 \quad (16)$$

and can be linearized in the following way:

$$\log(\mathcal{A}_b^{c_1c_2}) = \sum_{\{i,j\}\in b^{c_1c_2}} \log(\alpha_{ij}^0) \quad (17)$$

Then, the necessary availability of the primary path $p^{c_1c_2}$ to ensure (15) is given by:

$$\mathcal{A}_p^{c_1c_2} = \frac{\lambda - \mathcal{A}_b^{c_1c_2}}{1 - \mathcal{A}_b^{c_1c_2}} \quad (18)$$

Hence, the optimization subproblem to select the links of $\mathcal{S}$ for upgrade can be formulated as an ILP model. Consider the following decision variables:

- z_{ij}^k binary variable that is 1 if link $\{i,j\} \in \mathcal{S}$ is in level k, with $k = 0, 1, ..., \kappa$, and 0 otherwise

The ILP model is then given by:

$$\min \sum_{\{i,j\} \in \mathcal{S}} \sum_{k=1}^{\kappa} c_{ij}^k z_{ij}^k \quad (19)$$

s.t

$$\sum_{k=0}^{\kappa} z_{ij}^k = 1 \qquad \{i,j\} \in \mathcal{S} \quad (20)$$

$$\sum_{k=0}^{\kappa} \sum_{\{i,j\} \in p_{c_1 c_2}} z_{ij}^k \log(\alpha_{ij}^k) \geq \log(\mathcal{A}_p^{c_1 c_2}) \quad c_1, c_2 \in \mathcal{C} \quad (21)$$

$$z_{ij}^k \in \{0,1\} \qquad \{i,j\} \in \mathcal{S}, k = 0, 1, ..., \kappa \quad (22)$$

The objective function (19) minimizes the cost of upgrading the links of $\mathcal{S}$, and is the other objective function of the main joint optimization problem. Constraints (20) guarantee that each link $\{i,j\} \in \mathcal{S}$ is either not upgraded ($k = 0$) or is upgraded to one of the levels $k = 1, ..., \kappa$ (it cannot be in more one level simultaneously). Constraints (21) guarantee that the primary paths have the necessary availability $\mathcal{A}_p^{c_1 c_2}$. The path pair availability constraints (15) are not linearizable, but using the approach in (16)-(18), it is possible to linearize only the primary path availability constraints (21). Finally, constraints (22) are the variable domain constraints.

VI. Resolution Strategy

Since the main optimization model is impractical to solve (as explained in Section II), we employ an iterative decomposition approach where its constituent subproblems are solved sequentially in each iteration. This heuristic strategy is defined as follows:

(i) The CPP is the first subproblem to be solved using the ILP model given by (1)-(4), obtaining a set with the minimum number of controllers.
(ii) Once the controller placement is obtained in step (i), the intercontroller D-geodiverse path pairs are determined using the ILP model given by (7)-(13), which is solved for each pair of controllers.
(iii) Finally, $\mathcal{S}$ is defined as the set of links belonging to the primary paths obtained in step (ii). The ILP model given by (19)-(22) is solved, returning the minimum upgrade cost for $\mathcal{S}$.

Then the approach goes back to the CPP in step (i) to determine a new set of controllers $\mathcal{C}$, by adding constraints that remove the previously obtained controller sets from the search space. Assume that $C_{\min}$ is the minimum number of controllers given by the ILP model (1)-(4). Consider that the set of controllers $\gamma_1, ..., \gamma_{C_{\min}}$ was obtained. To obtain a new set of controllers, the additional constraint

$$y_{\gamma_1} + \cdots + y_{\gamma_{C_{\min}}} \leq C_{\min} - 1 \quad (23)$$

is added to the ILP model. A similar constraint is added for each previously obtained controller set with $C_{\min}$ controllers. Eventually, the new set of controllers will have an incremented number of controllers $C = C_{\min} + 1$, in which case, the set of constraints (23) changes to

$$\sum_{i \in N} y_i \geq C \quad (24)$$

$$y_{\gamma_1} + \cdots + y_{\gamma_C} \leq C - 1 \quad (25)$$

Constraint (25) is added for each previously obtained controller set with C controllers. The set of constraints is updated accordingly as C is incremented.

The subproblems are solved sequentially obtaining an upgrade cost for each controller set. The best trade-off solutions are stored. A best trade-off solution sol is such that if any other solution has a smaller number of controllers then its cost is greater than that of sol, and vice-versa. The process is repeated until one of the following stopping criteria is met:

(a) The upgrade cost is zero, i.e., the objective function (19) is zero meaning that no links of $\mathcal{S}$ needed to be upgraded. Since the controllers sets are generated with non-decreasing C in the CPP model, there cannot be any new solution with a lower cost and a lower C, so the heuristic can be stopped.
(b) The CPP model becomes infeasible meaning that all controller sets satisfying D_{sc} and D_{cc} have been obtained.
(c) A given maximum number M of consecutive solutions are generated that do not improve the upgrade cost. This criterium is useful when the CPP search space is too large and the zero upgrade cost is not achievable.

In Fig. 2, the cost266 network from SNDlib [18] is illustrated for given maximum D_{sc} and D_{cc} values. The best solutions obtained for $C = 2$ controllers are shown for link-disjoint path protection $D = 0$ (top network), and with geodiversity $D = 260$ km (bottom network).

The two controller locations are represented as red circles. The primary path is represented in red and the backup path in blue. The primary path has all its links upgraded in order to achieve $\lambda = 0.99999$, known as "five-nines" availability [4].

We considered $\kappa = 4$ levels of upgrade where $\varepsilon = 0.5$. In the top network, one link is upgraded to level $k = 3$, while the remaining three are upgraded to $k = 4$ (thicker links), with a cost of 1402.7. In this case although geodiversity is not imposed, the path pair is geographically separated 144 km (the pair of edges inducing this separation is shown in the graph by the arrows). In the bottom network, the geodiversity guarantee is of 260 km, making the paths quite long. So to achieve λ, all seven links composing the primary path are upgraded to level $k = 4$, with a cost of 2617.8. The maximum path geodiversity possible between this pair of controllers $D_{\max}^{c_1 c_2}$ is of 281 km which is greater than 260 km, meaning that the geodiversity of 260 km is indeed guaranteed.

VII. Computational Results

For our computational results, we considered the following networks: polska and cost266 from SNDlib [18], and spain from [19]. The topological characteristics of the networks are summarized in Table I, which shows the number of nodes, the number of edges, the average node degree and the graph

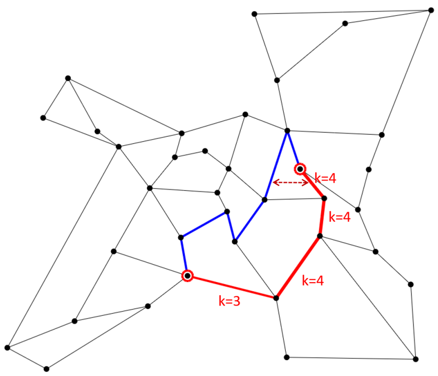

Without geodiversity constraints

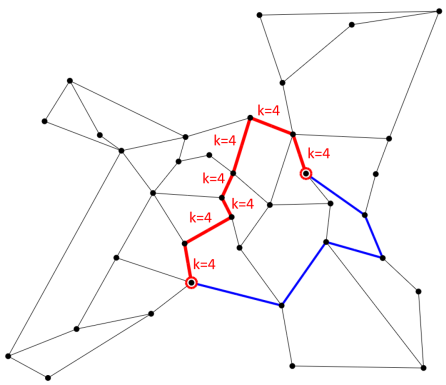

With geodiversity constraints: $D = 260$ km

Fig. 2. The best solutions for cost266 network with $C = 2$ without geodiversity constraints (top) and with geodiversity $D = 260$ km (bottom). The two controller locations are shown with red circles; the primary path is shown in red with upgraded links; the backup path is shown in blue.

diameter D_g (longest shortest path between any two nodes) for each network.

TABLE I
TOPOLOGICAL CHARACTERISTICS OF THE NETWORKS

Network	#nodes	#links	avg deg	D_g [km]
polska	12	18	3.00	811
spain	14	22	3.14	1034
cost266	37	57	3.08	4032

The heuristic approach was implemented in C/C++, using CPLEX 12.9 Callable libraries for solving the ILP models. The computational results were obtained on an Intel Core i7 laptop with 8 GB of RAM, running at 2.9 GHz.

The maximum delay values D_{sc} and D_{cc} are given as percentages of the graph diameter D_g [1], [10], [20]. Two sets of values were considered: set $\mathcal{D}_1$ where $D_{sc} = 0.40D_g$ and $D_{cc} = 0.70D_g$; and set $\mathcal{D}_2$ where $D_{sc} = 0.45D_g$ and $D_{cc} = 0.75D_g$. For the link upgrade, we consider $\kappa = 4$ levels of upgrade with $\varepsilon = 0.5$, and consider the intercontroller availability target of "five-nines" $\lambda = 0.99999$.

For geodiversity, we considered values between 100 km and 260 km with steps of 40 km in-between, i.e., $D = \{100, 140, 180, 220, 260\}$ in km. We compare the results for the different geodiversity values and also consider the case when $D = 0$. Furthermore, we also consider the case when $\mathcal{S}$ is the set of shortest paths between controllers (as done in [9]) and compare the results with those for $D = 0$, since in general the path pairs may be different. When D cannot be guaranteed between some pair of controllers c_1 and c_2, then $D^{\max}_{c_1 c_2}$ is considered.

The trade-off solutions for C and the upgrade cost are shown in Fig. 3 for polska, Fig. 4 for spain and Fig. 5 for cost266. The top chart refers to the set of delay bounds $\mathcal{D}_1$, while the bottom chart refers to the set $\mathcal{D}_2$. We note that in all the tested cases the trade-off solutions obtained with $D = 0$ are the same as those obtained with the shortest paths.

For polska (Fig. 3) with set $\mathcal{D}_1$ of delay bounds (top chart), we have obtained a global optimum with $C_{\min} = 3$ for the different values of D. When geodiversity is not imposed ($D = 0$), we obtain the ideal global optimum of zero upgrade cost (blue circle). When imposing $D = 100$ km the cost is increased to 85.2 (green diamond), and when $D = 140$ km the cost is further increased to 111.7 (orange triangle). From 140 km onwards $D^{\max}_{c_1 c_2}$ was considered for most controller pairs, and therefore, there is no cost increase.

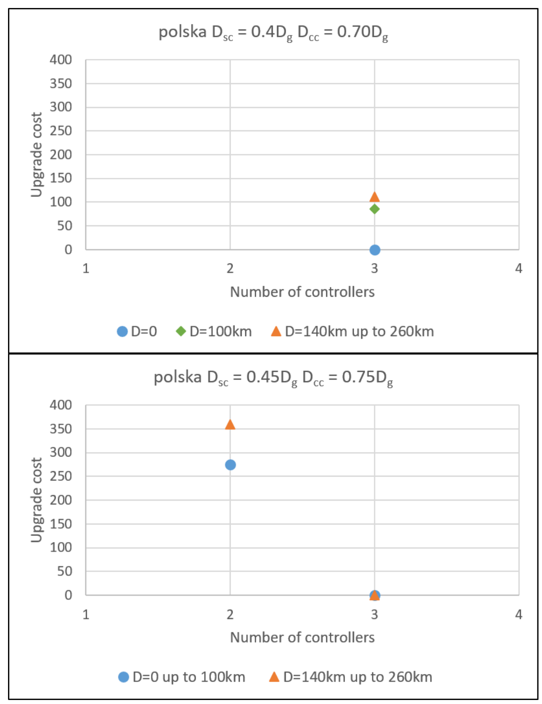

Fig. 3. Trade-off solutions between the number of controllers (x-axis) and the upgrade cost (y-axis) for polska

Considering set $\mathcal{D}_2$ of delay bounds (bottom chart), we have two trade-off solutions for the different values of D. In all cases, there is the zero cost solution with $C = 3$ meaning that despite the geodiversity imposed having 3 controllers is enough to ensure the target availability λ. There is also a trade-off solution for $C_{\min} = 2$. When $D = 0$ the cost of this solution is 274.9 (blue circle), which is the same for $D = 100$ km (since the path pairs for $D = 0$ are already geographically separated slightly more than 100 km). When $D = 140$ km the cost is increased to 360.0 (orange triangle) and is the same for $D = 180$ km. From 180 km onwards $D^{\max}_{c_1c_2}$ was considered (so there is no change in the cost).

For spain (Fig. 4) with set $\mathcal{D}_1$ of delay bounds (top chart), we have obtained two trade-off solutions for $D = 0$ which are the same for D up to 140 km: one for $C_{\min} = 4$ with a cost of 1031.4 and one for $C = 5$ with a cost of 771.5 (blue circles). When $D = 180$ km the cost of $C = 5$ is increased to 1713.6, making it no longer a trade-off solution (outlined orange triangle), because we have a solution for $D = 180$ km with $C = 4$ and cost 1031.4 (orange triangle overlapping blue circle). From 180 km onwards $D^{\max}_{c_1c_2}$ was considered.

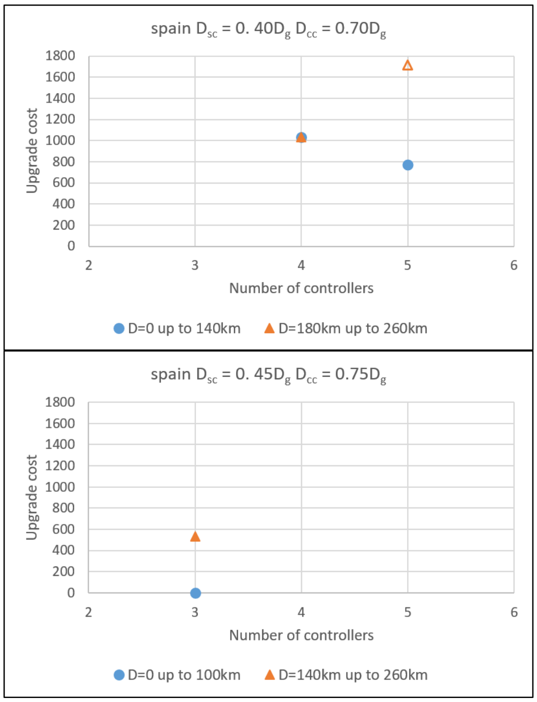

Fig. 4. Trade-off solutions between the number of controllers (x-axis) and the upgrade cost (y-axis) for spain (filled-in markers indicate trade-off solutions, while outlined ones indicate they are no longer trade-off solutions)

Considering set $\mathcal{D}_2$ of delay bounds (bottom chart), we have obtained a global optimum with $C_{\min} = 3$ for the different values of D. When $D = 0$ up to 100 km, we obtain the ideal global optimum of zero upgrade cost (blue circle). When imposing $D = 140$ km the cost is increased to 530.3 (orange triangle). Imposing higher geodiversity distances does not affect the upgrade cost, although $D^{\max}_{c_1c_2}$ is only considered from 220 km onwards.

For cost266 (shown in Fig. 5) with set $\mathcal{D}_1$ of delay bounds (top chart), we have obtained a global optimum with $C_{\min} = 2$ for the different values of D. When $D = 0$ the cost of this solution is 4136.1 which is the same up to $D = 220$ km (blue circle). Imposing $D = 260$ km the cost increases to 5173.9 (orange triangle). There was no need to use $D^{\max}_{c_1c_2}$, since D could be guaranteed for all path pairs. This case is the one illustrated in Fig. 2 where for $D = 0$ the path pair is geographically separated 144 km. When imposing geodiversity up to 220 km the best solution is with the same primary path and same upgraded link levels, although the backup path changes to ensure the desired geodiversity, but still guaranteeing the path pair availability of at least λ. When imposing $D = 260$ km both paths need to change increasing the upgrade cost.

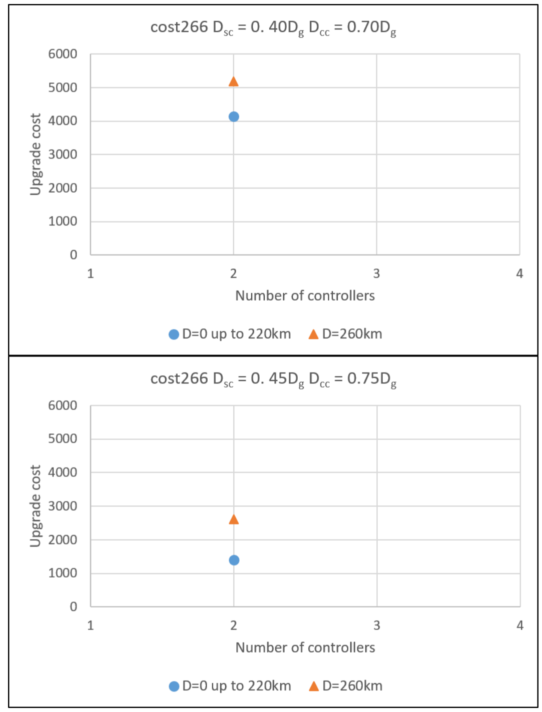

Fig. 5. Trade-off solutions between the number of controllers (x-axis) and the upgrade cost (y-axis) for cost266

Considering set $\mathcal{D}_2$ of delay bounds (bottom chart), we have also obtained a global optimum with $C_{\min} = 2$ for the different values of D. When $D = 0$ the cost of this solution is 1402.7 (blue circle) which is the same up to $D = 220$ km (because of a similar behaviour as the previous case). Imposing $D = 260$ km the cost increases to 2617.8 (orange triangle). There was no need to use $D^{\max}_{c_1c_2}$.

Finally, Fig. 6 shows the runtimes of the heuristic approach for the two sets $\mathcal{D}_1$ and $\mathcal{D}_2$ and for the three networks. The runtimes are shown for the different values of D and also when $\mathcal{S}$ is the set of shortest paths (for the case without geodiversity imposition). At first glance, the runtimes of the heuristic for the shortest paths are much smaller, significantly highlighted for the two larger networks. This is expected since the shortest paths are obtained much more efficiently than solving the ILP model (7)-(13) with $D = 0$.

For polska and spain, the runtimes are higher for set $\mathcal{D}_2$ of delay bounds since the search space of the CPP sets increases significantly in relation to $\mathcal{D}_1$, while for cost266 the search space does not increase so much. Moreover, the runtimes for the different values of D do not change significantly, meaning that the complexity of the ILP model for finding a pair of geodiverse paths is not particularly sensitive to D.

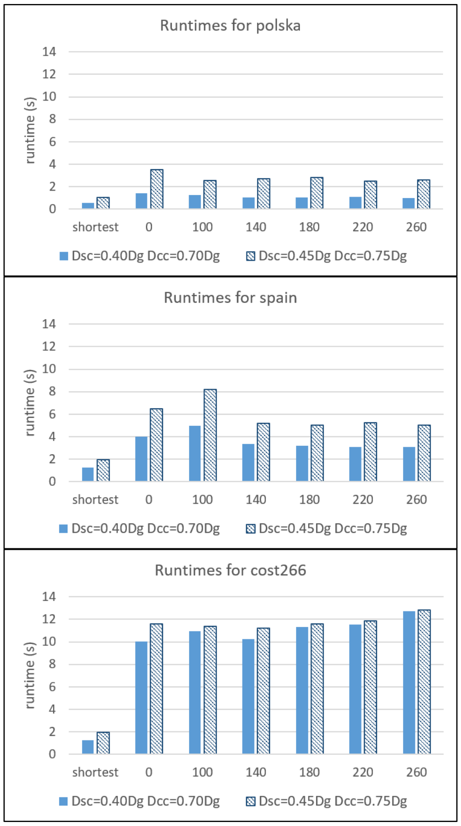

Fig. 6. Runtimes (in seconds) for he heuristic approach with different values of D and for the shortest paths

The runtimes are higher for the larger network, as expected, but still below 14 seconds.

VIII. Conclusions

In this paper, we presented the bi-objective optimization problem aiming to minimize both the number of controllers and the upgrade cost, while dealing with delay, availability and geodiversity constraints. Due to its complexity we have proposed a heuristic approach inspired by our previous work and adapted to include geodiversity. The best trade-off solutions were obtained and discussed.

The computational results show that the best trade-off solutions for $D = 0$ are the same as using the shortest paths for the primary paths, although in general this may not happen. As expected, the runtimes for the shortest paths are much lower. Moreover, when $D = 0$ and geodiversity is not imposed the path pairs are nevertheless separated by some distance D', meaning that the solution is valid for geodiversity values up to D'. It can still be valid for greater geodiversity values if only the backup paths need to change to satisfy the constraints. In the case, where the paths do change to accommodate the necessary geodiversity, the upgrade cost will increase because the paths will be longer. We note that, based on our previous work [9], the cost for the case without geodiversity imposition can be lower if trees are used instead of the shortest paths.

References

[1] B. Heller, R. Sherwood, and N. McKeown, “The controller placement problem,” in *ACM 1st Workshop on Hot topics in Software Sefined Networks (HotSDN)*, New York, USA, 2012, pp. 7–12.

[2] T. Zhang, A. Bianco, and P. Giaccone, “The role of inter-controller traffic in SDN controllers placement,” in *2nd IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, Palo Alto, USA, 2016, pp. 87–92.

[3] T. Das and M. Gurusamy, “Controller placement for resilient network state synchronization in multi-controller SDN,” *IEEE Communications Letters*, vol. 24, no. 6, pp. 1299–1303, 2020.

[4] F. J. Ros and P. M. Ruiz, “Five nines of southbound reliability in software-defined networks,” in *ACM HotSDN*, New York, USA, 2014, pp. 31–36.

[5] A. A. Seyedkolaei, S. A. H. S. A. Moradi, and R. Budiarto, “Cost-effective survivable controller placement in software-defined networks,” *IEEE Access*, vol. 9, pp. 129 130–129 140, 2021.

[6] B. Zhang, X. Wang, and M. Huang, “Multi-objective optimization controller placement problem in internet-oriented software defined network,” *Computer Communications*, vol. 123, pp. 24–35, 2018.

[7] Y. Cheng, M. M. Rahman, S. Gangadhar, M. J. F. Alenazi, and J. P. G. Sterbenz, “Cross-layer framework with geodiverse routing in software-defined networking,” in *11th International Conference on Network and Service Management (CNSM)*, 2015, pp. 348–353.

[8] A. de Sousa and D. Santos, “The minimum cost D-geodiverse anycast routing with optimal selection of anycast nodes,” in *13th International Conference on the Design of Reliable Communication Networks (DRCN)*, Coimbra, Portugal, 2019, pp. 1–7.

[9] D. Santos, J. P. Vidal, T. Gomes, and L. Martins, “Improving east/westbound SDN connectivity via enhanced availability,” in *12th International Conference on Network of the Future (NoF)*, Coimbra, Portugal, 2021, pp. 1–9.

[10] N. Perrot and T. Reynaud, “Optimal placement of controllers in a resilient SDN architecture,” in *12th International Conference on the Design of Reliable Communication Networks (DRCN)*, Paris, France, 2016, pp. 145–151.

[11] D. Santos, T. Gomes, and D. Tipper, “SDN controller placement with availability upgrade under delay and geodiversity constraints,” *IEEE Transactions on Network and Service Management*, vol. 18, no. 1, pp. 301–314, 2021.

[12] S. Song, H. Park, B. Choi, T. Choi, and H. Zhu, “Control path management framework for enhancing software-defined network (SDN) reliability,” *IEEE Transactions on Network and Service Management*, vol. 14, no. 2, pp. 302–316, 2017.

[13] A. Alashaikh, D. Tipper, and T. Gomes, “Embedded network design to support availability differentiation,” *Annals of Telecommunications*, vol. 74, no. 9-10, pp. 605–623, 2019.

[14] W. Grover and A. Sack, “High availability survivable networks: When is reducing MTTR better than adding protection capacity?” in *6th International Conference on Design and Reliable Communication Networks (DRCN)*, Oct 2007, pp. 1–7.

[15] 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, no. 3, pp. 908–921, 2015.

[16] A. de Sousa, D. Santos, and P. Monteiro, “Determination of the minimum cost pair of D-geodiverse paths,” in *DRCN*, Munich, Germany, 2017.

[17] M. Tarifeño-Gajardo, A. Beghelli, and E. Moreno, “Availability-driven optimal design of shared path protection in WDM networks,” *Networks*, vol. 68, no. 3, pp. 224–237, 2016.

[18] 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.

[19] R. Martinez, R. Casellas, R. Vilalta, and R. Muñoz, “GMPLS/PCE-controlled multi-flow optical transponders in elastic optical networks [invited],” *IEEE/OSA Journal of Optical Communications and Networking*, vol. 7, no. 11, pp. 71–80, 2015.

[20] M. Tanha, D. Sajjadi, R. Ruby, , and J. Pan, “Capacity-aware and delay-guaranteed resilient controller placement for software-defined WANs,” *IEEE Transactions on Network and Service Management*, vol. 15, no. 3, pp. 991–1005, 2018.