Ana I. Pereira · Andrej Košir ·
Florbela P. Fernandes · Maria F. Pacheco ·
João P. Teixeira · Rui P. Lopes (Eds.)

Communications in Computer and Information Science 1754

Optimization, Learning Algorithms and Applications

Second International Conference, OL2A 2022
Póvoa de Varzim, Portugal, October 24–25, 2022
Proceedings

Springer

OL2A
INTERNATIONAL
CONFERENCE

Communications in Computer and Information Science 1754

Ana I. Pereira · Andrej Košir ·
Florbela P. Fernandes · Maria F. Pacheco ·
João P. Teixeira · Rui P. Lopes (Eds.)

Optimization, Learning Algorithms and Applications

Second International Conference, OL2A 2022
Póvoa de Varzim, Portugal, October 24–25, 2022
Proceedings

Springer

Editors

Ana I. Pereira
Instituto Politécnico de Bragança
Bragança, Portugal

Andrej Košir
University of Ljubljana
Ljubljana, Slovenia

Florbela P. Fernandes
Instituto Politécnico de Bragança
Bragança, Portugal

Maria F. Pacheco
Instituto Politécnico de Bragança
Bragança, Portugal

João P. Teixeira
Instituto Politécnico de Bragança
Bragança, Portugal

Rui P. Lopes
Instituto Politécnico de Bragança
Bragança, Portugal

ISSN 1865-0929 ISSN 1865-0937 (electronic)
Communications in Computer and Information Science
ISBN 978-3-031-23235-0 ISBN 978-3-031-23236-7 (eBook)
https://doi.org/10.1007/978-3-031-23236-7

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

This CCIS volume 1754 contains the refereed proceedings of the Second International Conference on Optimization, Learning Algorithms and Applications (OL2A 2022), a hybrid event held during October 24–25, 2022.

OL2A 2022 provided a space for the research community on optimization and learning to get together and share the latest developments, trends, and techniques, as well as to develop new paths and collaborations. The conference had more than three hundred participants in an online and face-to-face environment throughout two days, discussing topics associated with optimization and learning, such as state-of-the-art applications related to multi-objective optimization, optimization for machine learning, robotics, health informatics, data analysis, optimization and learning under uncertainty, and Industry 4.0.

Five special sessions were organized under the following topics: Trends in Engineering Education, Optimization in Control Systems Design, Measurements with the Internet of Things, Advances and Optimization in Cyber-Physical Systems, and Computer Vision Based on Learning Algorithms. The OL2A 2022 program included presentations of 56 accepted papers. All papers were carefully reviewed and selected from 145 submissions in an single-blind process. All the reviews were carefully carried out by a scientific committee of 102 qualified researchers from 21 countries, with each submission receiving at least 3 reviews.

We would like to thank everyone who helped to make OL2A 2022 a success and hope that you enjoy reading this volume.

October 2022

Ana I. Pereira
Andrej Košir
Florbela P. Fernandes
Maria F. Pacheco
João P. Teixeira
Rui P. Lopes

Organization

General Chairs

Ana I. Pereira	Polytechnic Institute of Bragança, Portugal
Andrej Košir	University of Ljubljana, Slovenia

Program Committee Chairs

Florbela P. Fernandes	Polytechnic Institute of Bragança, Portugal
Maria F. Pacheco	Polytechnic Institute of Bragança, Portugal
João P. Teixeira	Polytechnic Institute of Bragança, Portugal
Rui P. Lopes	Polytechnic Institute of Bragança, Portugal

Special Session Chairs

João P. Coelho	Polytechnic Institute of Bragança, Portugal
Luca Oneto	University of Genoa, Italy

Technology Chairs

Alexandre Douplik	Ryerson University, Canada
Paulo Alves	Polytechnic Institute of Bragança, Portugal

Local Organizing Chairs

José Lima	Polytechnic Institute of Bragança, Portugal

Program Committee

Ana I. Pereira	Polytechnic Institute of Bragança, Portugal
Abeer Alsadoon	Charles Sturt University, Australia
Ala' Khalifeh	German Jordanian University, Jordan
Alexandre Douplik	Ryerson University, Canada
Ana Maria A. C. Rocha	University of Minho, Portugal
Ana Paula Teixeira	University of Trás-os-Montes and Alto Douro, Portugal
André Pinz Borges	Federal University of Technology – Paraná, Brazil
André R. da Cruz	Federal Center for Technological Education of Minas Gerais, Brazil

Andrej Košir	University of Ljubljana, Slovenia
Arnaldo Cândido Júnior	Federal University of Technology – Paraná, Brazil
António J. Sánchez-Salmerón	Universitat Politécnica de Valencia, Spain
António Valente	Universidade de Trás-Os-Montes e Alto Douro, Portugal
Bilal Ahmad	University of Warwick, UK
Bruno Bispo	Federal University of Santa Catarina, Brazil
Carlos Henrique Alves	CEFET/RJ, Brazil
Carmen Galé	University of Zaragoza, Spain
B. Rajesh Kanna	Vellore Institute of Technology, India
Carolina Gil Marcelino	Federal University of Rio de Janeiro, Brazil
Christopher E. Izquierdo	University of Laguna, Spain
C. Sweetlin Hemalatha	Vellore Institute of Technology, India
Damir Vrančić	Jozef Stefan Institute, Slovenia
Daiva Petkeviciute	Kaunas University of Technology, Lithuania
Dhiah Abou-Tair	German Jordanian University, Jordan
Diego Brandão	CEFET/RJ, Brazil
Dimitris Glotsos	University of West Attica, Greece
Diamantino Silva Freitas	University of Porto, Portugal
Eduardo Vinicius Kuhn	Federal University of Technology – Paraná, Brazil
Elizabeth Fialho Wanner	Federal Center for Technological Education of Minas Gerais, Brazil
Elaine Mosconi	Université de Sherbrooke, Canada
Esteban Clua	Federal Fluminense University, Brazil
Eric Rogers	University of Southampton, UK
Felipe N. Martins	Hanze University of Applied Sciences, Netherlands
Florbela P. Fernandes	Polytechnic Institute of Bragança, Portugal
Florentino F. Riverola	University of Vigo, Spain
Gaukhar Muratova	Dulaty University, Kazakhstan
Gediminas Daukšys	Kauno Technikos Kolegija, Lithuania
Glaucia Maria Bressan	Federal University of Technology – Paraná, Brazil
Glotsos Dimitris	University of West Attica, Greece
Humberto Rocha	University of Coimbra, Portugal
J. Marcos Moreno Veja	University of Laguna, Spain
João Paulo Carmo	University of São Paulo, Brazil
João P. Teixeira	Polytechnic Institute of Bragança, Portugal
Jorge Igual	Universidad Politécnica de Valencia, Spain
José Boaventura-Cunha	University of Trás-os-Montes and Alto Douro, Portugal
José Lima	Polytechnic Institute of Bragança, Portugal

Joseane Pontes	Federal University of Technology – Ponta Grossa, Brazil
Juan Alberto G. Esteban	Universidad de Salamanca, Spain
Juan A. Méndez Pérez	University of Laguna, Spain
Juani Lopéz Redondo	University of Almeria, Spain
Julio Cesar Nievola	Pontifícia Universidade Católica do Paraná, Brazil
João Paulo Coelho	Polytechnic Institute of Bragança, Portugal
Jorge Ribeiro	Polytechnic Institute of Viana do Castelo, Portugal
José Ramos	NOVA University Lisbon, Portugal
Kristina Sutiene	Kaunas University of Technology, Lithuania
Lidia Sánchez	University of León, Spain
Lino Costa	University of Minho, Portugal
Luis A. De Santa-Eulalia	Université de Sherbrooke, Canada
Luís Coelho	Polytecnhic Institute of Porto, Portugal
Luca Oneto	University of Genoa, Italy
Luca Spalazzi	Marche Polytechnical University, Italy
Maria F. Pacheco	Polytechnic Institute of Bragança, Portugal
Manuel Castejón Limas	University of León, Spain
Marc Jungers	Université de Lorraine, France
Marco A. S. Teixeira	Universidade Tecnológica Federal do Paraná, Brazil
Maria do R. de Pinho	University of Porto, Portugal
Marco A. Wehrmeister	Federal University of Technology – Paraná, Brazil
Markus Vincze	TU Wien, Austria
Martin Hering-Bertram	Hochschule Bremen, Germany
Mikulas Huba	Slovak University of Technology in Bratislava, Slovakia
Michał Podpora	Opole University of Technology, Poland
Miguel Ángel Prada	University of León, Spain
Nicolae Cleju	Technical University of Iasi, Romania
Paulo Lopes dos Santos	University of Porto, Portugal
Paulo Alves	Polytechnic Institute of Bragança, Portugal
Paulo Leitão	Polytechnic Institute of Bragança, Portugal
Paulo Moura Oliveira	University of Trás-os-Montes and Alto Douro, Portugal
Pavel Pakshin	Nizhny Novgorod State Technical University, Russia
Pedro Luiz de P. Filho	Federal University – Paraná, Brazil
Pedro Miguel Rodrigues	Catholic University of Portugal, Portugal
Pedro Morais	Polytechnic Institute of Cávado e Ave, Portugal
Pedro Pinto	Polytechnic Institute of Viana do Castelo, Portugal

Roberto M. de Souza	Federal University of Technology – Paraná, Brazil
Rui P. Lopes	Polytechnic Institute of Bragança, Portugal
Sabrina Šuman	Polytechnic of Rijeka, Croatia
Sani Rutz da Silva	Federal University of Technology – Paraná, Brazil
Santiago Torres Álvarez	University of Laguna, Spain
Sara Paiva	Polytechnic Institute of Viana do Castelo, Portugal
Shridhar Devamane	Global Academy of Technology, India
Sofia Rodrigues	Polytechnic Institute of Viana do Castelo, Portugal
Sławomir Stępień	Poznan University of Technology, Poland
Teresa P. Perdicoulis	University of Trás-os-Montes and Alto Douro, Portugal
Toma Roncevic	University of Split, Croatia
Uta Bohnebeck	Hochschule Bremen, Germany
Valeriana Naranjo-Ornedo	Universidad Politécnica de Valencia, Spain
Vivian Cremer Kalempa	Universidade Estadual de Santa Catarina, Brazil
Vitor Duarte dos Santos	NOVA University Lisbon, Portugal
Wynand Alkema	Hanze University of Applied Sciences, Netherlands
Wojciech Paszke	University of Zielona Gora, Poland
Wojciech Giernacki	Poznan University of Technology, Poland
Wolfgang Kastner	TU Wien, Austria

Analysis of Topological Characteristics Impacting the Allocation of Controllers in an SD-WAN Network

Vicente Padrenuestro[3] and André Mendes[1,2,3](✉)

[1] Research Centre in Digitalization and Intelligent Robotics (CeDRI), Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
a.chaves@ipb.pt

[2] Laboratório para a Sustentabilidade e Tecnologia em Regiões de Montanha (SusTEC), Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal

[3] Department of Systems Engineering and Automatics, Universidade de Vigo, 36.310 Pontevedra, Spain

Abstract. Finding a viable and optimal solution to the Software-Defined Networks (SDN) controller allocation problem in an open SDN network is a challenging task. In this context, this work was developed to expose a network's real characteristics, which may impact the choice of positioning an SDN controller. Furthermore, the experiments with the Pareto-based Optimal Controller-placement (POCO) tool and with the ONOS and Floodlight controllers are shown, which served as a comparison to assist in decision-making regarding the best positioning of the controller within the network. In summary, the results showed that not only the static aspects of the network should be considered, but also the dynamic characteristics and aspects related to the model and the purpose for which the controllers are used, factors that can impact the positioning.

Keywords: SDN · Controller placement problem · OpenFlow Controllers

1 Introduction

In recent years, since open Software-Defined Networks (SDN) emerge, using a protocol like OpenFlow to control the behaviour of virtual and physical switches at the data plane level, has gained a lot of attention and is earning more strength with the current scenarios. Thus, flexibility, programmability and scalability are some of the benefits brought by SDN compared to traditional network infrastructure.

In this sense, an important issue during the implementation of an SDN is the controller positioning in the network, i.e., deciding where to position an overview of the network, which allows managing resources in a simple and effective way,

A. I. Pereira et al. (Eds.): OL2A 2022, CCIS 1754, pp. 505–519, 2022.
https://doi.org/10.1007/978-3-031-23236-7_35

considering the use of programmable capacity to obtain efficient responses, with minimal latency between the nodes and the controller, reaching the maximum throughput between them.

The first approach on the subject was made in the work [4], followed by other researchers who added different issues in their observations. However, most of the research in the SDN area about the controller placement problem (*Controller Placement Problem*, CPP) focuses most of its efforts on the analysis of networks considering only their static characteristics data [5], such as considering geographical distances and latencies related to these distances.

Moreover, the network characteristics and the individual load caused by different nodes connected to a controller are still little considered in detail. Also quantifying how many controllers are needed, so that there is not a limiting point, and which controllers have the best performance according to the placement position in each scenario, are important issues in order to minimise the cost of network overlay and ensure good performance.

Given this, the main contributions of this work are in investigating, in a comparative way, which topological characteristics can impact the positioning of controllers in real SD-WAN networks, through a study on the use of the exhaustive method implemented by the POCO (Pareto-based Optimal COntroller) tool [5], which consists of, in a summarised way, calculating the geographical distances in relation to all $\mathcal{P}$ positions of topology, to connect the k controllers in a network containing n nodes.

Thus, the results defined by the POCO tool are compared in the case study of the *Ipê network*, of RNP (Brazilian National Network for Education and Research), considering its natural characteristics and also the analysis of ONOS and Floodlight controllers in a virtualised environment.

As a result, it is shown that in the studies about the controllers positioning, only the considerations about geographical distances and latencies defined by the POCO tool are not enough, therefore, the tests proved that the latency characteristics together with the bandwidth and traffic information, are also important factors in the positioning definition, in addition, the tests also showed that the aspects related to the model and the use of the controllers are important.

After this brief introduction, with thoughts on the theoretical background and motivation, the paper is organized as follows. Section 2 provides the related works, basic concepts and problems encountered when dealing with the placement controller problem. A methodology used and applied tools are presented in Sect. 3. Simulation and results are presented and discussed in Sect. 4, and finally, Sect. 5 concludes the paper and lists future works.

2 Related Work

Several studies have been carried out aiming to address the problem of controller positioning, other studies have worked towards comparing SDN controllers in an isolated way, focusing only on the controller performance. However, few studies have helped in decision-making to select a controller with the best performance, taking into account its proper position within a network with real characteristics.

The controller positioning problem in SDN architecture was first introduced by [4], where optimisation was performed in relation to the latency of the nodes up to the designated controller. In this paper, the author cites that for WANs, the best positioning depends, among other metrics, on latency.

In the field of the positioning problem, [4] shows the network performance by varying the position of controllers in the network. Finding the location and number of controllers that will be needed is a challenging task in a network architecture like SD-WAN.

In [6], POCO (*Pareto-based Optimal COntroller placement*) [5], a framework for finding the optimal controller placement so that the connectivity between the switches and the controller is maximized taking into account the capacity of the controller, is presented. First, the authors propose a brute-force algorithm, but it is valid only for small networks. Then, the authors propose heuristics to solve the controller placement problem in large networks. To this end, they use an algorithm based on *Pareto Simulated Annealing* [2]. In this work, the authors focus on maximising network resilience, where they considered the placement of controllers in a dynamic SDN network, in which there are latency variations between controllers and their switches. However, they do not consider the placement of controllers taking into consideration the dynamic network allocation.

The positioning needs to be chosen carefully. The *framework* POCO, has the ability to handle small and medium-size topologies, which provides the solution in seconds. However, for large-scale networks, the exhaustive evaluation needs a considerable amount of computational effort and memory usage. It is in this context that the search for a computational solution closer to real environments becomes necessary.

The work described in [9] performs a comparative analysis in terms of traffic capacity. The central objective of the work is to present an analysis of two controllers named Pox and Ryu respectively, in terms of traffic handling capacity. The Mininet emulator was used to emulate the environment of SDN controllers and thus monitor traffic performance. However, the work did not consider the real network characteristics, such as link capacity and occupancy.

As for the comparison of controllers, the work done by [11], one of the first comparative studies of SDN controllers, considered a limited number of controllers (NOX, NOX-MT, Beacon and Maestro) focusing only on controller performance. With the advancement of technologies, such controllers are already considered outdated.

The work described in [1] makes a qualitative comparison between two open-source SDN controllers, OpenDaylight and Open Network Operation System (ONOS). The study focuses on the Northbound interface of these devices.

More recent research, such as the one conducted by [8], studies and evaluates the performance of some popular open-source controllers such as ONOS, Ryu, Floodlight and OpenDaylight in terms of only latency as a metric, using an OpenFlow *benchmarking* tool called *Cbench*. In the paper by [3], a *survey* is presented with an in-depth survey of the SDN controller placement problem, which extensively classifies existing work from various perspectives.

The work described in [10] performs a qualitative evaluation of open source SDN controllers (MUL, Beacon, Maestro, ONOS, Ryu, OpenDaylight, Floodlight, NOX, IRIS, Libfluid-based and POX). The metrics evaluated are latency and throughput performed over a variable number of switches. The results obtained suggest that MUL and Libfluid-based have the best throughput performance while Maestro showed the best latency performance.

In this work, ONOS and Floodlight controllers will be compared and evaluated in two placement scenarios in an SD-WAN network, according to the methodology that will be presented in the next section.

3 Metodology

In this section, we show the base tool used, named POCO [5], an exhaustive method to compute the optimal placement of controllers, implemented in Matlab and available as open-source software, but considering the geographical distance and static latencies, only. Also, we used the Mininet [7] as an SDN emulator for performance analysis and to compare the impact of placement ONOS and Floodlight controllers in the operation of OpenFlow enabled networks. Figure 1, shows the RNP topology used in POCO's graphical interface.

The entire testing environment was configured on a virtual machine (VM), created with the Ubuntu 16.04.05 LTS operating system, containing the default kernel version 4.4.0-87-generic, 8 GB RAM, 8 processors and internal storage space of 30 GB. This VM was configured on a desktop with Windows 10 Home - 64 bits, Intel Core i7-9700 processor and 16 GB of RAM.

We choose the Ipê network, from RNP, which is the first Internet access network in Brazil and integrates more than 800 education and research institutions

Fig. 1. Graphic interface of POCO.

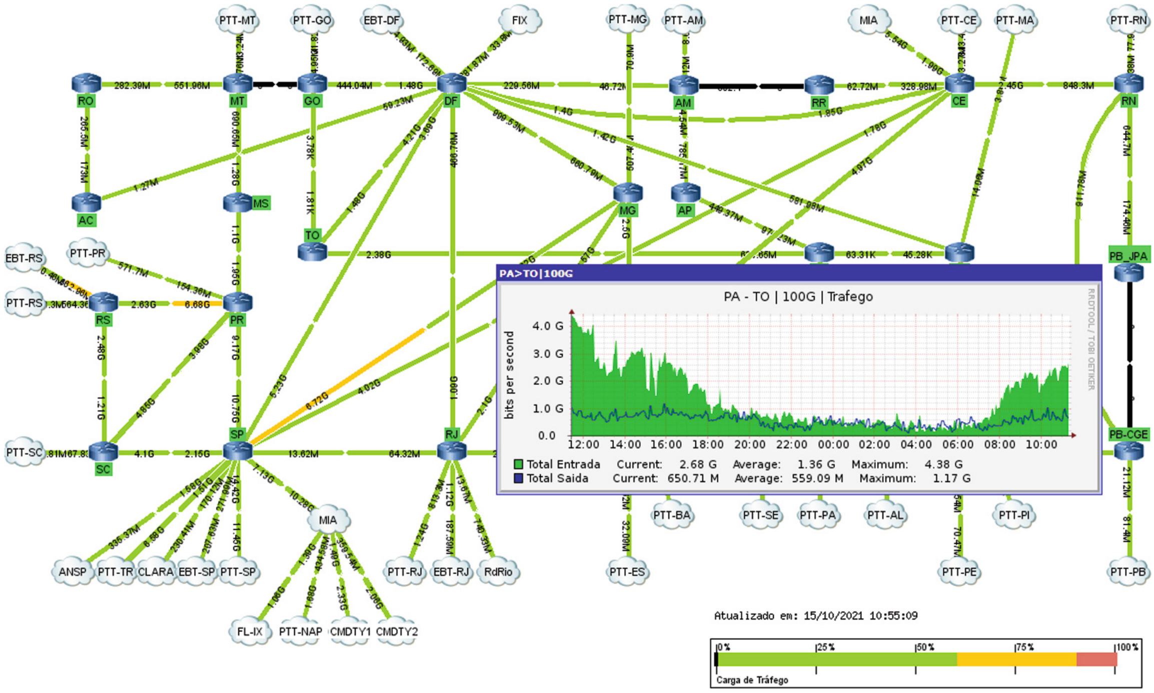

Fig. 2. RNP connections and traffic panorama.

in the country, with 28 nodes positioned in all states of the country territory for the experiments. All information used in the experiments regarding the connections between the links, bandwidth and incoming and outgoing traffic between the nodes were collected from the RNP's real topology traffic panorama, available on the institution's website, as a snapshot shown in Fig. 2.

Initially, a study of the main functionalities of the POCO tool for defining the controller positioning in an SD-WAN network was carried out. Besides that, ONOS and Floodlight controllers were also compared, which were evaluated in 2 placement scenarios, where in the first scenario only the geographical distances and latencies were considered (without link occupancy) and in the second scenario, link occupancy was considered. Then, all 28 possible positions for the controller in RNP's topology were evaluated, and these positions were defined as (i) optimal position and (ii) worst position. The optimal position is that where latency between nodes and controllers is minimised.

Thus, the placement positions indicated by Mininet were compared with the results from the exhaustive method tool – the POCO tool –, which only considers scenario 1 (without traffic). These comparisons considered not only static metrics such as the geographical distances between network nodes but also quantitative measures, such as data transfer, *jitter* and packet loss according to the scenarios.

4 Results

With the entire test environment configured, the tests were performed within Mininet, integrating the controllers in scenarios without traffic (scenario 1) and with traffic (scenario 2). Then, for each of the 28 RNP topology nodes, we use the *fping* command to obtain the latency data from the controller to the destination place in all the other nodes each time. Data gathered has a confidence interval of 95%. Results are shown with the controller positions – over these RNP topology nodes – on the x-axis.

To perform the tests in scenario 2, that considered network traffic (link occupation), the *iPerf* tool was used for measurements, considering the capacities of the links and the real traffic values extracted from the RNP's topology traffic panorama. After that, the test was also performed with the *fping*, in the same way as in the scenario without traffic.

This way, the results of the tests with the controllers in scenario 1 (without traffic), considering the best and worst position defined for the controller, were compared with the results obtained by the calculations of the POCO tool. Thus, these results could confirm the positions established by the POCO tool or show new ones, considering other factors such as the link capacity. The tests in scenario 2 (with traffic) showed the importance of considering metrics such as link capacity, data transfer capacity, *jitter* and packet loss, besides only the geographical distance between nodes.

For the execution of the tests with the POCO tool, the RNP's topology data was gathered during 3 different courses (November 2020, March 2021 and July 2021) after some changes were observed in the topology, for example, some links ceased or their capacities were extended, or other new links were established, creating, in fact, 3 different topologies according to these courses.

4.1 POCO

As the first results, it was calculated the placement positions defined by the POCO tool for the 3 topologies mentioned. For this, the fault-free scenario was considered, with only one controller (k = 1). In these results, it is possible to notice that the values of the average latencies are identical or very close for the same positions of the tested topologies. Therefore, it can be noticed that the results were maintained with the exhaustive method used by the POCO tool for the latencies calculations based only on the geographical distances. Thus, it can be concluded that even with the changes in the topologies links, not in the geographical positions of the nodes, the results did not show significant changes, that is, the distances and latencies were maintained, as shown in Fig. 3

Then, for the 3 topologies tested, position 7 (DF) was defined as the best position, with the lowest average latency, and position 1 (AC) as the worst position, with the highest latency. The position definitions are represented by larger circles in Fig. 4a and Fig. 4b.

In Fig. 4a, it is also possible to notice that the node defined with the lowest latency (7-DF), is also the most centralized node in the topology, therefore, the one with the shortest distance between two points, considering only the geographical distances between the nodes with established links. Also, it is verified that node 7 (DF), has a good amount of redundant links, which facilitates the connection with the other nodes of the network. From another perspective, the

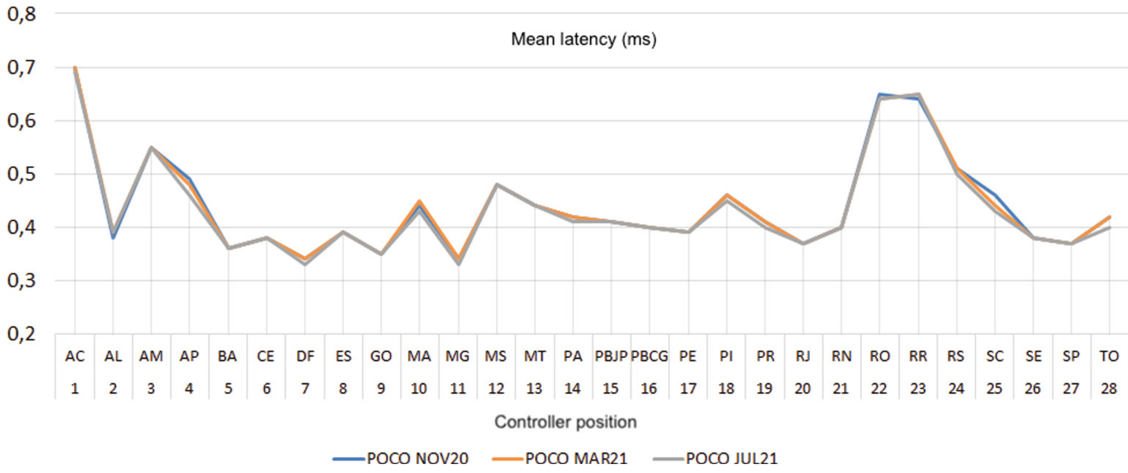

Fig. 3. Classification of positions – POCO.

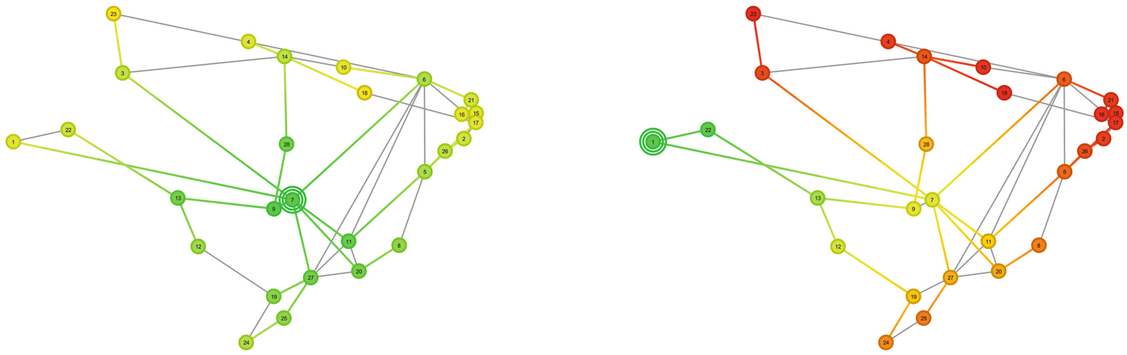

(a) Best position (b) Worst position

Fig. 4. POCO.

reason for node 1 (AC) to have been defined as the worst position can be justified by its geographical position being the most distant in relation to the other nodes of the topology. Thus, it becomes another indication that the geographical distance between the nodes has enough weight in the decisions through the POCO tool, which may not be sufficient for the decision of the controller positioning in a real environment.

4.2 Controllers' Classification

After the classification of the positions by the POCO tool, the topologies were also tested in Mininet using the ONOS and Floodlight controllers, considering the 2 proposed scenarios. Also, tests were performed considering critical performance metrics, such as *jitter* and data loss.

Finally, the results of the scenarios were compared, thus realizing the impacts of such metrics on the SDN controller placement.

Scenario 1 – No Link Occupancy – We show the comparisons of the latency metric between the POCO tool and the controllers in the scenario without traffic. For a matter of space, only the results of the November 2020 topology will be discussed.

First, the tests with the ONOS and Floodlight controllers performed without occupation of the links (scenario 1), showed that both controllers considered position 28 (TO) as the best position, contrary to the choice calculated by the POCO tool, which defined the position 7 (DF) as the best position, shown in Fig. 4a. In the definition of the worst position, the test results considered position 22 (RO), whereas the POCO tool considered position 1 (AC), as in Fig. 4b. These positions are identified by the markers in the curve of Fig. 5.

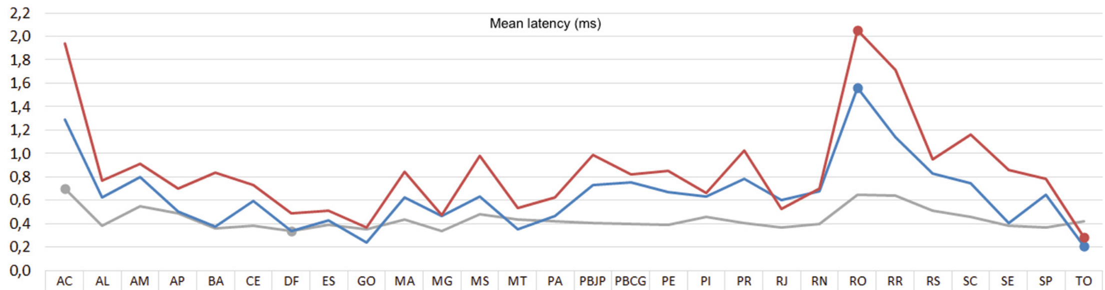

	AC 1	AL 2	AM 3	AP 4	BA 5	CE 6	DF 7	ES 8	GO 9	MA 10	MG 11	MS 12	MT 13	PA 14
POCO	0,70	0,38	0,55	0,49	0,36	0,38	0,34	0,39	0,35	0,44	0,34	0,48	0,44	0,42
ONOS	1,29	0,63	0,80	0,50	0,38	0,59	0,34	0,43	0,24	0,62	0,46	0,64	0,35	0,47
Floodlight	1,94	0,77	0,92	0,70	0,83	0,73	0,49	0,51	0,37	0,84	0,48	0,98	0,54	0,63

	PBJP 15	PBCG 16	PE 17	PI 18	PR 19	RJ 20	RN 21	RO 22	RR 23	RS 24	SC 25	SE 26	SP 27	TO 28
POCO	0,41	0,40	0,39	0,46	0,41	0,37	0,40	0,65	0,64	0,51	0,46	0,38	0,37	0,42
ONOS	0,73	0,75	0,67	0,64	0,78	0,60	0,68	1,56	1,14	0,83	0,75	0,41	0,65	0,21
Floodlight	0,99	0,82	0,85	0,66	1,02	0,53	0,70	2,06	1,71	0,95	1,16	0,86	0,79	0,29

Fig. 5. Topology RNP Nov/20 – no link occupancy (scenario 1).

So, it can be observed that with the inclusion of the controllers in the network other parameters were considered, besides the distances, such as for example, the processing of the controllers and the processes of network discovery and definition of the best paths. This way, in the scenario without traffic, the positions defined as the lowest and highest latency contradicted the POCO definitions. However, in a more careful analysis, position 28 (TO), defined as the best position in the tests with both controllers, in spite of also being a central position, may not be the most advantageous in a real scenario, with several types of applications and traffic, where it depends on several other factors. An indication of the disadvantage of this position could be the small number of redundant links (TO-GO; TO-PA), which may generate a bottleneck in the network.

Another fact observed in the tests was that the ONOS controller presented better performance compared to Floodlight regarding latency, shown in Fig. 6. In this graph, it is possible to see that the results with ONOS showed lower average latency in most of the tested positions, which means that this controller had the best use in the processing demanded by each received flow in the scenario without traffic.

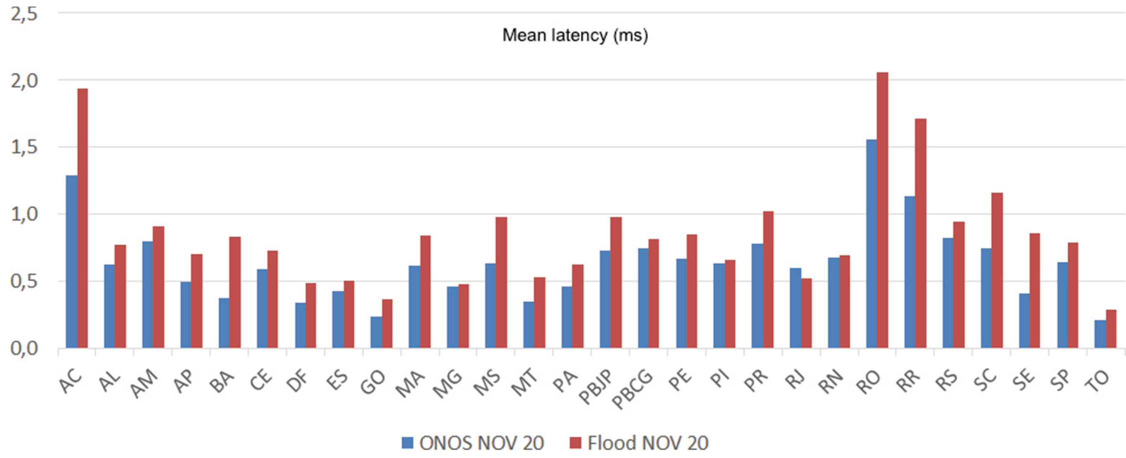

Fig. 6. Comparison of latency metric in scenario 1.

Scenario 2 – With Link Occupancy – In this scenario, the experiments were developed considering the traffic parameters of a real network that can significantly impact the placement of the SDN controller. Therefore, in this case, the results could not be directly compared with the placement defined by the POCO tool, because the tool does not consider the dynamic parameters characteristic of real networks. However, in the graphics presented in the results the definitions of the POCO tool calculations were included, just to demonstrate the differences and the importance of the considerations of the tested characteristics.

Initially, the experiments were performed with the November 2020 topology, with the occupancy of the links. The tests with the ONOS controller presented position 9 (GO) as the position with the lowest average latency and position 23 (RR) with the highest latency. Therefore, comparing these results to the tests with the same topology in a scenario without traffic, where positions 28 (TO) and 22 (RO) were defined, respectively, as the best and worst positions, one can notice that the definitions of the positions were changed, that is, in a scenario with traffic and an active controller in the network other factors were considered, such as, for example, packet traffic, *jitter* and data loss, besides the processing performed by the controller for the treatment of packet flows.

Furthermore, tests were performed with the Floodlight controller in the same topology, which also presented different results from the tests performed with the same controller in the scenario without traffic. In the first scenario, without traffic, Floodlight tests also defined positions 28 (TO) and 22 (RO) as the best and worst positions. In the scenario with traffic, position 9 (GO) was defined as the best position, which coincided with the definition of the test with ONOS in the same scenario, and position 18 (PI) as the worst position. Figure 7 shows the markers with the positions defined in the presented tests and the average latencies for comparison of the other positions.

Figure 7 also shows that the result for the definition of the best position (9-GO) in the tests with the two controllers was very close to the result of the position defined by the POCO tool (7-DF), however, if the position defined in the tests with the controllers in the POCO tool is simulated, that is, if the

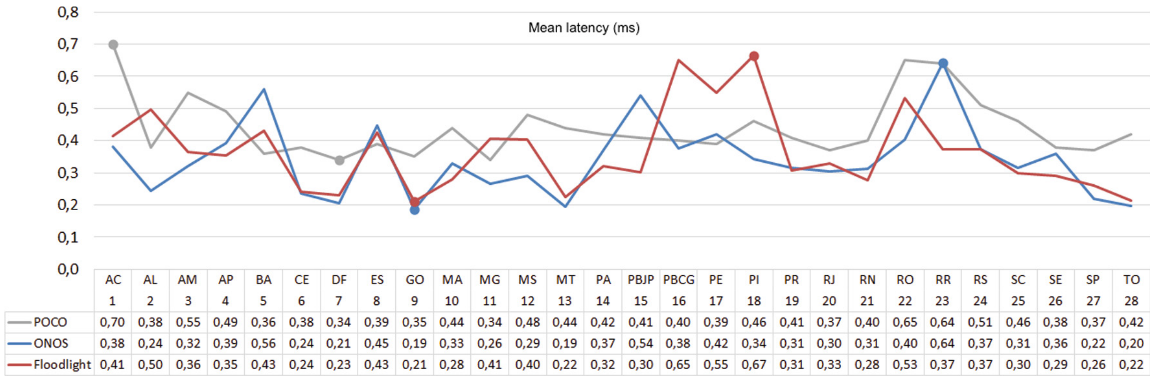

Fig. 7. Topology RNP Nov/20 – with link occupancy (scenario 2).

controller position is changed from position 7 (DF) to 9 (GO), other paths will be established, considering only the distances between the nodes.

As confirmation, the Table 1 brings the example of the comparison between two paths with the same origin and destination, having as origin node 9 (GO) and destination node 1 (AC). In this example, also observed in Fig. 8a, it is shown that the path with the smallest number of jumps (9-7-1) is the most distant, and for this reason, it was not chosen in the simulation with the tool, which defined the path with the greatest number of jumps (9-13-22-1), but with the smallest geographical distance. So, if the definition of the POCO tool is followed for the calculation of the best position, perhaps a path with the smallest number of hops between two points is not considered for the positioning of the controller, a fact that may not be true in a real network, where other factors must be considered, such as, for example, the routing protocol. For comparison, Fig. 8b shows the setting calculated by the POCO tool for the best position.

Table 1. Comparison of paths and distances.

Ways and distances									
Source	Dest.	Jump 1	Dist (Km)	Jump 2	Dist (Km)	Jump 3	Dist (Km)	Total jumps	Total dist.
9	1	(9,13)	764,35	(13,22)	946,38	(22,1)	437,81	3	2,148,54
9	1	(9,7)	161,40	(7,1)	2.234,77	–	–	2	2.396.17

Thus, these results also highlight the importance of considering other metrics, besides distances and latencies, for SDN controller positioning decision-making.

Finally, in this scenario, similar to the previous scenario, the ONOS controller also showed better overall performance with reference to average latencies, shown in Fig. 9.

Comparisons Between Scenarios 1 and 2. The scenario without traffic presented higher latency compared to the scenario with traffic. This fact can be justified by the network discovery and best path processes, besides the treatment

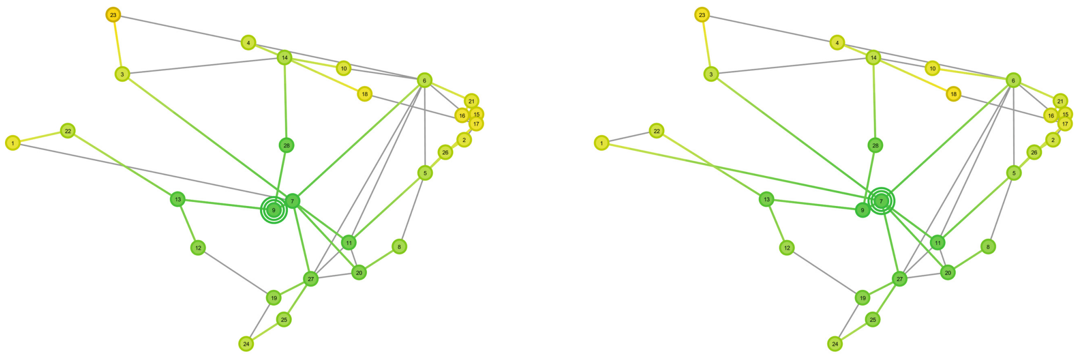

(a) Best ONOS position (b) Best position – POCO.

Fig. 8. POCO.

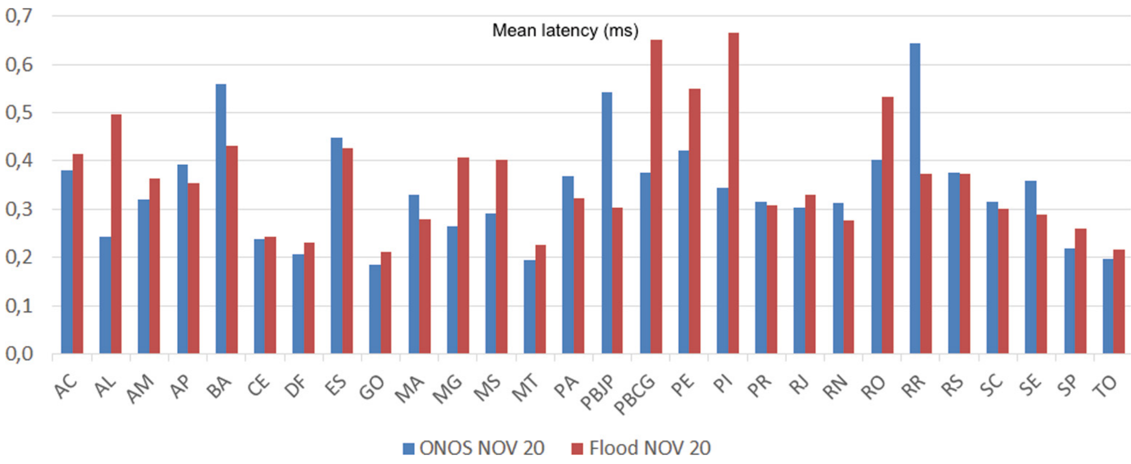

Fig. 9. Comparison of latencies in scenario 2.

of packet flows by the controllers, which caused this difference in the tests with the two scenarios. In other words, the way the controllers store data in RAM, in the *buffer* and in the *cache*, and also the processing demanded by each received flow makes the controllers have different performance. Therefore, these factors can also reduce the responsiveness of the controller and tend to degrade its performance.

In this sense, when a triggered flow is received by an OpenFlow switch and the destination is not found in its flow table, the switch forwards the first packet of the flow to the controller. Then, the controller after receiving the packet defines the forwarding rule to be used for the flow and forwards a message to the switch, which from then on will forward the next packets in the flow direction. Thus, it is expected that the first packet of each generated flow presents a higher latency (RTT – *Round-Trip Time*) compared to subsequent packets.

This fact occurred in the tests with both scenarios, however, in the scenario with traffic the network discovery and best-paths processes started with the first packets transmitted with the *iPerf* in the traffic generation phase between node pairs, that is, before the tests with the *fping*, unlike the scenario without traffic, that the network discovery and best-paths processes happened together with the latency test with the *fping* when firing the first flow, therefore, in this scenario the resulting latency was higher. Figure 10 shows the difference in testing with the ONOS controller in the November 2020 topology.

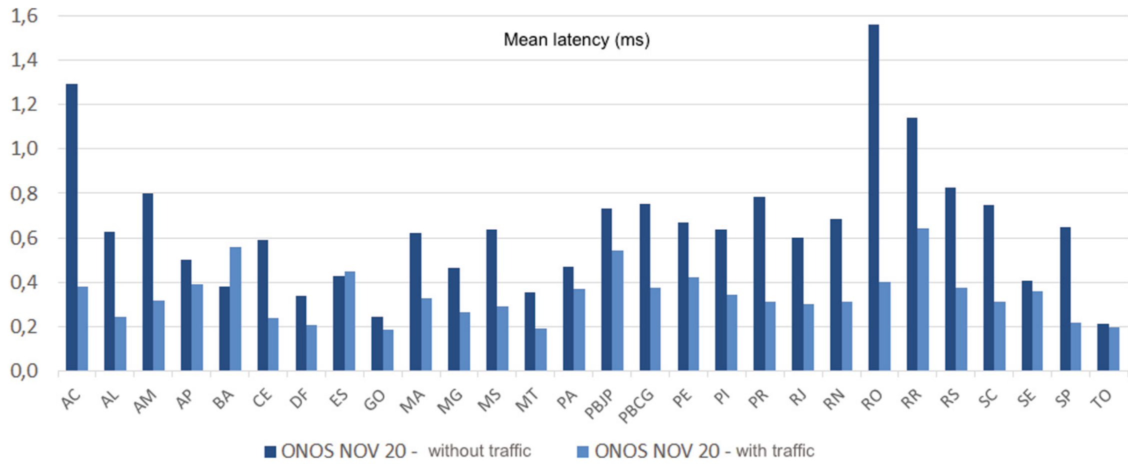

Fig. 10. ONOS – comparison of latencies in scenarios 1 and 2.

Furthermore, it was observed that the ONOS controller in the scenario with traffic did not show higher latencies in the first packets of the test flows with *fping* since, after the start of the traffic transfer with *iPerf* between the pairs of directly connected nodes, the controller maintained a connection to each OpenFlow switch and thus did not change the latency of the first packets of the test with *fping* and the switches were able to send and receive packets as long as there was space in the controller buffers.

The same tests were performed including the Floodlight controller in the surveyed topology. Then, in the same way, as occurred in the tests with ONOS, the tests in the scenario without traffic showed higher average latency than in the tests in the scenario with traffic. Figure 11, shows this difference.

Finally, with reference to the first packets of the generated flows, it is possible to conclude that comparing the two controllers in the tests with the scenario without traffic, the ONOS controller presented a worse performance in the process of calculating the best path, which generated higher latencies in the transmission of the first packets. However, it outperformed the Floodlight controller in the scenario with traffic, where it stabilised the connection with the network nodes after the generation of traffic for the occupation of links and did not generate delay in the first packets of the tests with *fping*.

Critical Performance Metrics. In this section, other tests are presented considering some critical metrics that can impact the positioning and also in controller operation.

The information displayed was obtained through the generation of network traffic using the *iPerf* tool, in the same way as performed in the tests with

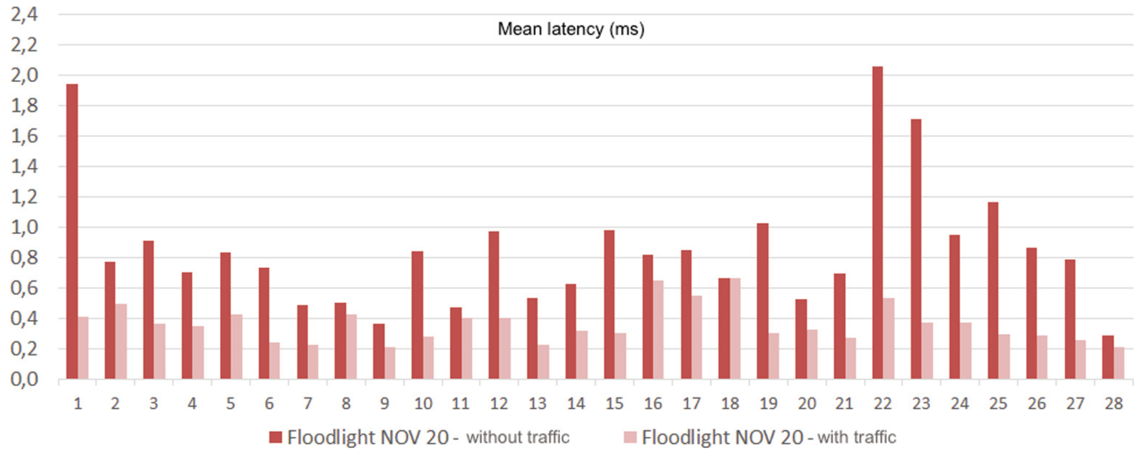

Fig. 11. Floodlight – comparison of latencies in scenarios 1 and 2.

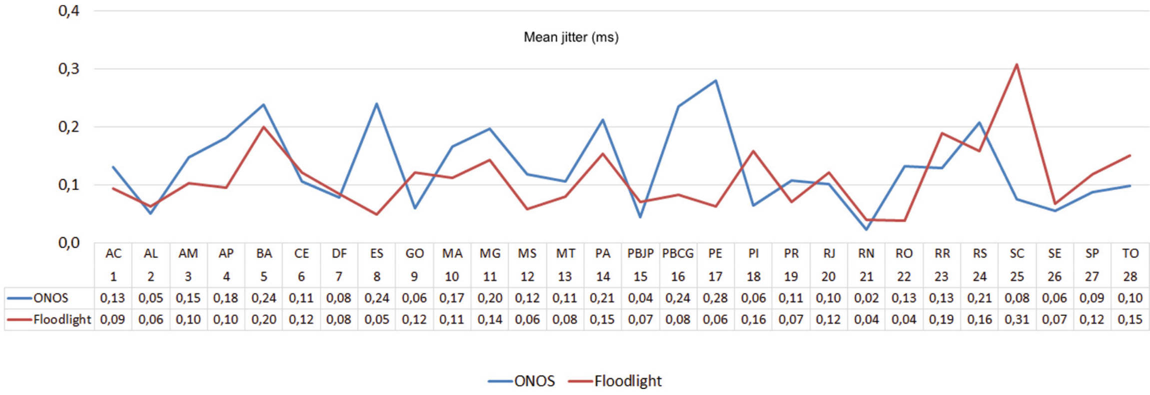

	AC	AL	AM	AP	BA	CE	DF	ES	GO	MA	MG	MS	MT	PA	PBJP	PBCG	PE	PI	PR	RJ	RN	RO	RR	RS	SC	SE	SP	TO
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
ONOS	0,13	0,05	0,15	0,18	0,24	0,11	0,08	0,24	0,06	0,17	0,20	0,12	0,11	0,21	0,04	0,24	0,28	0,06	0,11	0,10	0,02	0,13	0,13	0,21	0,08	0,06	0,09	0,10
Floodlight	0,09	0,06	0,10	0,10	0,20	0,12	0,08	0,05	0,12	0,11	0,14	0,06	0,08	0,15	0,07	0,08	0,06	0,16	0,07	0,12	0,04	0,04	0,19	0,16	0,31	0,07	0,12	0,15

Fig. 12. ONOS x Floodlight – *jitter* metric.

scenario 2, described in Fig. 12. This traffic was generated for a total time of 60 sec, and after that, the generated information was analysed and compared, as will be presented in the next Sections.

Jitter. Briefly, *jitter* is the variation in latency, i.e. this metric can be defined as the measure of variation or "fluctuation" in the time it takes for a data packet to go to a destination and back.

Figure 12 shows that the ONOS controller had a slight disadvantage in the isolated comparison of the metric *jitter* with the Floodlight controller, as it presented a higher variation of the total delay, considering all the tested positions. Therefore, other factors should also be analysed.

Packet Loss. These experiments were run on fault-free scenarios, so only packet transfer losses were considered, and no failures on the links or devices were simulated.

Figure 13 shows the comparison between controllers on the measured packet loss. The results show that the ONOS controller had a lower average packet loss than the loss tested with the Floodlight controller.

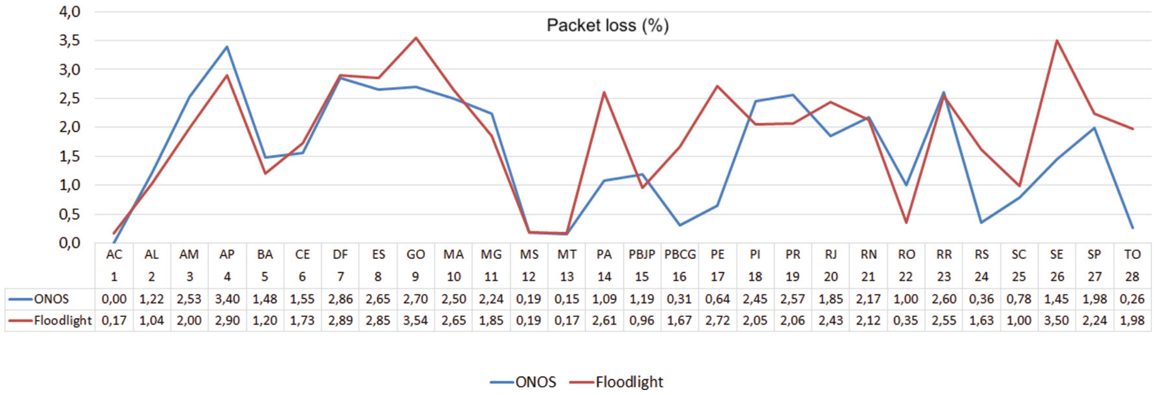

Fig. 13. ONOS x floodlight – packet loss metric.

In summary, in the comparison between the 2 controllers, the Floodlight controller showed less variation of the delay (*jitter*), however, higher packet loss.

Therefore, with the same amount of packets transferred in the two tests with the controllers, the fact that ONOS presented lower packet loss, can be justified by the better treatment given to the flow received in the controller, which is linked to its storage capacity. On the other hand, Floodlight presented lower *jitter*, which evidences the low storage of data in RAM, in the *buffer* and in the *cache*, however, this can cause higher data loss and can reduce the response capacity of the controller.

Thus, it is also noted that in addition to the metrics of distances, latencies and considerations of network traffic, it is also important to identify traffic priorities and decide which will be the most appropriate controller for this purpose.

5 Conclusions and Future Work

This work presented an analysis of the use of the POCO tool and simulation to shedding light on the problem of controller placement in SDN. In this context, it was described the implementation performed in the *framework* POCO that, in general, presented reliable results considering a static network topology, however, the dynamic characteristics of the network can impact significantly the choice of the controller placement. Also, it is possible to observe that the choice of a particular type of controller should be thought of in accordance with the desired purpose because its characteristics can also impact the placement position choice.

Also, considering the real characteristics of the tested topologies, the ONOS controller presented better results in both test scenarios (with and without traffic), regarding the latency parameter. In these results, both controllers diverged from the definitions and positions uncovered by using the POCO tool.

Finally, it is worth emphasising that efficient controller positioning tries to improve the performance of metrics such as latency, traffic priorities, loss and so on. However, the study for the controller positioning problem can still comprise several different solutions. Therefore, it is expected to extend the study to the development and improvement of tools for controller positioning based on performance measures of real networks, and that these tools can analyze the controller positioning problem in the presence of data traffic.

Acknowledgements. This work has been conducted under the project "BIOMA – Bioeconomy integrated solutions for the mobilization of the Agri-food market" (POCI-01-0247-FEDER-046112), by "BIOMA" Consortium, and financed by the European Regional Development Fund (FEDER), through the Incentive System to Research and Technological Development, within the Portugal2020 Competitiveness and Internationalization Operational Program.

The authors are grateful to the Foundation for Science and Technology (FCT, Portugal) for financial support through national funds FCT/MCTES (PIDDAC) to CeDRI (UIDB/05757/2020 and UIDP/05757/2020) and SusTEC (LA/P/0007/2021).

References

1. Bondkovskii, A., Keeney, J., van der Meer, S., Weber, S.: Qualitative comparison of open-source SDN controllers. In: NOMS 2016–2016 IEEE/IFIP Network Operations and Management Symposium, pp. 889–894. IEEE (2016)
2. Czyzżak, P., Jaszkiewicz, A.: Pareto simulated annealing-a metaheuristic technique for multiple-objective combinatorial optimization. J. Multi-Criteria Dec. Anal. **7**(1), 34–47 (1998)
3. Das, T., Sridharan, V., Gurusamy, M.: A survey on controller placement in SDN. IEEE Commun. Surv. Tutorials **22**(1), 472–503 (2019)
4. Heller, B., Sherwood, R., McKeown, N.: The controller placement problem. ACM SIGCOMM Comput. Commun. Rev. **42**(4), 473–478 (2012)
5. Hock, D., Gebert, S., Hartmann, M., Zinner, T., Tran-Gia, P.: Poco-framework for pareto-optimal resilient controller placement in sdn-based core networks. In: 2014 IEEE Network Operations and Management Symposium (NOMS), pp. 1–2. IEEE (2014)
6. Lange, S., et al.: Heuristic approaches to the controller placement problem in large scale SDN networks. IEEE Trans. Network Serv. Manage. **12**(1), 4–17 (2015)
7. Lantz, B., Heller, B., McKeown, N.: A network in a laptop: rapid prototyping for software-defined networks. In: Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks, pp. 1–6 (2010)
8. Mamushiane, L., Lysko, A., Dlamini, S.: A comparative evaluation of the performance of popular SDN controllers. In: 2018 Wireless Days (WD), pp. 54–59. IEEE (2018)
9. Rastogi, A., Bais, A.: Comparative analysis of software defined networking (SDN) controllers-in terms of traffic handling capabilities. In: 2016 19th International Multi-Topic Conference (INMIC), pp. 1–6. IEEE (2016)
10. Salman, O., Elhajj, I.H., Kayssi, A., Chehab, A.: SDN controllers: a comparative study. In: 2016 18th Mediterranean Electrotechnical Conference (MELECON), pp. 1–6. IEEE (2016)
11. Tootoonchian, A., Gorbunov, S., Ganjali, Y., Casado, M., Sherwood, R.: On controller performance in software-defined networks. In: 2nd {USENIX} Workshop on Hot Topics in Management of Internet, Cloud, and Enterprise Networks and Services (Hot-ICE 12), pp. 1–6 (2012)