Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 00 (2021) 000–000

Procedia Computer Science

www.elsevier.com/locate/procedia

International Conference on Industry Sciences and Computer Sciences Innovation

Multiclass data plane recovery using different recovery schemes in SDN: a simulation analysis

Luisa Jorge[a,d,*], Paulo Melo[b,d], Teresa Gomes[c,d]

[a]*Research Centre in Digitalization and Intelligent Robotics (CeDRI), Instituto Politécnico de Bragança, 5300-253 Bragança, Portugal*
[b]*Centre for Business and Economics Research (CeBER), Faculty of Economics, University of Coimbra, 3004-512 Coimbra, Portugal*
[c]*University of Coimbra, Department of Electrical and Computer Engineering, Polo 2, 3030-290 Coimbra, Portugal*
[d]*INESC Coimbra, University of Coimbra, Department of Electrical and Computer Engineering, Polo 2, 3030-290 Coimbra, Portugal*

Abstract

To provide dependable services SDN networks need to be resilient to link or switching node failures. This entails, when faults occur, ensuring differentiated types of recovery, according to carried traffic, to routing paths. However, the choice of the recovery scheme best suited to each traffic class is not direct, nor is obvious the impact of the combination of various recovery schemes, according to traffic classes. We explore the usage of different recovery schemes for traffic with distinct requirements Simulation analysis confirms that using different recovery schemes for distinct types of traffic does create differentiated effects in terms of traffic carried and bandwidth usage.

Keywords: SDN; Data-plane recovery; protection; rerouting; simulation

1. Introduction

Software-Defined Networking (SDN) refers to a class of networks in which the control plane (supporting routing decisions and detecting the most convenient path for data to flow in the network) is separated from the data plane (which concerns itself with efficiently delivering data between individual switching devices).

In this work, data plane protection is the recovery of the data traffic being routed. This recovery can be predetermined before the failure occurs (called precomputed protection) so that at failure time the recovery involves

[*] Corresponding author. Tel.: +351-273-303041 ; fax: +351-273-325405 .
E-mail address: ljorge@ipb.pt

only the switches (if that is possible); or it can be performed (by the controller) when the failure occurs (called rerouting). In SDN networks, much research has recently emerged addressing the problem of protection in the control plane (which among other things includes the problem of protecting the paths between the controller devices - hereafter referred to simply as controllers, and the controlled switching devices - hereafter referred to simply as switches). As stated in [1], the same attention has not been paid to data plane protection (problem of protecting the routing of data between the controlled switches).

To support traffic with demanding service requirements, SDN networks must be able to offer reliable services with guaranteed QoS, where the recovery mechanism used must be fast enough to recover from faults while maintaining the pre-agreed service level. However, usually not all traffic requires this kind of service, with most network applications being able to accept non-guaranteed service (best effort).

Several mechanisms like MPLS labels can be used in SDN (see [2]) to help the transition from MPLS-related mechanisms to SDN mechanisms, but direct application of MPLS-based schemes may not be very efficient (according to [1], it may require additional redirections or not be able to protect all traffic). It will therefore be necessary to use SDN's intrinsic capabilities to achieve similar functionality. Although routing protection is not a recent issue, the need for efficient protection schemes is increasing, due to the substantial amounts of traffic carried and the distinct types of services supported by communication networks, including many critical services.

When a fault happens in the Active Path (AP, path where traffic flows when there is no failure, also called primary path), the recovery mechanism must redirect the traffic to a Recovery Path (RP, the path by which the traffic is restored after the occurrence of a fault, also called backup path) which bypasses the fault (called path restoration). The two basic recovery models used to redirect traffic are usually called rerouting and protection [3]. These approaches differ mainly on when the RP is established. When rerouting is used the RP is computed only upon fault detection in the AP. Protection switching pre-establishes a RP before any failure detection in the protected AP. The recovery scope is usually characterized as global or local. Global recovery intends to protect against any link or node fault in a path, whereas local recovery intends to protect against a link or node fault. Local recovery is attempted by the node immediately upstream of the fault [4]. Between local and global protection, segment protection can be used to protect paths defined as composed of segments (sub-paths) that can be recovered independently.

When a recovery path is pre-established, its resources can be pre-reserved (or, in SDN, set aside by the controller) so they can be guaranteed to be available upon fault. The reserved bandwidth (BW) may be dedicated to a single resource, however to provide more efficient resource usage, pre-reserved resources can be shared by multiple primary resources that are not considered probable to fail at the same time [5]. When the backup bandwidth of different AP is shared, this is called Inter-demand sharing. This type of approach is used for 100% protection in single failure scenarios. When an AP is protected using several RP and they share bandwidth, we say intra-demand sharing takes place.

When the path used to route traffic between source and destination is affected by a fault it must be recovered, preferably in a way that should not be noticeable to the service using it. In SDN networks, even if recovery paths have been pre-computed, this may either be performed using pre-computed information already present in the switching device flow tables or require intervention of the controller.

Section 2 explores recovery in the SDN data plane, and the characteristics of different recovery schemes than can be applied to this task. In section 3 is described how to handle traffic with different attributes and needs, and how this can be supported in terms of data-plane recovery. In section 4 a simulation environment created to explore the effect of the interaction of recovery schemes with traffic from different class-types is described, and results are presented from its application on combinations of networks, recovery schemes and class-types, and patterns discussed. Finally, section 5 proposes some concluding remarks.

2. Protecting the SDN data plane

In this section, we describe different approaches that can be used to protect the data traffic in the network, when using SDN. Notice that while control traffic (between controllers and switches and between controllers themselves if more than one is present) would also require protection, in this paper we are not directly addressing this problem

2.1. Recovery in SDN data plane

Similarly to standard routing, recovery in SDN networks can be performed solely in the controller (henceforth called standard recovery on the control-plane), with the switching devices consigned to detecting failures and acting on (updated) switching information provided by the controllers. The response times provided by this kind of recovery are probably large, but the quality of paths will be better suited to the current network state. This approach can be considered an extension of the reactive paradigm [6] in the OpenFlow protocol.

To limit delays, an approach like fast-reroute [7, 8] of the MPLS or IP kind can also be used [1, 9] henceforth called data-plane fast-reroute, where alternative protection paths can be pre-computed on the controller and applied by the switches immediately upon failure detection for a directly connected link or node. This allows for usually (very) fast protection, but the resulting paths can be of low quality or even unable to provide actual protection, due to routing loops. This approach has similarities to proactive [6] flows in OpenFlow.

Due to the different tradeoffs, the choice of which approach to use may depend on the traffic characteristics and operator preference. Recovery schemes in SDN can traditionally be classified according to where the switching between AP and RP occurs (either near the fault - local schemes, or at the traffic origin - global schemes). They can also be classified as either providing (pre-computed) protection or using rerouting. This creates four combinations of recovery schemes that can be implemented:

- Local protection schemes rely on switches upon fault detection, to switch the traffic to a pre-computed recovery path. Several mechanisms can be used to provide this capability, frequently using the switch flow tables. Global protection in SDN must usually be performed by the controller, which upon receiving information from a switch adjacent to the fault updates the source switch tables to use a (pre-computed) alternative route for the affected flow(s). Since this may take some time and lead to dropped traffic until tables are updated, a tunnel may be created from the fault-adjacent switch to the source switch from where the traffic follows the alternative path to the destination (using a remote Loop-Free Alternate [7] path). If the system supports it, receiving traffic from that tunnel may itself trigger the switching in the source switch to use the backup path as standard, thereby reducing the need for direct action from the controller.
- Unlike precomputed protection, rerouting schemes must be implemented upon fault by the controllers (using a reactive model) to be able to incorporate system state information in path computation. Local rerouting schemes and global rerouting schemes are implemented in similar ways, but local rerouting schemes can limit the path change to a minimum, to provide faster calculation and reduce the number of tables to be updated, while global rerouting schemes may consider all the network state information to create optimized paths.

2.2. Recovery schemes

Many schemes have been proposed to support recovery from failure, with variations depending on the network technologies used (IP; MPLS, WDN). Since the number of schemes/variants proposed over time is very high, some attempts of consolidation and comparative analysis of them have also been frequently proposed, for specific areas [10, 11]. Particularly in the area of Fast Recovery (Fast Reroute) much research has been developed, namely [1, 3, 12]. Whereas data-plane protection schemes in SDN assume pre-computation of recovery paths or segments, other recovery technologies fully dependent on reactive rerouting only upon failure have been studied in the past but are less likely to be considered in today’s research for SDN networks.

In this work, we studied several recovery schemes that try to address most of the dimensions presented in the previous subsection. We do not claim the implemented schemes are either “best-in-kind” or incorporate all variations proposed for these approaches, instead we try to present adequate representatives for different recovery scheme characteristics. Table 1 presents the list of the recovery schemes that were implemented. The acronym/implementation name used to refer to each scheme in the analysis indicates the characteristics of the type of scheme to which it corresponds: L – Local and G – Global; P – Protection, R – Rerouting.

Table 1 - Recovery Schemes Implemented

Recovery Path Setup Method	Recovery Scope	Resource Reservation for Recovery Paths	Implementation Name
Protection	Global	Inter-demand Sharing	GP1
		No Reservation	GP2
	Local	Intra-demand Sharing	LP1
		Inter-demand and Intra-demand Sharing	LP2
		Intra-demand Sharing	LP3
Rerouting	Global	-	GR
	Local	-	LR

Some of the proposed schemes allow for sharing of reserved resources for the RP. Inter-demand sharing is the sharing of RP between requests whose AP do not share the same link being protected (thus even if the AP are not fully disjoint it may still be possible to have this kind of sharing). Intra-demand sharing is the sharing of links among several RP protecting different links of the same AP (this is possible only in local recovery). In [13] is shown that the amount of sharing possible is dependent on the amount of available information regarding link usage.

In Table 1, GP1 corresponds to an implementation of the recovery scheme proposed in [13]. It is a global protection scheme in which the AP and RP are determined simultaneously, and the determination is done using inter-demand reservation sharing. GP2 is also a global protection scheme, but without any reservation, where the AP and RP are disjoint paths determined between the source and destination using Dijkstra's algorithm.

Schemes LP1 and LP3 are simplified implementations of Fast Reroute (One-to-one Backup variant) [8, 14] with intra-demand BW sharing and without inter-demand BW sharing. The difference between them is in the paths selection, in LP1 Dijkstra's algorithm is used to determine the single AP and the various RP, while in LP3 more opportunities are given to find disjoint paths by determining additional AP (considering all AP with at most two more links than the minimum) if RP cannot be found for all links in the originally determined AP. LP2 in Table 1 corresponds to the local protection implementation proposed in [5] with the improvements published in [15]. It is a local protection scheme in which the AP and the several RP are determined simultaneously, and the determination is done using inter-demand and intra-demand BW sharing.

The schemes GR and LR in Table 1 correspond to rerouting schemes. In all schemes the AP are determined between the source and destination using Dijkstra's algorithms and the same algorithm is also used in determining the RP for requests that are affected by failures. In GR the RP are determined between the source and destination of the request, while in LR they are determined between the node before the failure and the destination node (or any other node in the AP between the node after the failure and the destination node). These schemes have been implemented both in a reservation variant (with resources reserved immediately before they are used) and no-reservation variant.

3. Protection of several traffic classes / class-types

An inescapable fact of networking is that although all packets in a network may possess similar structure, the requirements of the applications that create them will vary greatly. This means that providing the same kind of protection for all applications is not necessarily the best choice when designing network protection mechanisms.

While it is common to try to measure the factual characteristics of a telecommunications service that bear on its ability to satisfy stated and implied needs of the user of the service – Quality of Service (QoS) [16], more recently focus has moved to determine the actual degree of delight or annoyance of the user of an application or service, resulting from the fulfilment of his or her expectations concerning the utility and/or enjoyment of the application or service in the light of the user's personality and current application/service state – Quality of Experience (QoE) [17]. Many mechanisms have been proposed to support QoS and QoE in SDN, as seen in [18, 19], either using the support present in the different versions of the OpenFlow specification or making use of specific capabilities on SDN controllers.

A mechanism to support differentiated QoS/QoE in a network with several traffic classes is the usage of so-called Bandwidth Constraints Models (BCM) which attributes defined bandwidth to particular traffic classes or sets of traffic classes. This is particularly common in MPLS (see [20]) but lately, this kind of approach has been suggested also for SDN networks (as in [21, 22]), and is the approach followed in this work. To support traffic with distinct characteristics sharing the same network, it is usual to assign traffic with similar handling in terms of bandwidth constraints to a particular Class-Type (CT), defined as "the set of Traffic Trunks crossing a link that is governed by a specific set of Bandwidth Constraints [where a CT] is used for [the purposes of] link bandwidth allocation, constraint-based routing and admission control [and] a given Traffic Trunk belongs to the same CT on all links" [23]. It should be noticed that to be able to always provide the required bandwidth even in face of failures, the system must use some sort of reservation/preemption mechanism to ensure that the needs of higher priority traffic can override those of lower priority. To prevent the overly strict application of reservation from inhibiting traffic even when sufficient capacity is present, it is possible to set aside a certain bandwidth to be used even over reservation limits (until a certain total occupation threshold is reached), as is the case of the Max Allocation with Reservation Bandwidth Constraints Model (MAR) model (defined in [24]).

It is expectable that the actual QoE may depend both on the BCM model used and on the particular parameters applied within this model. In this study, we selected a particular BCM model (MAR) with a set of fixed parameters chosen after an initial viability simulation.

4. Simulation results and discussion

From the previous description it can be expected that the different characteristics of recovery schemes may create distinct traffic experiences when applied to a particular network. Moreover, the use of different schemes for dissimilar CT may also provide unequal results. To test this, we developed a simulation environment, and performed a set of experiments to try to match traffic outcomes to the recovery schemes used for each CT.

4.1. Simulation environment

A simulation environment was built in OMNeT++ [25], to analyse the effect of recovery scheme combinations on some network traffic characteristics. The simulation study used two different network topologies KL-15 [5, 26] and COST-239 [27], with sufficiently distinct characteristics, however, due to space limitations we will only show here the KL-15 (with 15 nodes and 56 directed links) network results. For the KL-15 network link capacity was defined as 60 units for most links and 240 for a small core set of links, as shown in [5]. This large difference in link capacity and network topologies allowed us to use a large set of load configurations to try to find the load best suited to the cases under study.

The requests arrived in the network (one at a time) with bandwidth uniformly distributed in the range [1, 6]. The origin and destination paths were chosen randomly from all nodes. The CT associated with the request was also randomly chosen from a uniform distribution using the Monte Carlo Method. The arrivals follow a Poisson process of intensities 0.1 and 0.2 connection requests per unit time, respectively for the KL-15 and COST 239 networks. The duration of the requests is exponentially distributed with a mean value equal to 1/0.001 time units. To generate distinct load situations the intensity of arrival of requests was multiplied by successive integer values (in the range [1, 6], henceforth called load factors). Each simulated failure affected a single directed link randomly chosen from all directed links in the network (failures also generated as Poisson processes). Faults occurred with a frequency of 0.0005 per time unit and were repaired 0.001 per time unit, regardless of the load conditions.

In each of the performed runs, simulation was executed for 2 800 000 requests with warm-up and cooldown times allowing for 100 000 requests each to minimize transitory effect. For each of the cases, 10 independent replications were run (which allowed obtaining the confidence intervals presented for the different results).

To model different types of traffic and their interaction, we created 4 traffic classes (CT0 to CT3, with CT0 modelling best-effort traffic, CT1 and CT2 normal priority traffic with different values of jitter tolerance and CT3 high priority traffic). The proportion of generated traffic was 50%, 20%, 20% and 10%, respectively for CT0 to CT3. A MAR bandwidth reservation model was used for all runs, reserving 30% for each of the CT1 to CT3 traffic, with no reservation performed for best-effort traffic.

4.2. Analysis

The simulator was used in a set of experiments, trying to assess the impact of using a particular recovery scheme for different traffic classes/CT. In all experiments presented the same recovery scheme is used for all CT except for CT0 (Best Effort traffic). Reservation will not be considered for best-effort traffic, so only rerouting schemes without reservation will be considered admissible for this traffic. The remaining CT, however, for these experiments will all share the same recovery scheme (chosen among schemes with reservation). With this approach, we intend to investigate the impact of different recovery schemes and reservation approaches on CT.

Table 2 - Traffic schemes used for each experiment

Experiment	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
CT0	GR	GR	GR	GR	GR	GR	GR	LR	LR	LR	LR	LR	LR	LR
CT1 - CT3	GP1	GP2	GR	LP1	LP2	LP3	LR	GP1	GP2	GR	LP1	LP2	LP3	LR

Table 3 – Blocking upon establishment - CT1/CT3, load factors 1 to 4

		LR	**GR**	**LP2**	**GP1**	**GP2**	**LP3**	**LP1**
1	CT1			0,63±0,02%	0,68±0,01%	0,85±0,02%	1,25±0,03%	1,50±0,03%
	CT3			0,53±0,01%	0,55±0,01%	0,61±0,01%	0,73±0,01%	0,83±0,02%
2	CT1	0,11±0,00%	0,11±0,00%	6,68±0,04%	9,69±0,04%	12,33±0,02%	19,00±0,03%	19,16±0,04%
	CT3	0,05±0,00%	0,04±0,00%	2,92±0,02%	3,58±0,03%	5,38±0,04%	8,57±0,05%	8,84±0,05%
3	CT1	1,52±0,02%	1,55±0,02%	18,59±0,05%	26,35±0,07%	27,24±0,05%	33,38±0,07%	33,44±0,05%
	CT3	0,46±0,01%	0,46±0,01%	8,24±0,05%	11,20±0,07%	14,24±0,06%	18,88±0,09%	19,08±0,09%
4	CT1	6,43±0,04%	6,50±0,04%	27,51±0,07%	37,48±0,08%	36,67±0,06%	41,91±0,07%	41,93±0,06%
	CT3	1,73±0,03%	1,76±0,02%	13,71±0,08%	19,11±0,09%	21,86±0,08%	26,79±0,10%	26,97±0,11%

Table 4 – Blocking upon fault - CT1/CT3, load factors 2 to 5, rerouting schemes

		GR	**LR**
2	CT1	2,64% ± 0,16%	3,63% ± 0,19%
	CT3	1,45% ± 0,15%	2,28% ± 0,24%
3	CT1	14,88% ± 0,35%	17,75% ± 0,38%
	CT3	9,26% ± 0,41%	12,20% ± 0,47%
4	CT1	28,11% ± 0,59%	30,92% ± 0,79%
	CT3	17,19% ± 0,51%	20,75% ± 0,53%
5	CT1	36,69% ± 0,62%	38,87% ± 0,69%
	CT3	23,26% ± 0,37%	26,59% ± 0,71%

The total number of experiments presented is 14, which corresponds to the possible combinations of 2 recovery schemes without reservation with 7 recovery schemes with reservation, as shown in Table 2. Since the same recovery scheme is in each case applied to CT1, CT2 and CT3, the experiments could be run with only two CT, but splitting traffic among 3 CT allows us to analyse the effect of using BCM on the different CT.

Regarding BCM, it was found in the KL-15 network for small load factors that even after trying to adjust the MAR parameters, CT0 often gets lower blocking probabilities at path establishment than the other CT (only when the other CT use protection schemes). This occurs due to the conjunction of the other CT using protection (and as such to establish a request they may require more than twice the BW that is needed for the AP) and that in the KL-15 network some links have excess bandwidth (which is not the case in COST 239). In all result tables the values are presented as mean value ± half of the confidence interval with a 95% degree of confidence. In Tables 3 to 7, CT2 results are not shown since they are approximately equal to CT1.

Table 3 shows the probability of request rejection upon the path establishment, refused requests divided by total requests. This data is presented by recovery scheme used, for both CT1 and CT3, when CT0 uses the LR scheme. Table 3 results are presented for various load factors, ranging between 1 and 4. The results are only presented for CT0 using LR scheme since it was found that the results do not change significantly if the GR scheme is used by CT0 instead of LR. As can easily be seen Table 3 can be sorted by column in increasing order of blocking (except for very low loads for LR and GR) meaning this characteristic is scheme dependent. Notice however that while LR blocking is assumed to be lower than GR blocking, there are overlaps of the confidence intervals in almost all load factors.

Table 4 presents the probability of disconnection upon fault, defined as number of fault-terminated requests divided by number of requests affected by fault when the GR and LR rerouting scheme is used (due to simulation design, only in rerouting schemes disconnection at fault can occur), by both CT1 and CT3, with CT0 using the LR scheme, for the reasons presented before. Results are presented for load factor values between 2 and 5, with GR performing better than LR consistently.

Table 5 – Active Path average path length - CT1/CT3, load factors 1 to 4

		GP2	**LP1**	**LP3**	**LP2**	**GP1**	**LR**	**GR**
1	CT1	2,143 ± 0,001	2,150 ± 0,001	2,152 ± 0,001	2,467 ± 0,001	2,774 ± 0,001	2,142 ± 0,001	2,142 ± 0,001
	CT3	2,138 ± 0,001	2,142 ± 0,001	2,142 ± 0,001	2,436 ± 0,001	2,653 ± 0,001	2,141 ± 0,000	2,141 ± 0,000
2	CT1	2,177 ± 0,001	2,172 ± 0,001	2,183 ± 0,001	2,522 ± 0,001	2,925 ± 0,002	2,173 ± 0,001	2,173 ± 0,001
	CT3	2,163 ± 0,001	2,169 ± 0,001	2,176 ± 0,000	2,495 ± 0,000	2,791 ± 0,002	2,155 ± 0,001	2,156 ± 0,001
3	CT1	2,170 ± 0,001	2,149 ± 0,001	2,162 ± 0,001	2,576 ± 0,001	2,968 ± 0,001	2,282 ± 0,001	2,283 ± 0,001
	CT3	2,175 ± 0,001	2,168 ± 0,001	2,178 ± 0,001	2,543 ± 0,001	2,875 ± 0,002	2,213 ± 0,001	2,214 ± 0,001
4	CT1	2,154 ± 0,001	2,128 ± 0,001	2,140 ± 0,001	2,581 ± 0,002	2,967 ± 0,003	2,410 ± 0,001	2,411 ± 0,001
	CT3	2,174 ± 0,001	2,159 ± 0,001	2,170 ± 0,001	2,570 ± 0,001	2,921 ± 0,002	2,295 ± 0,001	2,297 ± 0,001

Table 6 – Active Path average bandwidth - CT1/CT3, load factors 1 to 4

		GP2	LP3	LP1	LP2	GP1	LR	GR
1	CT1	5,494 ± 0,003	5,504 ± 0,003	5,509 ± 0,003	6,335 ± 0,004	7,115 ± 0,005	5,490 ± 0,003	5,490 ± 0,003
	CT3	5,495 ± 0,005	5,497 ± 0,005	5,498 ± 0,004	6,420 ± 0,007	6,932 ± 0,005	5,494 ± 0,005	5,494 ± 0,005
2	CT1	5,781 ± 0,003	5,639 ± 0,003	5,589 ± 0,004	6,376 ± 0,004	7,684 ± 0,006	5,514 ± 0,003	5,514 ± 0,003
	CT3	5,622 ± 0,005	5,693 ± 0,004	5,682 ± 0,005	6,352 ± 0,004	7,137 ± 0,007	5,502 ± 0,004	5,502 ± 0,004
3	CT1	5,413 ± 0,003	5,016 ± 0,003	4,979 ± 0,004	6,875 ± 0,004	7,521 ± 0,006	5,713 ± 0,004	5,714 ± 0,004
	CT3	5,818 ± 0,006	5,724 ± 0,005	5,681 ± 0,005	6,550 ± 0,008	7,592 ± 0,006	5,578 ± 0,005	5,577 ± 0,005
4	CT1	4,932 ± 0,005	4,527 ± 0,005	4,499 ± 0,005	6,573 ± 0,005	7,085 ± 0,009	6,343 ± 0,008	6,354 ± 0,006
	CT3	5,751 ± 0,005	5,471 ± 0,005	5,423 ± 0,004	6,753 ± 0,007	7,741 ± 0,007	5,789 ± 0,005	5,790 ± 0,006

Table 7 – Recovery Path average bandwidth - CT1/CT3, load factors 1 to 4

		GP1	LP2	GP2	LP1	LP3	LR	GR
1	CT1	1,744 ± 0,003	2,131 ± 0,003	7,734 ± 0,005	11,024 ± 0,006	10,969 ± 0,006	7,320 ± 0,035	7,920 ± 0,037
	CT3	2,116 ± 0,007	2,280 ± 0,007	7,735 ± 0,005	10,965 ± 0,008	10,932 ± 0,009	7,385 ± 0,054	7,988 ± 0,050
2	CT1	1,702 ± 0,005	1,978 ± 0,003	8,436 ± 0,007	12,113 ± 0,008	12,213 ± 0,007	7,510 ± 0,037	7,981 ± 0,031
	CT3	1,748 ± 0,004	2,137 ± 0,004	7,964 ± 0,004	11,708 ± 0,011	11,702 ± 0,010	7,476 ± 0,070	7,986 ± 0,073
3	CT1	1,518 ± 0,003	1,888 ± 0,004	8,523 ± 0,009	11,238 ± 0,010	11,311 ± 0,008	8,380 ± 0,047	8,498 ± 0,049
	CT3	1,650 ± 0,004	2,032 ± 0,004	8,495 ± 0,009	12,215 ± 0,013	12,293 ± 0,009	8,211 ± 0,054	8,375 ± 0,033
4	CT1	1,306 ± 0,008	1,695 ± 0,004	8,007 ± 0,005	10,300 ± 0,013	10,345 ± 0,012	9,031 ± 0,074	9,014 ± 0,052
	CT3	1,565 ± 0,006	2,003 ± 0,004	8,716 ± 0,007	11,975 ± 0,010	12,064 ± 0,012	8,957 ± 0,080	8,936 ± 0,064

Table 5 presenting average AP path length (number of links in the AP) is ordered by schemes considering only nominal load, but this ranking is not regular for other load factors, since GP2, LP3 and LP1 schemes, although presenting very close AP lengths, do not maintain their relative order with the increase of the load factors. The LP1 scheme presents smaller AP lengths for more load factors, while LP2 and GP1 present significantly higher AP lengths. There is practically no distinction between schemes LR and GR in terms of AP length (grey background).

To verify the ability of schemes to establish new paths, Table 6 presents the average bandwidth used by established request (the bandwidth used in the AP by all established requests divided by the number of established requests). For protection schemes, for load factors higher than 2, generally the BW that is used for AP decreases with increasing load (Table 6), which signals network congestion (with only "smaller" BW requests being established). This does not occur with rerouting schemes (grey background. There is a significant increase in BW per AP in the schemes that use bandwidth sharing (LP2 and GP1) over the ones that do not share (GP2, LP1 and LP3). Among the schemes with sharing, BW per AP for GP1 is significantly higher than LP2. From Table 6, LP3 and LP1 present overlapping confidence intervals for nominal load but for the other load factors the BW consumption by the AP is higher in LP3 than in LP1.

When researching average bandwidth usage by RP it can be seen from Table 7 that there is a large difference between the average BW for the RP of schemes that do share BW (GP1 and LP2) relative to that of those that do not share (GP2, LP3 and LP1). It is also visible that in the schemes with sharing, the average BW used by RP for GP1 is lower than LP2. For rerouting schemes (grey background), we see a higher BW consumption for the RP of requests in GR than in LR (except for load factor 4, where the confidence intervals overlap).

4.3. Discussion

Analysis was carried out on two separate networks, with in general consistent results obtained on both networks. From the results in Table 3, Table 4, and Table 5, we can say that no scheme is superior in all the measures presented in those tables (no overall dominant scheme). However, partial dominance can be found: LP2 is superior to GP1 for all load factors and GP2 is superior to LP3 and LP1. Note however that, as mentioned, the superiority of GP2 in terms of AP length only occurs for nominal load.

When the COST 239 network is used, while AP path length shows lower values than those presented with the KL-15 network, in general the relations between the various schemes are maintained (that is, when there are schemes with significantly larger measures than others, those relations are maintained; when the schemes present measures close to each other in one network, those measures in the other network will also be close to each other). The same is true, in most cases, for the remaining studied measures, and as such the results are omitted from this paper.

From Table 3 on the probability of rejecting requests at the establishment, we can observe that schemes which share the BW reserved for protection have lower request rejection probability than those without sharing. From the same table we can also confirm that local recovery schemes have higher blocking than global recovery schemes (for instance comparing GP2 with LP1 and LP3). However, this statement is only valid when there is no inter-demand sharing. In schemes where there is inter-demand sharing, the existing BW sharing efficiency may make the local perform better than global in terms of blocking, as is the case for LP2 and GP1. Another factor contributing to the higher probability of blocking in the establishment of GP1 is that it uses AP with significantly higher number of links than the one used by LP2 (see Table 5).

The analysis of the results in Table 3 also shows that LP2 is better than GP1 in terms of blocking at establishment, which can be justified because LP2 uses AP that are significantly shorter than GP1, and this is not compensated by GP1 using RP that consume less BW (from the results in Table 5 and Table 7). This seems to show that GP1 tries to over-optimise sharing at the expense of creating longer APs, which leads to is worse results than sharing slightly less but using fewer resources in APs. This makes in the overall results LP2 generally preferred to GP1. The blockage in establishment is lower in the schemes with reserve sharing, LP2 and GP1, than in those without sharing (GP2, LP3 and LP1). However, the opposite is true for the AP path length, i.e., the schemes with sharing are the ones that use paths with a larger number of links (and therefore require larger BW for the AP). On the other hand, concerning the BW consumption for the RP, LP2 and GP1 spend less, as expected.

Regarding the probability of request rejection / blocking in the establishment for the forwarding schemes, we observe that this probability is similar using either local or global recovery, but the same does not occurs for the disconnection upon failure (see Table 4).

5. Conclusions

The article presented a comparative analysis of the effect of several data-plane recovery schemes regarding effect of using particular recovery scheme combinations on network traffic belonging to different traffic classes when subject to different levels of traffic load. This simulation study allowed the collection of statistics regarding request rejection upon path establishment, disconnection at fault, AP average path length, AP average BW usage, and RP average BW usage. An analysis was conducted on two networks for 14 different recovery scheme combinations with traffic belonging to 4 different traffic classes and for several load factors.

From the results collected it could be seen that no particular combination of recovery schemes provides the “best” results on every measure for each traffic class, but that for individual measures some recovery schemes combinations provide significantly better results than others. Therefore, we can see that it can be expected that selecting a particular combination of recovery schemes would be best suited to optimize a particular traffic characteristic, be it ability to overcome faults or resource usage. It could also be seen from the results that the particulars of the network used are not a significant factor affecting outcomes.

Acknowledgements

The work has been partially supported by European Regional Cohesion and Portuguese Funds, under project Resilient Network Design - enhancing availability for critical services (CENTRO-01-0145-FEDER-029312). The support by the Portuguese Science and Technology Foundation (FCT) Grants UIDB/00308/2020 (INESC Coimbra), UIDB/05757/2020 (CeDRI), and UIDB/05037/2020 (CeBER) is also acknowledged.

References

1. Merling, D., Braun, W., Menth, M.: Efficient Data Plane Protection for SDN. In: 2018 4th IEEE Conference on Network Softwarization and Workshops (NetSoft). pp. 10–18. IEEE (2018). https://doi.org/10.1109/NETSOFT.2018.8459923.
2. ONF: OpenFlow Switch Specification Version 1.5.1. (2015). https://opennetworking.org/wp-content/uploads/2014/10/openflow-switch-v1.5.1.pdf.
3. Ganchev, I., Rak, J., Cinkler, T., O’Droma, M.: Taxonomy of Schemes for Resilient Routing. In: Guide to Disaster-Resilient Communication Networks. pp. 455–482 (2020). https://doi.org/10.1007/978-3-030-44685-7_18.
4. Sharma, V., Hellstrand, F., Mack-Crane, B., Makam, S., Owens, K., Huang, C., Weil, J., Cain, B., Anderson, L., Jamoussi, B., Chiu, A., Civanlar, S.: Framework for multi-protocol label switching (MPLS)-based recovery, (2003). https://www.rfc-editor.org/info/rfc3469.
5. Kodialam, M., Lakshman, T.V.: Restorable dynamic quality of service routing. IEEE Commun. Mag. 40, 72–81 (2002). https://doi.org/10.1109/MCOM.2002.1007411.
6. Fernandez, M.P.: Comparing OpenFlow Controller Paradigms Scalability: Reactive and Proactive. In: 2013 IEEE 27th International Conference on Advanced Information Networking and Applications (AINA). pp. 1009–1016. IEEE (2013). https://doi.org/10.1109/AINA.2013.113.

7. Atlas, A.K., Zinin, A.: Basic Specification for IP Fast Reroute: Loop-Free Alternates. (2008). https://doi.org/10.17487/rfc5286.
8. Bryant, S., Filsfils, C., Previdi, S., Shand, M., So, N.: Remote Loop-Free Alternate (LFA) Fast Reroute (FRR). (2015). https://doi.org/10.17487/RFC7490.
9. Chiesa, M., Kamisinski, A., Rak, J., Retvari, G., Schmid, S.: A Survey of Fast-Recovery Mechanisms in Packet-Switched Networks. IEEE Commun. Surv. Tutorials. 23, 1253–1301 (2021). https://doi.org/10.1109/COMST.2021.3063980.
10. Ramirez, W., Masip-Bruin, X., Marin-Tordera, E., Sànchez-López, S.: Managing resilience in carrier grade networks: Survey, open issues and trends. Comput. Commun. 61, 1–16 (2015). https://doi.org/10.1016/j.comcom.2015.02.015.
11. Jorge, L., Gomes, T.: Survey of recovery schemes in MPLS networks. Proc. Int. Conf. Dependability Comput. Syst. DepCoS-RELCOMEX 2006. 110–118 (2006). https://doi.org/10.1109/DEPCOS-RELCOMEX.2006.52.
12. Lemeshko, O., Yeremenko, O., Sleiman, B., Yevdokymenko, M.: Fast ReRoute Model with Realization of Path and Bandwidth Protection Scheme in SDN. Adv. Electr. Electron. Eng. 18, (2020). https://doi.org/10.15598/aeee.v18i1.3548.
13. Kodialam, M., Lakshman, T.V.: Dynamic routing of restorable bandwidth-guaranteed tunnels using aggregated network resource usage information. IEEE/ACM Trans. Netw. 11, 399–410 (2003). https://doi.org/10.1109/TNET.2003.813044.
14. Pan, P., Swallow, G., Atlas, A.: Fast Reroute Extensions to RSVP-TE for LSP Tunnels. (2005). https://doi.org/10.17487/rfc4090.
15. Jorge, L., Gomes, T.: An on-line routing algorithm of locally protected paths with exact reservations. In: Proceedings of the 6th IASTED International Conference on Communication Systems and Networks (CSN 07). pp. 76–83 (2007). http://hdl.handle.net/10198/1988.
16. ITU: P.10 : Vocabulary for performance, quality of service and quality of experience. (2017). https://www.itu.int/rec/T-REC-P.10-201711-I/en.
17. Brunnström, K., Beker, S.A., Moor, K. de, Dooms, A., Egger, S., Garcia, M.-N., Hossfeld, T., Jumisko-Pyykkö, S., Keimel, C., Larabi, M.-C., Lawlor, B.: Qualinet white paper on definitions of quality of experience (2013) v1.2. European Network on Quality of Experience in Multimedia Systems and Services (COST Action IC 1003), Lausanne, Switzerland (2013). https://hal.archives-ouvertes.fr/hal-00977812/document.
18. Binsahaq, A., Sheltami, T.R., Salah, K.: A Survey on Autonomic Provisioning and Management of QoS in SDN Networks. IEEE Access. 7, 73384–73435 (2019). https://doi.org/10.1109/ACCESS.2019.2919957.
19. Karakus, M., Durresi, A.: Quality of Service (QoS) in Software Defined Networking (SDN): A survey. J. Netw. Comput. Appl. 80, 200–218 (2017). https://doi.org/10.1016/J.JNCA.2016.12.019.
20. Le Faucher, F., Lai, W.: Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering. (2003). https://doi.org/10.17487/rfc3564.
21. Mendiola, A., Astorga, J., Jacob, E., Higuero, M., Urtasun, A., Fuentes, V.: DynPaC: A Path Computation Framework for SDN. Proc. - Eur. Work. Softw. Defin. Networks, EWSDN. 119–120 (2015). https://doi.org/10.1109/EWSDN.2015.77.
22. Torres, E., Reale, R., Sampaio, L., Martins, J.: A SDN/OpenFlow Framework for Dynamic Resource Allocation based on Bandwidth Allocation Model. IEEE Lat. Am. Trans. 18, 853–860 (2020). https://doi.org/10.1109/TLA.2020.9082913.
23. Lai, W., Le Faucheur, F.: Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering. (2005). https://doi.org/10.17487/rfc4125.
24. Ash, G.: Max Allocation with Reservation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering \& Performance Comparisons. RFC Editor (2005). https://doi.org/10.17487/RFC4126.
25. Varga, A.: A Practical Introduction to the OMNeT++ Simulation Framework. EAI/Springer Innov. Commun. Comput. 3–51 (2019). https://doi.org/10.1007/978-3-030-12842-5_1.
26. Agarwal, S., Kodialam, M., Lakshman, T. V.: Traffic engineering in software defined networks. Proc. - IEEE INFOCOM. 2211–2219 (2013). https://doi.org/10.1109/INFCOM.2013.6567024.
27. Barradas, A.L., Medeiros, M.C.R.: An OMNeT++ Model for the Evaluation of OBS Routing Strategies. In: Proceedings of the First International ICST Conference on Simulation Tools and Techniques for Communications Networks and Systems. ICST (2008). https://doi.org/10.4108/ICST.SIMUTOOLS2008.3017.