Enduring Node Failures through Resilient Controller Placement for Software Defined Networks

Maryam Tanha, Dawood Sajjadi, and Jianping Pan Department of Computer Science, University of Victoria, Victoria, BC, Canada

GLOBECOM, Dec. 2016

By Sedef Savas Networks Lab, April 14, 2017

UCDAVIS

1

Control Plane Mapping Problem

For high resiliency \rightarrow solve controller placement (both number & placement.)

CP affects all of the resilience disciplines

- survivability (as a superset of fault tolerance),
- dependability (as a superset of reliability),
- security,
- performability,
- traffic tolerance, and
- disruption tolerance.

Focus is on the controllers' failures while taking into some factors which are missing in existing works such as **flow setup latency**, **switches' demands**, and **the capacity of the controllers**.

Related Work

A. Controller Placement

Round-trip propagation latency, calculated by using the length of the shortest path between s-c

- Facility Location Problem (FLP)
- Capacitated k-center problem: positioning the controllers to minimize the propagation latency while considering capacity of the controllers and demand of the switches.
- Pareto-based Optimal COntroller (POCO) placement framework provides pareto-optimal placements with regard to different measures, including switch-controller latency, controller-controller latency and controller load imbalance



Related Work (cont.)

B. Resilient Controller Placement

Mainly concerned with resilient controller placement to improve the resilience of the southbound connections (controller-switch connections) in SDN, resilient controller placement to enhance the resilience of the control plane (dealing with controller node failures).

Fault Tolerant Controller Placement (FTCP) problem to achieve high south- bound reliability. In the proposed formulation, each switch is required to satisfy a reliability constraint in a way that the operational route to any of its connected controllers remains with at least a given probability. The simulation outcome of applying the proposed heuristic algorithm to several network topologies, demonstrated that being connected to **two controllers** suffices for each switch.[*]

[*] F. J. Ros and P. M. Ruiz, "Five Nines of Southbound Reliability in Software-defined Networks," in *Proceedings of the Third ACM SIG- COMM Workshop on Hot Topics in Software Defined Networking*, 2014, pp. 31–36.

Problem Definition

To meet the resilience constraints as well as to address the performance (mostly related to the propagation latency), the cost, and capacity limitation.

- Cost limitations are associated with the number of required controllers or having a budget in terms of the the number of controller instances, and the inherent cost of deployments (e.g., CAPEX and OPEX).
- Also, the capacity constraints assist in dealing with the load on a controller (CPU, memory, and access bandwidth).

If controller becomes overloaded, processing latency will go up, affect the latency between a switch and the controller (it becomes a non-negligible part of the total latency).

Moreover, overloaded controllers have a higher probability of failure [6].



Assumptions

All nodes are suitable for controller deployment.

Switches have certain load, controllers have certain capacity. All controllers have a uniform failure probability and they fail independently.



Multi-level backup controller list

Objective: minimize the switch- controller re-assignment costs after the possible controller(s)' failures. In RCP, r denotes the resilience level at which a controller serves a given switch.

For instance, r = 0 indicates a primary assignment of a controller to a switch, and r = 1 denotes the assignment of the first backup controller to a switch.



To do capacity management after the reassignment – reassigns every switch.



Performance Evaluation

Network topologies of US continental tier- 1 service providers obtained from the Internet Topology Zoo.

Point of Presence (PoP)-level topologies are Sprint, ATT North America (two maps, one before 2008 and another in 2008), PSINeT (now part of Cogent Communications), and UUNET (now part of Verizon Business).

	Sprint	ATT NA (1)	PSINet	ATT NA (2)	UUNET
Network size	11	12	24	25	42
Links	18	21	25	56	77
Average node degree	3.27	3.5	2.08	4.48	3.6

TABLE I: Information about the chosen network topologies.



Cost of a Controller

The higher the degree of a potential node (location) for the controller, the less the cost is.

• nodes with better connectivity are reachable from more nodes, and placing the controllers on such nodes most probably decreases the cost in terms of the number of controllers.

In-band control, i.e., no dedicated links between controllers and switches for control traffic. Also, shortest-path length between a switch and a controller is the sum of the propagation latencies of all the links along the path.

3 scenarios

1) homogeneous switches (i.e., switches with homogeneous demands equal to 500 kilo req/s) and controllers (i.e., controllers with homogeneous capacities equal to 5, 000 kilo req/s). The probability of the node failure (i.e., pf) is assumed to be randomly chosen from [0.01 0.25]

2) The corresponding demands of the switches are configured with the minimum and maximum values of 150 kilo req/s and 1,000 kilo req/s (with step size of 50), respectively. The capacities of the controllers and pf are the same as the first scenario.

3) demands of the switches are uniformly distributed in [200 1,000] kilo req/s and controllers' capacities have values in [1, 800 8, 000] kilo req/s with fixed pf as 0.05.

Results

- 1) number of assigned controllers,
- 2) propagation latency between s-c,
- 3) controller locations at different resilience levels, and
- 4) distribution of the loads among the controllers.



Results (1) - # of assigned controllers



Fig. 1: Average number of the required controllers in each topology for all the three scenarios.

It can be seen that having a higher resilience level is more cost-effective (in terms of the number of required controllers) for the UUNET which has larger network size as well as more redundant paths and higher node degree TABLE II: Required number of controllers as a percentage of network size for m = 2.

Topology	Scenario 1	Scenario 2	Scenario 3
Sprint	36%	27%	40%
ATT NA (1)	33%	41%	35%
PSINet	33%	33%	28%
ATT NA (2)	32%	32%	30%
UUNET	30%	26%	33%

Regardless of the assigned probability of failure to the nodes, the number	o
required controllers are quite similar in the first and second scenarios.	

TABLE I: Information about the chosen network topologies.

	Sprint	ATT NA (1)	PSINet	ATT NA (2)	UUNET
Network size	11	12	24	25	42
Links	18	21	25	56	77
Average node degree	3.27	3.5	2.08	4.48	3.6

Considering the highest resilience level (i.e., m = 2), the maximum propagation latencies for all of the topologies are below 50 ms, which is far less than the latency threshold and leaves the room for other contributors of the flow-setup latency, including transmission delay, processing delay and probably the delay incurred by congestion in the network.



Results (2)

Fig. 2: CDF of the propagation latency for all the topologies in scenarios 2 and 3.

Result (2) cont.



Fig. 3: Maximum propagation latency for scenario 3 (ATT NA (2)).

In most cases, when the resilience level is increased (e.g., from m = 0 to m = 1), the maximum latency between a switch and its assigned controller in a topology goes up.

But this is not always the case; for scenario 3 and ATT NA (2), the maximum latency decreases when the resilience level increases from 1 to 2 in the second experiment.

Results (3) - Controller location



Fig. 4: Controller locations for Sprint with m = 2 in scenario 2.

ł

Kansas City (common in all scenarios) has a strategic location in most of topologies since it connects east and west of the US in the maps.

Due to their higher connectivity (higher node degree) and subsequently, better reachability from other nodes. This is reflected by our assumed location-dependent deployment cost (fc) which increases in inverse proportion to the node degree.

Results (4) – controller load distribution

While increasing the resilience level results in more load imbalance for some topologies in different scenarios, it leads to less load imbalance for the others. This again confirms the reliance of the solution on the network topology.

	Resilience	Sprint	ATT NA (1)	PSINet	ATT NA (2)	UUNET
Scenario 1	m=0	(3, 5, 3.66, 1.15)	(5, 7, 6, 1.41)	(4, 10, 8, 3.56)	(7, 10, 8.33, 1.52)	(5, 10, 8.40, 2.30)
	m=1	(5, 10, 7.33, 2.51)	(7, 10, 8, 1.73)	(8, 10, 9.60, 0.89)	(10, 10, 10, 0)	(6, 10, 9.22, 1.30)
	m=2	(5, 10, 8.25, 2.36)	(6, 10, 9, 2)	(4, 10, 9, 2.13)	(8, 10, 9.37, 0.74)	(7, 10, 9.69, 0.85)
Scenario 2	m=0	(3, 8, 5.50, 3.53)	(2, 10, 6, 5.65)	(4, 11, 8, 3.60)	(7, 10, 8.33, 1.52)	(5, 13, 10.50, 3.78)
	m=1	(11, 11, 11, 0)	(5, 10, 8, 2.64)	(6, 13, 9.60, 2.70)	(8, 12, 10, 2)	(6, 14, 10.50, 2.67)
	m=2	(11, 11, 11, 0)	(3, 9, 7.20, 2.49)	(4, 13, 9, 2.56)	(8, 12, 9.37, 1.92)	(8, 14, 11.45, 2.11)
Scenario 3	m=0	(3, 8, 5.50, 3.53)	(2, 7, 4, 2.64)	(4, 10, 8, 3.46)	(3, 9, 6.25, 2.50)	(3, 12, 7, 3.34)
	m=1	(3, 10, 7.33, 3.78)	(4, 12, 8, 4)	(4, 16, 12, 5.65)	(3, 14, 8.33, 3.72)	(4, 16, 9.33, 3.74)
	m=2	(4, 11, 6.60, 2.96)	(5, 12, 9, 2.94)	(7, 18, 14.40, 4.27)	(3, 14, 9.37, 3.50)	(4, 15, 9.69, 3.72)

4-tuple shows the minimum, maximum, mean and standard deviation for the number of switches managed by the controllers in a given topology.

Conclusion

Resilient controller placement problem, takes into account the capacity of the controllers as well as the demands of the switches.

Minimizing the total cost (including the **cost of deployment, the propagation latency**, **and the number of required controllers**) achieved while considering different resilience levels to enhance the resilience of the controller plane.

Future research directions involve designing heuristic/approximation algorithms to deal with large network sizes as well as testing more random/synthetic topologies to gain useful insights on the results.