



A k -Cover Model for Reliability-Aware Controller Placement in Software-Defined Networks

Gabriela Schütz^{1,2} 

¹ Institute of Engineering of the University of Algarve,
Campus da Penha, Faro, Portugal
gschutz@ualg.pt

² Center for Electronic, Optoelectronic and Telecommunications (CEOT),
Campus de Gambelas, Faro, Portugal

Abstract. The main characteristics of Software-Defined Networks are the separation of the control and data planes, as well as a logically centralized control plane. This emerging network architecture simplifies the data forwarding and allows managing the network in a flexible way. Controllers play a key role in SDNs since they manage the whole network. It is crucial to determine the minimum number of controllers and where they should be placed to provide low latencies between switches and their assigned controller. It is worth to underline that, if there are long propagation delays between controllers and switches, their ability of reacting to network events quickly is affected, degrading reliability. Thus, the Reliability-Aware Controller Placement (RCP) problem in Software-Defined Networks (SDNs) is a critical issue. In this work we propose a k -cover based model for the RCP problem in SDNs. It simultaneously optimizes the number and placement of controllers, as well as latencies of primary and backup paths between switches and controllers, providing reliable networks against link, switch and controller failures. Although RCP problem is NP-hard, the simulation results show that reliabilities greater than 97%, satisfying low latencies, were obtained and the model can be used to find the optimum solution for different network topologies, in negligible time.

Keywords: Software-defined network · Controller placement · Reliability · k -cover problem

1 Introduction

Software-Defined Network (SDN) decouples control and data planes simplifying the data forwarding and allowing the network management in a flexible way. The SDN control plane is crucial to the network performance [9]. It handles

Supported by the UID/MULTI/00631/2013 project of the Portuguese Science and Technology Foundation (FCT).

© Springer Nature Switzerland AG 2019

J. M. F. Rodrigues et al. (Eds.): ICCS 2019, LNCS 11536, pp. 604–613, 2019.

https://doi.org/10.1007/978-3-030-22734-0_44

state distribution, control applications and network connectivity for propagating events to switches and also between multiple controllers. Network failures that disconnect the control and data planes could block requesting instructions from switches to controllers and may cause packet loss and network unsatisfactory performance [4]. The optimal number and placement of controllers, as well as the assignment of controllers to switches, play a very important role towards performance and reliability of SDN [2].

Although a switch can detect a control path failure, it has no capacity to establish a new route and connection will be lost until a backup control path is found. The distance between switches and their assigned controller affects propagation latency and restoration time, so low latency paths are required. Providing in advance backup control paths with acceptable latency allows quick restoration of the control plane against path failure, since a switch can initiate its backup path as soon as it detects a control path failure [7]. Similarly, planning in advance, for each switch, low latency connections to two different controllers over two disjoint paths allows quick restoration of the control plane against controller failure or congestion.

The above reasoning grounds our approach. In this work we propose a mathematical model based on the k -cover problem to plan a reliable SDN, enhancing the protection of the control plane against link, switch, and controller failures. It determines the optimum number of controllers and their placements, constrained to: (i) every switch must be connected to two different controllers, a primary and a backup controller, over two disjoint control paths; (ii) every switch must be connected to its assigned primary controller over two disjoint paths; (iii) control paths (primary and backup) latencies must be below a given threshold.

The remainder of this article is organised as follows: Sect. 2 is a short review of related works. The proposed model is described in Sect. 3, which includes the mathematical formalization. The experimental simulation is described in Sect. 4, while Sect. 5 presents and analyses the results. Section 6 draws some conclusions from the obtained results.

2 Related Work

Reliability-Aware Controller Placement (RCP) is a particular case of the Control Placement Problem (CPP). Several works have already addressed different issues related to the CPP. This section briefly overviews some works on fault tolerant and reliable controller placement towards the improvement of network resilience. Network elements failures may cause the disconnection between controllers and switches. Zhang et al. [8] call lost nodes to these switches that are unable to connect the controller due to failures. They minimize the number of lost nodes using a min-cut based controller placement algorithm to obtain a partition of the network, such that inside each partition switches and the respective controller are well connected. Hock et al. [3] define performance and resilience metrics in the controller placement problem and implement a framework to evaluate the entire solution space. In [4] the expected percentage of control path loss, defined as the

number of broken control paths due to network failures, is used to characterize the SDN reliability and for a given number of controllers they maximize the SDN reliability through a binary integer programming. Muller et al. [5] formulate the problem as a binary integer programming to maximize the average number of disjoint paths between switches and controllers. They also propose heuristics for defining lists of backup controllers to deal with controller failure. Ros and Ruiz [6] develop a heuristic for the fault tolerant controller placement problem where reliability thresholds must be satisfied. Their results show that if each node connects to two or three controllers, it can provide more than five nines reliability and also that, generally, ten controllers are enough, being its number more related to the network topology than to the network size. Vizarreta et al. [7] present two controller placement strategies for a resilient control plane. One strategy considers that switches have to be connected to a controller over two disjoint paths and the other considers that switches have to be connected to two different controllers over two disjoint paths. They evaluate their two approaches in comparison to the unprotected case.

In this work the planning of primary and backup control paths in advance, as in [7], is also considered, but this approach is different because it determines the minimum number of controllers that ensure, simultaneously, disjoint primary and backup control paths providing required latencies and also primary and backup controllers for each switch, as mentioned in Sect. 1, modelled and formalized as a k -Cover Problem.

3 Problem Formalization

3.1 Problem Overview

When deploying multiple controllers, the reliability and resilience of SDN reside on a controller placement highly fault tolerant. Clearly, more controllers can increase the control network reliability, but also imply on more communications to exchange information, harder network management and overall cost increase [6]. It is advisable to place as few controllers as possible, taking into consideration that too few controllers would increase latency and decrease reliability. Thus the main goal is to find the appropriate number and locations of controllers to ensure control plane reliability and satisfy low propagation delay between switches and their assigned controllers.

This approach achieves the above goal. It minimizes the number of controllers ensuring the existence of at least two disjoint paths, primary and backup, between each switch and one controller and the assignment of two controllers, primary and backup, to each switch. All paths between switches and controllers satisfy the required latency. It is assumed that each switch communicates with the primary controller over the primary control path. If the primary controller fails than communication is quickly restored to the backup controller over a disjoint path, meaning that disconnection is avoided. If the primary control path to the primary controller fails (due to a link or switch failure) then the disjoint backup path to that controller is promptly initiated. Thus, for each switch, we

compute its control paths reliability, denoted by (R_s) , as the probability of no communication disconnection between switch and controller, as follows:

$$R_s = r_s \cdot [r_{C_p} \cdot (1 - (1 - r_p)(1 - r_b)) + (1 - r_{C_p}) \cdot r_{C_b} \cdot r_{dp}] \quad (1)$$

where, for simplicity, failure probability of a path, a link, a switch or a controller is denoted by f_* , where $*$ is equal to l for link; s for switch; C_p and C_b for primary and backup controllers; p and b for primary and backup disjoint paths to primary controller; dp for disjoint backup path to backup controller, being the reliability of a component $r_* = 1 - f_*$. The failure probability of a path from i to j is computed as

$$f_{p(i,j)} = 1 - \prod_{l \in p(i,j)} (1 - f_l) \prod_{s \in p(i,j) - \{i,j\}} (1 - f_s). \quad (2)$$

Average network reliability, denoted by R , is calculated as the average of the switches control paths reliabilities,

$$R = \frac{\sum_s R_s}{\text{number of switches}} \quad (3)$$

Obviously, there are two disjoint paths between switches and a controller if the degree of every node is equal or greater than two and the network has no articulation points. Next, we introduce some cover definitions applied to this problem.

Definition 1: A switch is covered by a controller if the path between them provides the required latency.

Definition 2: A switch is k -covered if it is covered by at least k different controllers.

Definition 3: A network is k -covered if every switch is k -covered and k is the degree of the coverage.

It is clear that the number of controllers to achieve a k coverage degree increases directly with k . Hence, the bigger is k the more controllers are needed. The network must be at least 2-covered to assure connectivity in the presence of a controller failure.

3.2 Mathematical Formalization

In the following mathematical formalization, the network is represented as an undirected graph $G(V, E)$, where $V = \{1, 2, \dots, N\}$ is the set of nodes (switches) and E is the set of edges (bidirectional links) connecting nodes. $V_c \subseteq V$ denotes the subset of switches ($v \in V$) hosting a controller. We assume a uniform demand and equal amount of traffic forwarded between switches and controllers. Since the propagation latency is the largest part of latency and the length of a communication link is proportional to the propagation delay it introduces, we assume that path length is equivalent to path latency. The latency of a primary control path between controller i and switch j is the length of the shortest path between them and is denoted by d_{ij}^p , the respective disjoint backup control path latency,

to the primary controller or to the backup controller, denoted by d_{ij}^b or d_{ij}^{dp} is the length of the shortest path between controller i and switch j in the sub-graph obtained by removing the links and intermediate nodes of the primary control path. Δ_p and Δ_b are, respectively, primary and backup control paths latency threshold.

We define, below, constants $a_{ij}, \forall i, j \in V$ used to ensure that j can be covered by a controller placed in i only if the two disjoint shortest paths between i and j satisfy the respective required latencies.

$$\forall i, j \in V, a_{ij} = \begin{cases} 1, & \text{if } d_{ij}^p \leq \Delta_p \wedge d_{ij}^b \leq \Delta_b \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

The coefficients of the objective function are equal to the average of the weighted sum of the distances of primary and backup paths, calculated as follows:

$$f_i = \frac{\sum_{j \in V} (\alpha d_{ij}^p + \beta d_{ij}^b) a_{ij}}{\sum_{j \in V} a_{ij}}, \forall i \in V \tag{5}$$

The binary decision variables are:

$$x_i = \begin{cases} 1, & \text{if the location of switch } i \in V \text{ is chosen to place a controller} \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

The RCP is formalized as a 2-Cover Problem, as follows:

$$\min \sum_{i \in V} f_i x_i. \tag{7}$$

subject to:

$$\sum_{i \in V} a_{ij} x_i \geq 2, \forall j \in V \tag{8}$$

$$x_i \in \{0, 1\}, \forall i \in V. \tag{9}$$

The objective function (7) minimizes the number of controllers weighted by the average of primary and backup path lengths. Constraints (8) ensure that every switch is, at least, covered by 2 different controllers, using disjoint paths, both providing feasible latencies. Constraints (9) define variables as binary.

The optimum solution of this formalization obtains the minimum number of controllers and their locations, such that each switch is connected, at least, to two controllers by two disjoint paths towards each controller, complying with the required latencies. The assignment of controllers to switches is implicit, since for each switch the primary controller is the nearest one; the backup controller is the nearest controller that can be connected over a disjoint path. Primary and backup paths between a switch and a controller will be disjoint by construction of $a_{ij}, \forall i, j \in V$, given in (4). Therefore, considering the obtained set of controllers $V_c = \{i \in V : x_i = 1\}$, we define the binary assignment variables of primary and backup controllers to each switch, as follows: $\forall j \in V, y_{ij}^{Cp} = 1$, if $arg(\min_{i \in V_c} d_{ij}^p) = i$ and $y_{ij}^{Cb} = 1$, if $arg(\min_{i \in V_c} d_{ij}^{dp}) = i$.

4 Experimental Setup

4.1 Network Topologies

This approach was tested in all topologies with lower node degree greater than or equal to 2, available online in SNDlib database [1] (so, abilene, brain, ta2 and zib54 topologies were excluded). Euclidean distances were computed and associated to the links. We define, as usually, the *Diameter* of a network as the maximum shortest path between any two nodes in the network. Table 1 summarizes networks parameters. Five small networks (dfn-bwin, dfn-gwin, di-yuan, geant and nobel-us) were not included in Table 1 because no feasible solution exists for the 2-cover problem given in (7)–(9). In fact, for at least one switch, j , $\sum_{i \in V} a_{ij} < 2$ thus constraints (8) can not be satisfied.

Table 1. Network parameters

Ref.	Topology	N	$ E $	Average degree	Diameter	
					Edges	Length (km)
1	pdh	11	34	6.18	3	670.36
2	polska	12	18	3.00	4	746.15
3	atlanta	15	22	2.93	5	612.79
4	newyork	16	49	6.12	3	522.75
5	nobel-germany	17	26	3.06	6	1893.81
6	ta1	24	55	4.58	4	1012.42
7	france	25	45	3.60	5	1215
8	janos-us	26	84	6.46	8	2172.61
9	norway	27	51	3.78	7	1255.56
10	sun	27	102	7.56	7	1068.92
11	nobel-eu	28	41	2.93	8	1807.7
12	india35	35	80	4.57	7	1582.08
13	cost266	37	57	3.08	8	2039.76
14	janos-us-ca	39	122	6.26	10	2776.39
15	giul39	39	172	8.82	6	1079.67
16	pioro40	40	89	4.45	7	1774.95
17	germany50	50	88	3.52	9	2583.56

4.2 Parameters

Propagation latency depends on the shortest paths and network topology, so we have considered for latency limits, in Eq. (4), $\Delta_p = 0.5 \times \text{diameter}$ and $\Delta_b = 0.6 \times \text{diameter}$. Without loss of generality, we assumed equal weight for primary and backup paths, so we used $\alpha = \beta = 0.5$ in Eq. (7).

5 Results

The optimum solution was obtained by exactly solving the 2-Cover Problem formalized above, (7)–(9), and afterwards, as described in Sect. 3.2, a procedure to assign primary and backup controllers to switches was applied. The algorithm was implemented in Matlab, taking an average execution time of 1.019 s, and varying between 0.217 s and 2.69 s, on an Intel i5-3210M CPU. The results obtained are presented and analysed below, concerning the optimum number of controllers, reliability and latency.

5.1 Number of Controllers

The optimum number of controllers obtained by the 2-Cover Problem, is shown in Table 2. Four controllers are enough in 58.8% of the networks. It is worth to underline that this number of controllers ensure, for each switch, two disjoint paths for two different controllers and also two disjoint paths for the primary controller, being all paths within the required latency. It is noted that the number of controllers depends greatly on the topology, mainly to cope with the requirement of connect two controllers to each switch over two disjoint control paths. The main difficulty occurs with topologies where two adjacent switches have degree 2 and both are adjacent to another switch, forming a triangle, because the latter is an articulation point and therefore creates a bottleneck, implying that is necessary to place a controller in one of the two-degree switches, to cover only these two switches. For instance, networks 6 to 11 have a similar number of switches but only network 7, which presents two bottlenecks (three switches forming a triangle twice), needed 6 controllers. Results also show that there is not a linear relation between the minimum number of controllers and the number of switches or the link density or the diameter. However, there is a tendency for networks with more switches, smaller link density and larger diameter to require fewer controllers (networks 13 to 17) and vice versa (networks 1 to 4).

Table 2. Number of controllers

Network Ref.	5, 6, 8, 9, 11, 13, 14, 15, 16, 17	1, 4, 10, 12	2, 3, 7
$ V_c $	4	5	6

5.2 Reliability

For each network, we compute the average network reliability, R , given in (3). It was considered, as in [7], the same failure probability for all the switches (including those hosting a controller) and the same failure rate per length of the link. We defined switch failure probabilities of 0.5%, 1% and 2% and link failure probabilities of 0.1% and 0.5% per 100 km. Figure 1 plots R for these 6 scenarios.

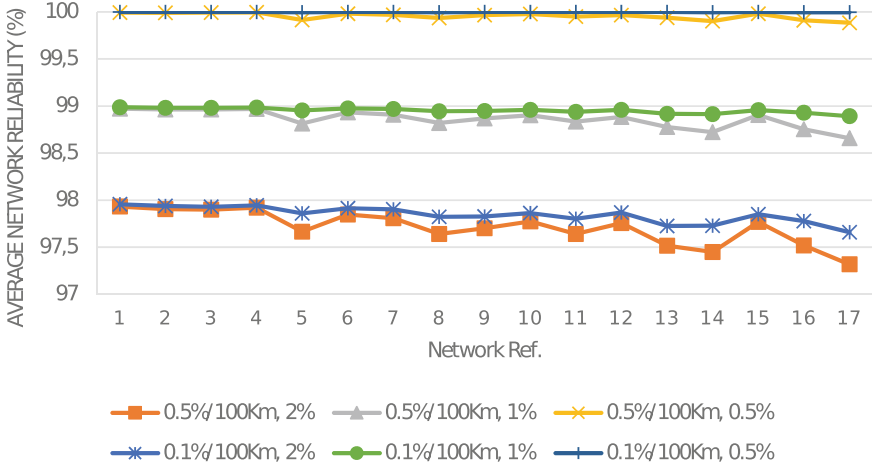


Fig. 1. Average network reliability, considering 6 scenarios combining 2 probabilities for link failure per 100 km with 3 probabilities for switch failure (f_l, f_s).

We have obtained R values ranging from 97.32% to 99.99% being all networks average equal to 98.89%. Therefore our approach ensures high average reliability. Figure 1 shows that, for each network, R presents almost the same values when only link failure probability varies. For the same link failure probability, R decreases around 1% with each switch failure probability increment. So, we conclude that switches (including controller) failure probabilities have greater impact on the network reliability than link failure probabilities.

The Average Control Path Availability used in [7] is the equivalent to R for their strategies. They only plot results for network 13, considering link failure probability of 0.1% per 100 km. Our results for that network outperformed theirs since we obtained reliabilities equal to 97.73%, 98.92% and 99.998% for node failure probabilities of 2%, 1% and 0.5%, respectively.

5.3 Propagation Latencies

As stated in Sect. 3.2 we have considered that the propagation latency is measured by the control path length. The average control path length to: primary controller over primary path ($L_p^{C_p}$) and over backup path ($(L_b^{C_p})$); to backup controller ($(L_{dp}^{C_b})$), over a disjoint backup path with length d^{dp} , are computed as follows:

$$L_p^{C_p} = \frac{1}{N} \sum_{i \in V_c} \sum_{j \in V} d_{ij}^p y_{ij}^{C_p} \tag{10}$$

$$L_b^{C_p} = \frac{1}{N} \sum_{i \in V_c} \sum_{j \in V} d_{ij}^b y_{ij}^{C_p} \tag{11}$$

$$L_{dp}^{C_b} = \frac{1}{N} \sum_{i \in V_c} \sum_{j \in V} d_{ij}^{dp} y_{ij}^{C_b} \tag{12}$$

Figure 2 plots these average control path lengths and the diameter.




Fig. 2. Average control path length

As expected the primary path has the smallest average length. Only Networks 1 and 12 present a backup path to the primary control slightly lower than the backup path to the backup controller. For each network the average control path length ratios to diameter were computed and for the 3 control paths (p, b and dp) they range from 9.77% to 24.12%, 26.01% to 40.22% and 16.45% to 35.19%, respectively. Therefore we can state that planning in advance paths concerning the protection against failures can be obtained with low latencies for disjoint primary and backup paths.

Vizarreta et al. [7] plot average latency results considering 2 and 4 controllers for five selected SNDlib topologies, 2, 12, 13, 14, and 17, considering their 2 strategies as mentioned in Sect. 2. Our approach needed more than 4 controllers for networks 2 and 12 thus path length results are not comparable. Comparing average control paths lengths, considering 4 controllers, we can see that our results present lower values for networks 13 (321, 599 and 528 km) and 14 (388, 801 and 773 km) and higher values for network 17 (462, 818 and 741).

6 Conclusion

In this article we have presented a 2-cover based approach for the RCP problem in SDN. It was able to find the minimum number of controllers, their placement and the assignment of controllers to switches, satisfying low propagation latencies, below defined limits, while ensuring for each switch the assignment of two

controllers (primary and backup) and the existence of at least two disjoint paths between each switch and its assigned controllers. Thus, it is foreseen the quickly restoration of communication to the backup controller over a disjoint path, when primary controller fails and also a promptly initialization of a disjoint backup path when the primary path fails. Results show that the proposed approach is able to determine, in all tested topologies, a highly reliable controller placement with low latencies. The approach proved also to be computationally efficient and scalable, as its performance is independent of network dimension. Therefore, it can be used to efficiently solve the considered RCP problem, under the assumptions discussed in this article and it can be easily extended to consider different amounts of traffic between switches and controllers, capacity constraints and switches with different protection levels.

References

1. Snd-lib. <http://sndlib.zib.de>
2. Heller, B., Sherwood, R., McKeown, N.: The controller placement problem. In: Proceedings of the First Workshop on Hot Topics in Software Defined Networks, HotSDN 2012, pp. 7–12. ACM, New York (2012). <https://doi.org/10.1145/2342441.2342444>. <http://doi.acm.org/10.1145/2342441.2342444>
3. Hock, D., Hartmann, M., Gebert, S., Jarschel, M., Zinner, T., Tran-Gia, P.: Pareto-optimal resilient controller placement in SDN-based core networks. In: Proceedings of the 2013 25th International Teletraffic Congress (ITC), pp. 1–9, September 2013. <https://doi.org/10.1109/ITC.2013.6662939>
4. Hu, Y., Wang, W., Gong, X., Que, X., Cheng, S.: On reliability-optimized controller placement for software-defined networks. *China Commun.* **11**(2), 38–54 (2014). <https://doi.org/10.1109/CC.2014.6821736>
5. Müller, L.F., Oliveira, R.R., Luizelli, M.C., Gasparly, L.P., Barcellos, M.P.: Survivor: an enhanced controller placement strategy for improving SDN survivability. In: 2014 IEEE Global Communications Conference, pp. 1909–1915, December 2014. <https://doi.org/10.1109/GLOCOM.2014.7037087>
6. Ros, F.J., Ruiz, P.M.: Five nines of southbound reliability in software-defined networks. In: Proceedings of the Third Workshop on Hot Topics in Software Defined Networking, HotSDN 2014, pp. 31–36. ACM, New York (2014). <https://doi.org/10.1145/2620728.2620752>. <http://doi.acm.org/10.1145/2620728.2620752>
7. Vizarreta, P., Machuca, C.M., Kellerer, W.: Controller placement strategies for a resilient SDN control plane. In: 2016 8th International Workshop on Resilient Networks Design and Modeling (RNDM), pp. 253–259, September 2016. <https://doi.org/10.1109/RNDM.2016.7608295>
8. Zhang, Y., Beheshti, N., Tatipamula, M.: On resilience of split-architecture networks. In: 2011 IEEE Global Telecommunications Conference - GLOBECOM 2011, pp. 1–6, December 2011. <https://doi.org/10.1109/GLOCOM.2011.6134496>
9. Zhang, Y., Cui, L., Wang, W., Zhang, Y.: A survey on software defined networking with multiple controllers. *J. Netw. Comput. Appl.* **103**, 101–118 (2018). <https://doi.org/10.1016/j.jnca.2017.11.015>. <http://www.sciencedirect.com/science/article/pii/S1084804517303934>