

Optimizing primary and backup SDN controllers' placement resilient to node-targeted attacks

Mariusz Mycek, Michał Pióro, Artur Tomaszewski
Institute of Telecommunications
Warsaw University of Technology
 Warsaw, Poland
 {m.mycek,m.pioro,a.tomaszewski}@tele.pw.edu.pl

Amaro de Sousa
Instituto de Telecomunicações
DETI, Universidade de Aveiro
 Aveiro, Portugal
 asou@ua.pt

Abstract—In Software Defined Networks (SDNs), a number of controllers are placed in a given data plane network. In a standard logically centralized control plane, each controller acts simultaneously as a primary controller for some switches and as a backup controller for other switches, and the controller placements must meet given switch-controller (SC) and controller-controller (CC) delay bounds. Then, the SDN should be resilient to network disruptions such as node-targeted attacks. To improve the SDN resilience to this kind of disruptions, we assume that some controllers are deployed only as backup controllers so that they take over the functions of primary controllers only in case of disruption. We propose an optimization model that solves a relevant primary and backup controller placement problem, where a minimum number of primary controllers minimizing the maximum SC delay is first established, and then a joint primary and backup controller placement maximizing the resilience of the SDN against a list of the most dangerous node-targeted attacks is determined. A numerical study illustrating the merits of the proposed optimization methodology is presented.

Index Terms—Targeted node attacks, SDN controllers, network resilience optimization

I. INTRODUCTION

In Software Defined Networks (SDNs), multiple controllers are placed in a given data plane network composed of switches and interconnecting links. When a packet of a new traffic flow reaches a switch, it queries its primary controller (PC in short), which then replies with the routing decision on how to forward the new flow. So, the maximum switch-controller (SC) delay (i.e., the delay between each switch and its PC) must be bounded so that the switches do not wait too long for routing decisions. In a standard logically centralized control plane [1], all controllers have a complete view of the network state, and the PC of each switch is the closest controller in terms of delay. So, all controllers act simultaneously as primary for some switches and backup controllers (BC in short) for others, and each new routing decision taken by a controller is sent to all other controllers so that the complete network state

This work was supported by the National Science Centre, Poland, under the grant no. 2017/25/B/ST7/02313: “Packet routing and transmission scheduling optimization in multi-hop wireless networks with multicast traffic”, and by POB Research Centre Cybersecurity and Data Science of Warsaw University of Technology within the Excellence Initiative Program - Research University (ID-UB). Amaro de Sousa was supported by ERDF Funds through the Centre’s Regional Operational Program, Portugal, and by Portuguese Funds, through FCT, under the project CENTRO-01-0145-FEDER-029312.

information is constantly updated. The time interval until the complete network state view is reached is, in the worst case, the longest delay among all pairs of controllers, and hence the maximum controller-controller (CC) delay must also be bounded. Finally, the controller placement must be resilient to network disruptions, such as natural disasters [2], multiple link failures [3] or node-targeted attacks [4]. To improve the SDN resilience to these kinds of disruptions, in general additional controllers are required. However, the SC and CC delay bounds limit the potential gains that can be reached when all controllers act as PCs.

Here we consider that some controllers are deployed only as BCs [5]; they are not active in the normal state and they take over the functions of PCs only in case of disruption. Since these controllers acting solely as BCs do not participate in the routing decisions in the normal state, they do not require significant processing power and their placement does not need to meet the SC and CC delay bounds. We address a primary and backup controller placement problem in which only the PC placement must meet given SC and CC delay bounds. We propose an optimization model where a minimum number of PCs minimizing the maximum SC delay is first established, and then a joint primary and backup controller placement maximizing the resilience of the SDN against a list of the most dangerous node-targeted attacks is determined.

The paper is organized as follows. Section II describes the basic concepts behind the addressed problem and how the resilience to node-targeted attacks is measured. Sections III and IV describe the optimization methodology and Section V presents numerical results illustrating its merits. Finally, concluding remarks are given in Section VI.

II. NETWORKS, CONTROLLER PLACEMENTS, ATTACKS, AVAILABILITY MEASURES

The data plane network is modeled by an undirected connected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with the set of (switching) nodes \mathcal{V} (of size V), and the set of links $\mathcal{E} \subseteq \mathcal{V}^2$. The controllers are collocated with a selected subset of the nodes in \mathcal{V} , denoted by \mathcal{S} (where $\mathcal{S} \subseteq \mathcal{V}$). A node with a collocated controller is named a controller node and each node can be equipped with at most one controller (either primary or backup).

TABLE I
 SUMMARY OF GENERAL NOTATION

\mathcal{V}	set of nodes ($V := \mathcal{V} $)
\mathcal{E}	set of links ($E := \mathcal{E} $)
\mathcal{S}	set of nodes with controllers
\mathcal{A}	set (family) of attacks
$\mathcal{V}(a)$	set of nodes affected by attack $a \in \mathcal{A}$ ($V(a) := \mathcal{V}(a) $)
$\mathcal{C}(a)$	set (family) of components c resulting from attack $a \in \mathcal{A}$
$\mathcal{V}(c)$	set of nodes of component c ($V(c) := \mathcal{V}(c) $)
$Y(a, c)$	binary indicator equal to 1 iff $\mathcal{V}(c) \cap \mathcal{S} \neq \emptyset$ ($c \in \mathcal{C}(a)$)
$f(c)$	metric of component c
$\mathcal{M}(a)$	network availability measure defined for attack a
$\mathcal{M}(\mathcal{A})$	network availability measure defined for family of attacks \mathcal{A}
$\mathcal{X}^{[2]}$	set of all 2-element subsets of a given set \mathcal{X}
$ \mathcal{X} $	number of elements in set \mathcal{X}

In this paper, we consider a set \mathcal{A} of *node-targeted attacks* aiming to shutdown a number of nodes. Each attack $a \in \mathcal{A}$ is identified with the set of nodes $\mathcal{V}(a) \subseteq \mathcal{V}$, of size $V(a)$. When attack a is conducted, all nodes v in $\mathcal{V}(a)$ are deleted from the network together with their sets of incident links (if a deleted node v is a controller node, the collocated controller is also deleted). So, the surviving network graph, denoted by $\mathcal{G}(a) = (\mathcal{W}(a), \mathcal{E}(a))$, is the maximal subgraph of \mathcal{G} induced by the set of surviving nodes $\mathcal{V}'(a) := \mathcal{V} \setminus \mathcal{V}(a)$, and hence $\mathcal{E}(a) := \mathcal{E} \cap \mathcal{W}(a)^{[2]}$. After attack a the set of surviving controllers is equal to $\mathcal{S}(a) := \mathcal{S} \cap \mathcal{V}'(a)$.

Note that as a consequence of attack a , the surviving graph $\mathcal{G}(a)$ is in general split into a set $\mathcal{C}(a)$ of (disjoint) connected components. Such a component c is identified with the set of its nodes $\mathcal{V}(c)$ (of size $V(c)$). The resilience impact of each attack $a \in \mathcal{A}$ is expressed by the *network availability measure* $\mathcal{M}(a)$ (referred to as *NA measure* or simply *NAM*):

$$\mathcal{M}(a) := \sum_{c \in \mathcal{C}(a)} f(c)Y(a, c). \quad (1)$$

In this definition, $Y(a, c)$ is a binary indicator equal to 1 when a given component $c \in \mathcal{C}(a)$ contains a controller node (i.e., when $\mathcal{V}(c) \cap \mathcal{S} \neq \emptyset$), and to 0 otherwise (in the former case, the component is named a *surviving component*). The quantity $f(c)$, in turn, is a *component metric* defined for the component (in fact, for the subsets of \mathcal{V}) assuming that an higher value of $\mathcal{M}(a)$ represents a lower (negative) impact of attack a . Here, we consider two particular metrics $f(c)$, both depending only on the size $V(c)$ of the component c :

$$f(c) := V(c) \quad f(c) := \binom{V(c)}{2}. \quad (2)$$

Using the left (*linear*) metric, the quantity $\mathcal{M}(a)$ measures the overall number of nodes in the surviving components, while using the right (*quadratic*) metric, $\mathcal{M}(a)$ measures the overall number of node-pairs in the surviving components.

Finally, for a given set of attacks \mathcal{A} , we define two NA measures, namely the *average* NAM and the *worst-case* NAM:

$$\mathcal{M}(\mathcal{A}) := \sum_{a \in \mathcal{A}} w(a)\mathcal{M}(a) \quad \mathcal{M}(\mathcal{A}) := \min_{a \in \mathcal{A}} \mathcal{M}(a). \quad (3)$$

In the left measure, $w(a)$, $a \in \mathcal{A}$, are given attack weights representing, for example, a given probability distribution characterizing the relative frequency of the considered attacks.

 TABLE II
 NOTATION USED IN OPTIMIZATION FORMULATIONS

P', P''	minimum and maximum number of PCs
$D(1), D(2)$	upper bounds on the SC and CC delay
$d(v, w)$	length (delay) of shortest paths between nodes v and w
$\mathcal{W}(v)$	set of nodes w with $d(v, w) \leq D(1)$
\mathcal{U}	set of node pairs $\{v, w\}$ with $d(v, w) > D(2)$ ($\mathcal{U} \subseteq \mathcal{V}^{[2]}$)
d	variable expressing the upper bound on the SC delays
y_v	binary variable equal to 1 iff node v contains a PC
z_{vw}	binary variable equal to 1 iff the controller serving switch v is placed at node w
$S(a, c)$	binary coefficient equal to 1 iff component $c \in \mathcal{A}$ contains a PC
\mathcal{P}	set of generated PC placements
$Y(v, p)$	binary coefficient equal to 1 iff node v belongs to placement $p \in \mathcal{P}$
u_p	binary variable equal to 1 iff placement $p \in \mathcal{P}$ is selected
P, B	number of PCs and BCs, respectively
x_v	binary variable equal to 1 iff node v contains a BC
Y_v	binary variable equal to 1 if node v contains a controller (primary or backup)
s_{ac}	binary variable equal to 1 iff component $c \in \mathcal{C}(a)$ contains a controller (primary or backup)
Z	variable expressing the value of objective function
\mathbb{B}, \mathbb{R}	sets of binary and real numbers

III. PRIMARY CONTROLLER PLACEMENT

This section presents two IP (integer programming) formulations that optimize placements of PCs taking into account SC and CC delays and the number of the installed controllers, i.e., solve the *PC placement problem* (PCPP). The notation used throughout this section (and beyond) is summarized in Table II.

A. Primary controller placement for min-max SC delay

We first consider the problem of finding a PC placement that minimizes the maximum SC delay while keeping the CC delay below the assumed upper bound $D(2)$. For that we introduce the following integer programming (IP) formulation.

PCPP/SC:

$$\min d \quad (4a)$$

$$P' \leq \sum_{v \in \mathcal{V}} y_v \leq P'' \quad (4b)$$

$$y_v + y_w \leq 1, \quad \{v, w\} \in \mathcal{U} \quad (4c)$$

$$\sum_{w \in \mathcal{W}(v)} z_{vw} = 1, \quad v \in \mathcal{V} \quad (4d)$$

$$z_{vw} \leq y_w, \quad v \in \mathcal{V}, w \in \mathcal{W}(v) \setminus \{v\} \quad (4e)$$

$$z_{vv} = y_v, \quad v \in \mathcal{V} \quad (4f)$$

$$d \geq \sum_{w \in \mathcal{V}} d(v, w)z_{vw}, \quad v \in \mathcal{V}. \quad (4g)$$

$$y_v \in \mathbb{B}, v \in \mathcal{V}; z_{vw} \in \mathbb{B}, v \in \mathcal{V}, w \in \mathcal{W}(v); d \in \mathbb{R}. \quad (4h)$$

Above, binary variables y_v determine the PC locations, i.e., $y_v = 1$ iff a PC is placed at node v . Thus, constraint (4b) assures that the number of located controllers is between the assumed values P' and P'' . Then, constraint (4c) (where \mathcal{U} is the set of all node-pairs whose distance, measured as the delay, exceeds the upper bound $D(2)$) enforces compliance with the CC delay requirements. The next group of constraints determine, using binary variables z_{vw} , the assignment of

nodes to controllers. Assuming that $z_{vw} = 1$ iff node v is assigned to the controller at node w , constraint (4d) assures that every node is assigned to exactly one controller from the set $\mathcal{W}(v)$ (i.e., the set of nodes for which the delay from node v is not greater than the upper bound $D(1)$), while constraint (4e) implies that if $z_{vw} = 1$ then node w must contain a controller. Constraint (4f), in turn, makes sure that the nodes equipped with controllers are assigned to themselves. Finally, constraint (4g) assures that the value of variable d is greater than or equal to the SC delay of each node. As this variable is minimized by objective (4a), its final value will minimize the maximum SC delay over the network nodes. \square

Below, the sets $\mathcal{W}(v)$ defined for $D(1) = d^*$, where d^* is the optimal value of d resulting from solving problem (4), are denoted by $\mathcal{W}^*(v)$, $v \in \mathcal{V}$.

B. Minimizing the number of primary controllers

After finding the min-max SC-delay d^* by means of formulation (4) and defining the sets $\mathcal{W}^*(v)$, $v \in \mathcal{V}$, we can now find the minimum number of PCs that assures the min-max SC delay $D(1) = d^*$ and the assumed CC delay $D(2)$. The appropriate IP formulation is as follows.

PCPP/NP:

$$\min \sum_{v \in \mathcal{V}} y_v \quad (5a)$$

$$\sum_{v \in \mathcal{V}} y_v \geq P' \quad (5b)$$

$$y_v + y_w \leq 1, \{v, w\} \in \mathcal{U} \quad (5c)$$

$$\sum_{w \in \mathcal{W}^*(v)} y_w \geq 1, v \in \mathcal{V}. \quad (5d)$$

$$y_v \in \mathbb{B}, v \in \mathcal{V}. \quad (5e)$$

Above, only constraint (5d) needs explanation: it simply forces that, for each $v \in \mathcal{V}$, at least one node in $\mathcal{W}^*(v)$ contains a controller. Note that $v \in \mathcal{W}^*(v)$, $v \in \mathcal{V}$. \square

Below, the resulting minimal value (5a) is denoted by P^* .

IV. JOINT OPTIMIZATION OF PRIMARY AND BACKUP CONTROLLER PLACEMENT

In this section we describe a two-phase approach to joint optimization of primary and backup controller placement aiming at maximizing network resilience to attacks from a given list. In the first phase we prepare a list of feasible PC placements and then, in the second phase, we select a placement from this list and find a corresponding BC placement that together maximize the resilience in question.

A. Phase 1: preparing lists of primary controller placements

Having found the min-max SC delay d^* (primary criterion, (4)) and the minimum number P^* of PCs that assure this min-max SC delay (secondary criterion, (5)), we are in a position to prepare a list of PC placements feasible with respect to these criteria (and the original CC delay bound). This is done by means of the following IP formulation.

PCPP/LPCP:

$$\min \frac{1}{V} \sum_{v \in \mathcal{V}} \sum_{w \in \mathcal{W}^*(v)} d(v, w) z_{vw} \quad (6a)$$

$$\sum_{v \in \mathcal{V}} y_v = P^* \quad (6b)$$

constraints (4c)-(4f)

$$\sum_{v \in \mathcal{S}(p)} y_v \leq P^* - 1, p \in \mathcal{P}, \quad (6c)$$

$$y_v \in \mathbb{B}, v \in \mathcal{V}; z_{vw} \in \mathbb{B}, v \in \mathcal{V}, w \in \mathcal{W}^*(v) \quad (6d)$$

where \mathcal{P} denotes the current list of PC placements, and $\mathcal{S}(p)$ denotes the set of nodes equipped with PCs in placement $p \in \mathcal{P}$. Hence, constraint (6c) forces the generated placement to differ from the placements on list \mathcal{P} by at least one node. The rest of the constraints, together with objective (6a), assure that the constructed solution is feasible (with respect to d^* , P^* and $D(2)$) and minimizes the average SC delay. \square

Formulation PCPP/LPCP is used iteratively, starting with an empty placement list \mathcal{P} . After each iteration, a newly found placement is added to the list and the formulation is resolved. The iterations are stopped when an assumed number of placements is generated or the formulation becomes infeasible. Note that the parameters d^* and P^* assumed in the above formulation could be slightly increased if doing so will result in a noticeable reduction in the average SC delay minimized by (6a).

B. Phase 2: primary/backup controller placement problem

In the second phase of the considered procedure we solve the joint primary and BC placement problem by means of one of the following IP formulations (which are extensions of the formulation presented in [5], [6]).

The first formulation maximizes the worst-case NA measure (3) for the assumed list of attacks \mathcal{A} .

PBCPP/WCNA:

$$\max Z \quad (7a)$$

$$\sum_{p \in \mathcal{P}} u_p = 1 \quad (7b)$$

$$\sum_{v \in \mathcal{V}} x_v = B \quad (7c)$$

$$Y_v = x_v + \sum_{p \in \mathcal{P}} Y(v, p) u_p, v \in \mathcal{V} \quad (7d)$$

$$s_{ac} \leq \sum_{v \in \mathcal{V}(c)} Y_v, a \in \mathcal{A}, c \in \mathcal{C}(a) \quad (7e)$$

$$Z \leq \sum_{c \in \mathcal{C}(a)} f(V(c)) s_{ac}, a \in \mathcal{A} \quad (7f)$$

$$u_p \in \mathbb{B}, p \in \mathcal{P}; x_v, Y_v \in \mathbb{B}, v \in \mathcal{V} \quad (7g)$$

$$s_{ac} \in \mathbb{B}, a \in \mathcal{A}, c \in \mathcal{C}(a); Z \in \mathbb{R}. \quad (7h)$$

In the formulation, binary variable u_p is equal to 1 iff placement $p \in \mathcal{P}$ is selected, and constraint (7b) assures that exactly one placement is actually selected. Constraint (7c), in turn, assures that exactly B BCs are deployed. Next, constraint (7d), where each $Y(v, p)$ is a binary coefficient equal to 1 iff placement p contains node v (i.e., iff $v \in \mathcal{S}(p)$), sets the value of binary variable Y_v to 1 iff node v contains a BC or a PC from the selected placement p with $u_p = 1$ (and assures that at most one controller is placed in v). Constraint (7e) forces the value of binary variable s_{ac} to 0 when component $c \in \mathcal{C}(a)$ does not contain any controller, and constraint (7f) assures that the value of variable Z is greater than or equal to value of the NA measure for each attack a . Finally, the optimization goal is achieved because Z is maximized by objective (7a). \square

Note that when the list \mathcal{P} contains just one element, then PBCPP/NA reduces to finding a BC placement maximally enhancing a given placement of PCs. In this case the binary variables u_p can be skipped.

The second formulation, a modification of PBCPP/WCNA, maximizes the average NA measure (3).

PBCPP/ANA:

$$\max \sum_{a \in \mathcal{A}} w(a) \sum_{c \in \mathcal{C}(a)} f(V(c)) s_{ac} \quad (8a)$$

constraints (7b)-(7e), (7g)

$$s_{ac} \in \mathbb{B}, a \in \mathcal{A}, c \in \mathcal{C}(a). \quad (8b)$$

V. NUMERICAL STUDY

Below we discuss numerical results obtained for the network *cost266* described in SNDlib [7] consisting of $V = 37$ nodes and $E = 57$ links, and depicted in Figure 1. In the following, the delay between a pair of nodes is assumed to be proportional to the length (in km) of their shortest path.

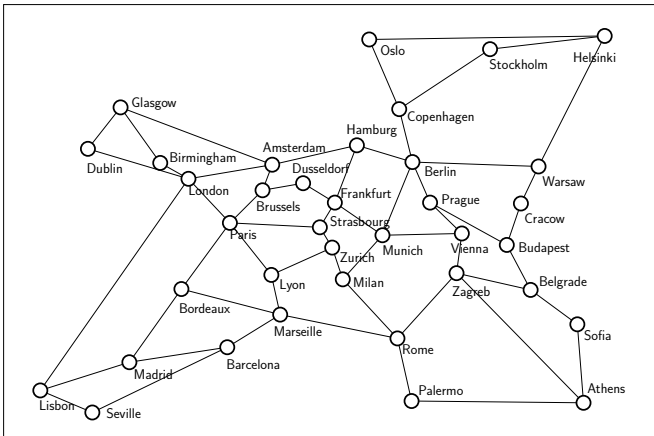


Fig. 1. An example network topology.

Table III characterizes PC placements for two CC delay bounds from column 1. Column 2 shows the min-max SC delay d^* optimized via PCPP/SC (formulation (4) with $P' = 2$, $P'' = 8$), achieved for the number of controllers P and the average SC delay, given in columns 3 and 4, respectively. Columns 5-6 show the results of PCPP/NP (formulation (5)): the minimum number of PCs P^* needed to meet the SC delay bound d^* (column 5), and the resulting average SC delay (column 6). As expected, PCPP/NP decreases the number of controllers (from $P = 8$ to $P^* = 3$ for CC delay 1500, and from $P = 8$ to $P^* = 5$ for CC delay 2000). The last two columns show the results from PCPP/LPCP (formulation (6)) for fixed d^* and P^* . The generated lists consist of 5 and 8 placements (column 7), and contain all feasible solutions of PCPP/LPCP. The minimized average SC delays of the first (the best in this aspect) and of the last (the worst) placement in each list are shown in column 8. Even in the worst cases the average SC delays are substantially smaller than those shown in column 6.

Note also that the admissible CC delay equal to 2000 allows the PCs to be further apart in the network as compared to CC delay 1500. Hence, as seen in Table III, the larger CC delay bound reduces the maximum (column 2) and the average (column 8) SC delays, but at the expense of requiring more PCs (column 5).

In all variants of controller placements considered below we used a list of $|\mathcal{A}| = 12$ worst-case node attacks, each targeting a set of $V(a) = 6$ nodes. To compute the first (most-dangerous) attack $a = 1$, we solve the *critical node detection* (CND) problem that finds a set of A nodes (in our case $A = 6$) whose removal minimizes the number of node pairs that can communicate in the surviving network (we used one of the CND IP formulations described in [8]). Then, the IP formulation is extended by a constraint that eliminates the current list of the sets of critical nodes from the feasible solution set of the problem (this constraint is analogous to (6c)) and solved to compute the nodes of the next attack. This process is repeated iteratively until 12 attacks are generated. Table IV presents the sets of nodes targeted by the attacks together with the number $\sum_{c \in \mathcal{C}(a)} \binom{V(c)}{2}$ of node pairs that are able to communicate after each attack $a = 1, 2, \dots, 12$.

Tables V and VI show the results of joint controller placement optimization for the above described attack list \mathcal{A} . All four cases of network availability measures (NAMs) implied by definitions (1)-(3) are considered (see column 2 in Tables V-VI). These are: the average (“avg”) or the worst-case (“w-c”) NAM with either linear (“L”) or quadratic (“Q”) metric. (The weights $w(a)$ used in the average NAM were uniformly set to 1.)

We solved PBCPP for two settings indicated in column 1 of Tables V-VI: “S” (single) where the list \mathcal{S} contains only one element (i.e., the placement with the minimum average SC delay obtained in the first iteration of PCPP/LPCP), and “F” (full) with all feasible placements. In both tables, the columns, starting with the third, show the optimized values of NA measures for the increasing number B of the deployed BCs.

Before analyzing the results, let us note that 31 is the upper bound on the optimal values of both linear NAMs (avg/L and w-c/L) since 31 is the number nodes remaining after removing the 6 attacked nodes. For the worst case quadratic NAM (w-c/Q) the upper bound is 124 since, as seen in Table IV, as a result of the most dangerous attack ($a = 1$) there are in total 124 node pairs in the surviving components. Finally, for the average quadratic NAM (avg/Q), its upper bound is 134.92 – the average of the 12 values in the second column of Table IV. Accordingly, Tables V-VI present the results up to the values of B for which these upper bounds are reached.

Looking at the third column (“ $B = 0$ ”) and the four rows marked with “S” in Table V, we notice that the levels of protection against the considered attacks provided by the optimal PC placement without BCs are substantially lower than the upper bounds (31,31,134.92, and 124, respectively). Yet, as shown in the remaining columns, adding BCs gradually improves these levels to finally reach the upper bounds in

TABLE III
RESULTS FOR THE PRIMARY CONTROLLER PLACEMENTS

1	2	3	4	5	6	7	8
CC delay bound	min-max SC delay (d^*)	P (4)	avg. SC delay (4)	P^*	avg. SC delay (5)	list size ($ \mathcal{P} $)	avg. SC delay (6)
1500	1529	8	788.0	3	810.8	5	656.3 – 727.5
2000	1168	8	446.7	5	708.5	8	517.9 – 559.9

TABLE IV
THE 12 WORST-CASE NODE-TARGETED ATTACKS WITH $|\mathcal{V}(a)| = 6$

a	$\sum_c \binom{V(c)}{2}$	$\mathcal{V}(a)$					
1	124	Berl	Paris	Buda	Marseil	Frankf	London
2	127	Berl	Paris	Buda	Lisbon	Marseil	Frankf
3	130	Berl	Paris	Buda	Marseil	Frankf	Amst
4	132	Berl	Paris	Milan	Buda	Marseil	Frankf
5	132	Berl	Crac	Paris	Marseil	Frankf	London
6	135	Berl	Crac	Paris	Lisbon	Marseil	Frankf
7	136	Berl	Crac	Paris	Milan	Marseil	Frankf
8	138	Berl	Crac	Paris	Marseil	Frankf	Amst
9	138	Berl	Paris	Buda	Zurich	Marseil	Frankf
10	142	Berl	Paris	Wars	Milan	Marseil	Frankf
11	142	Berl	Paris	Wars	Marseil	Frankf	London
12	143	Berl	Crac	Paris	Zurich	Marseil	Frankf

TABLE V
OPTIMAL PLACEMENTS (CC DELAY BOUND $D(2) = 1500$, MAX SC DELAY $D(1) = 1529$, NUMBER OF PRIMARY CONTROLLERS $P = 3$)

	NAM	$B=0$	$B=1$	$B=2$	$B=3$	$B=4$	$B=5$
S	avg/L	20.92	26.17	29.58	30.50	30.83	31.00
S	w-c/L	14.00	22.00	28.00	29.00	30.00	31.00
S	avg/Q	113.08	124.50	134.00	134.75	134.92	
S	w-c/Q	78.00	109.00	124.00			
F	avg/L	24.33	29.58	30.50	30.83	31.00	
F	w-c/L	22.00	28.00	29.00	30.00	31.00	
F	avg/Q	122.58	134.00	134.75	134.92		
F	w-c/Q	109.00	124.00				

question (the consecutive bounds are reached when 5, 5, 4 and 2, respectively, controllers are added). When the full list of 5 PC placements is considered (the four rows marked with “F”), the optimized combination of primary and backup controllers improves the protection levels for all values of B , and the upper bounds are reached using a smaller number of BCs as compared to the single PC placement case “S”.

TABLE VI
OPTIMAL PLACEMENTS (CC DELAY BOUND $D(2) = 2000$, MAX SC DELAY $D(1) = 1168$, NUMBER OF PRIMARY CONTROLLERS $P = 5$)

	NAM	$B=0$	$B=1$	$B=2$	$B=3$	$B=4$
S	avg/L	29.42	30.33	30.67	30.83	31.00
S	w-c/L	28.00	28.00	30.00	30.00	31.00
S	avg/Q	134.00	134.75	134.92		
S	w-c/Q	124.00				
F	avg/L	29.75	30.67	30.83	31.00	
F	w-c/L	28.00	30.00	30.00	31.00	
F	avg/Q	134.17	134.92			
F	w-c/Q	124.00				

Analogous observations are valid for the second CC delay case illustrated in Table VI. In this case, however, the number of BCs needed to reach the upper bounds of different NAMs is smaller than for the case considered in Table V. This is because now the number of PCs (equal to 5) is greater than

before (equal to 3).

All the considered optimization problems solve very quickly. Using AMPL/CPLEX 20.1 and a Hewlett-Packard HP DL380 G9 server (Xeon 10C processors) with access to 6 logical processors and 64 GB RAM, all the problem instances reported in Tables V-VI were solved to optimality in less than several tens of milliseconds. (For the cases with $|\mathcal{V}(a)| > 6$ the running times never exceeded 1 second.) As far the size of the problem is concerned, for the instances examined in the upper part of Table V there are 138 variables and 103 constraints. Finally, we note that for a network with 75 nodes and 99 links that we also examined, analogous values were 210 and 150 while the running times never exceed 5 seconds.

VI. CONCLUDING REMARKS

In this paper, we proposed a set of optimization models to solve the primary and backup controller placement problem where some controllers are only deployed as BCs. For a given CC delay bound, we first find a PC placement that minimizes the maximum SC delay and find the minimum number of PCs needed to reach that min-max SC delay. Next, we generate a list of feasible PC placements and find the particular placement that, together with a given number of BCs, maximizes resilience (evaluated for four different variants of network availability measures) of the SDN against a given list of node-targeted attacks. Finally, we presented a numerical study illustrating the merits of the proposed methodology for a well-known network topology.

REFERENCES

- [1] T. Hu, Z. Guo, P. Yi, T. Baker, and J. Lan, “Multi-controller based software-defined networking: A survey,” *IEEE Access*, vol. 6, pp. 15980–15996, 2018.
- [2] S. S. Savas, M. Tornatore, F. Dikbiyik, A. Yayimli, C. U. Martel, and B. Mukherjee, “RASCAR: Recovery-aware switch-controller assignment and routing in SDN,” *IEEE Trans. on Network and Service Management*, vol. 15, no. 4, pp. 1222–1234, 2018.
- [3] S. Yang, L. Cui, Z. Chen, and W. Xiao, “An efficient approach to robust SDN controller placement for security,” *IEEE Trans. on Network and Service Management*, vol. 17, no. 3, pp. 1669–1682, 2020.
- [4] D. Santos, A. de Sousa, C. Mas-Machuca, and J. Rak, “Assessment of Connectivity-Based Resilience to Attacks Against Multiple Nodes in SDNs,” *IEEE Access*, vol. 9, pp. 58 266–58 286, 2021.
- [5] E. Calle, D. Martínez, M. Mycek, and M. Pióro, “Resilient backup controller placement in distributed SDN under critical targeted attacks,” *Int. J. of Critical Infrastructure Protection*, vol. 33, pp. 12–260, 2021.
- [6] M. Pióro, M. Mycek, and A. Tomaszewski, “Network protection against node attacks based on probabilistic availability measures,” *IEEE Trans. on Network and Service Management*, vol. Early Access, pp. 1–22, 2021.
- [7] S. Orłowski, M. Pióro, A. Tomaszewski, and R. Wessály, “SNDlib 1.0 – Survivable network design library,” *Networks: An International Journal*, vol. 55, pp. 276–286, 2010.
- [8] D. Santos, A. de Sousa, and P. Monteiro, “Compact Models for Critical Node Detection in Telecommunication Networks,” *Electronic Notes in Discrete Mathematics*, vol. 64, pp. 325–334, 2018.