SafeGuard: Congestion and Memory-aware Failure Recovery in SD-WAN

Meysam Shojaee[*], Miguel Neves[‡] and Israat Haque[*]
[*]Dalhousie University
[‡]UFRGS University

Outline

1. SD-WANs and network failure
2. Failure recovery challenges
3. SafeGuard’s model and architecture
4. Results analysis
5. Conclusion

SDN deployment examples

- Large service providers such as Amazon, Facebook, Google, and Microsoft are deploying SDN in their WAN [1].

B4: Google’s private SD-WAN [2]

- Network failures are common in wide area networks [3].
- 80% of the network component failures last from 10 to 100 minutes which leads to intensive packet loss [4].

Why failure recovery is challenging is SD-WAN?

- Providing short backup routes **conflicts with** making a balanced link utilization

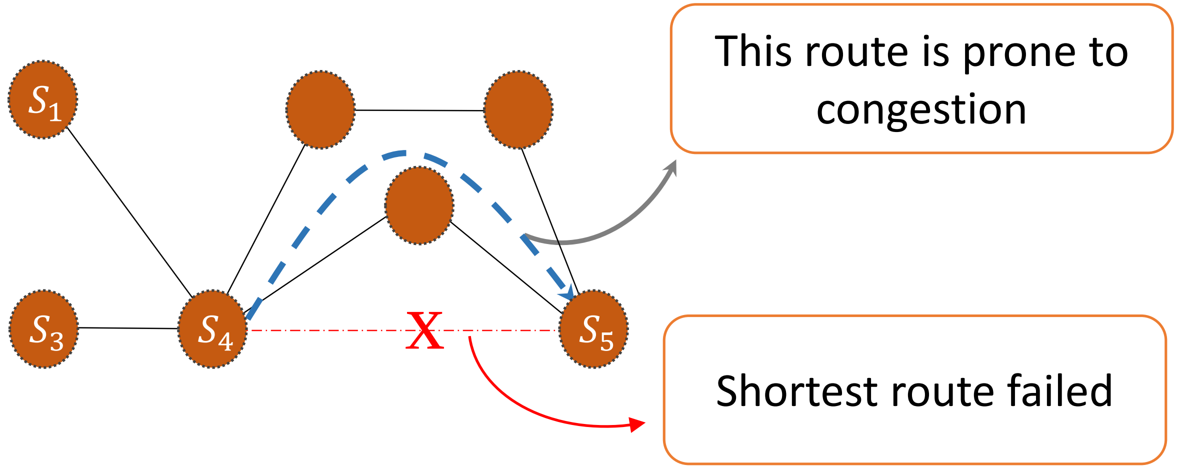

- Bandwidth capacity of links are limited
- The memory of the switches are limited
- The route and traffic rate need to be configured in small Traffic Engineering (TE) cycles

- TCAMs are well-known to be power-hungry and to have limited capacity [5]

Problem formulation

- We formulated the failure recovery problem as a multi-objective MILP optimization problem

Objectives
- Minimizing the maximum link utilization
- Minimizing the backup route length

- Objective function:

$$\min \sum_{f\in F_e} \sum_{p\in P'_f \cup Q_f} (\underbrace{\alpha x_{fp} l_p}_{1} + \beta \underbrace{\sum_{(i,j)\in E} \frac{r_{ijp} y_{fp}}{c_{ij}}}_{2})$$

<u>Term 1</u>: The length of the backup route

<u>Term 2</u>: The load on each link

Constraints of the designed model

- We have considered a number of constrains in our model:

1. Link bandwidth capacity

2. Switch memory capacity

3. Flow satisfaction

- It takes more than one hour to solve the model in CPLEX
- The model is an MILP problem which is known to be NP-hard
- We developed a heuristic to solve it in a reasonable time

Operation of the heuristic

Key ideas of the heuristic:

1) Assigns **the larger flows** to the **shorter** routes **first**
2) When get to the memory usage threshold: best-fit strategy is applied

Both flows are affected by the failure

- f_1 from switch S_1 to S_5 has a demand of 25Mbps
- f_2 from switch S_2 to S_5 requires 20Mbps
- The remaining capacity of all links after the allocation of the flows: 20Mbps except link S_4 –S_5 whose capacity is 5Mbps

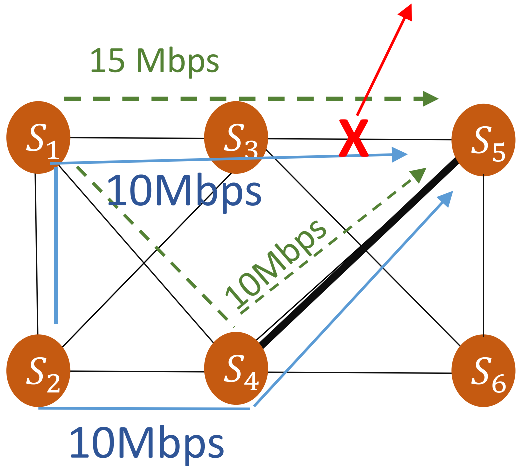

Operation of the heuristic

- routes are allocated to f_1 with higher demand

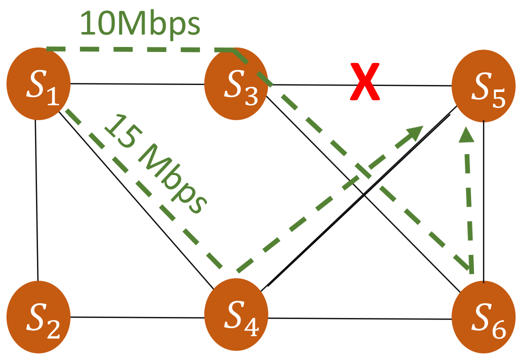

- routes are allocated to f_2 after f_1

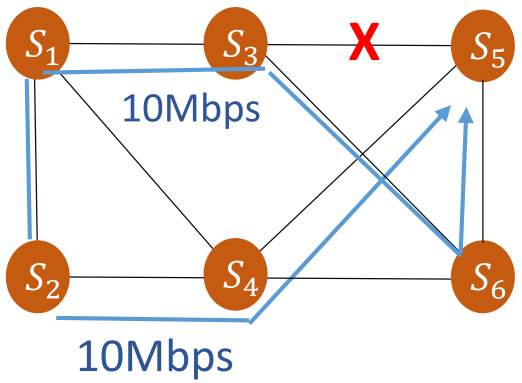

SafeGuard’s architecture

- **Input**: a traffic matrix, the network topology, and a set of forwarding routes (primary routes)

- Produces a primary allocation that will be used for forwarding traffic under normal conditions

- SafeGuard then applies the heuristic to augment the primary allocation with backup routes for all flows for each possible single link failure.

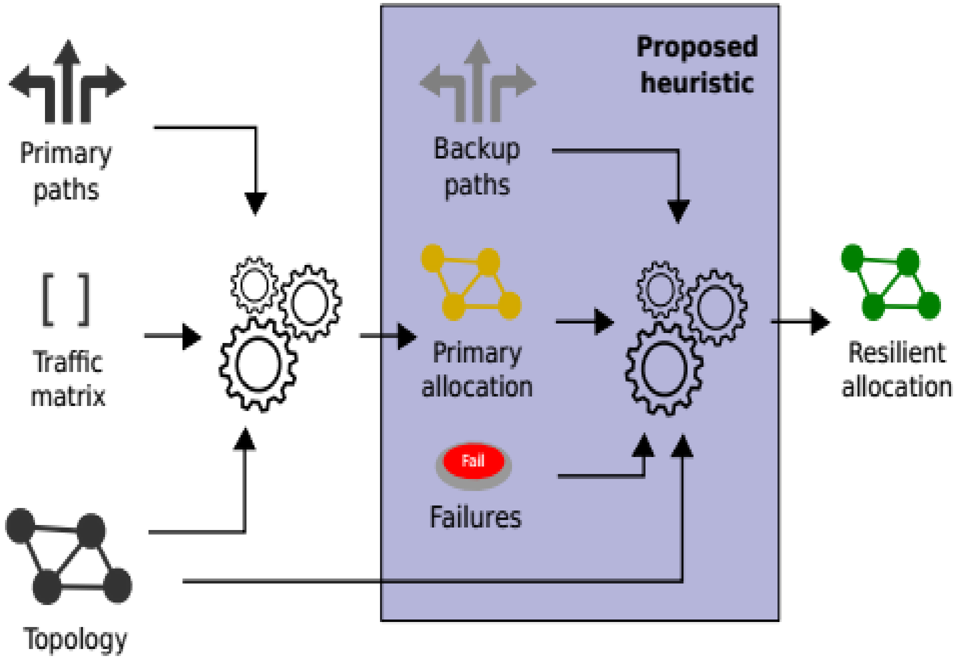

SafeGuard’s work-flow

Step A: SafeGuard computes and installs all primary tunnels and splitting weights for every flow

Step B: proactively installs backup routes and computes the weights for allocated flows

Step C: When a failure happens:

- the failing switch activates the corresponding backup tunnels
- sends a message reporting the failure to the network controller

Step D: the network controller adjusts splitting weights for all affected flows at their respective ingress switches.

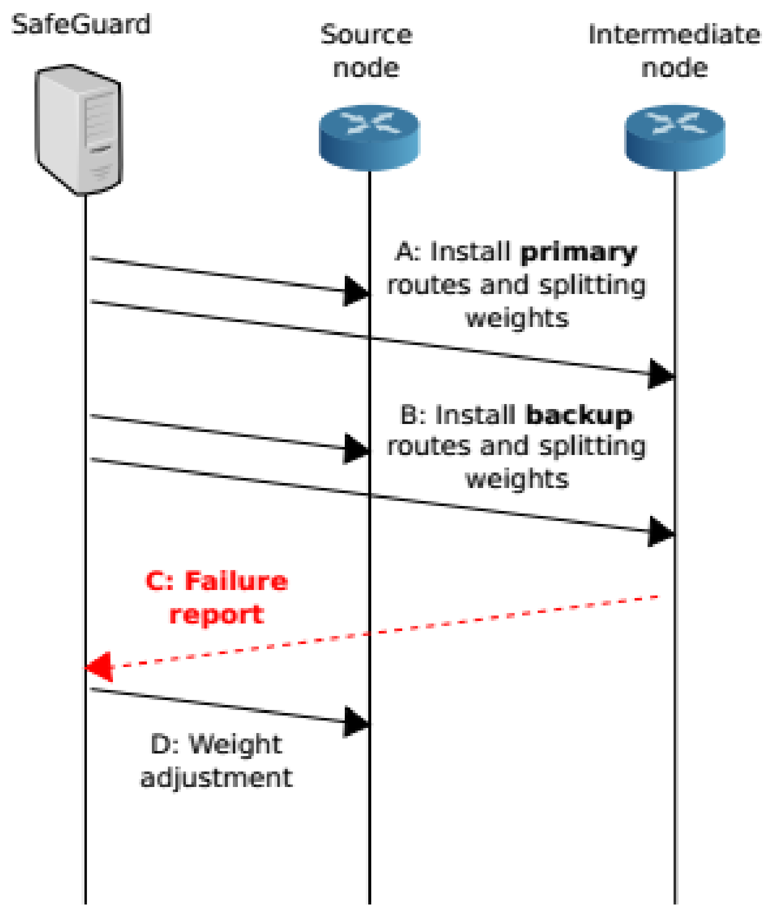

Evaluation setup

Implementation:

- We implemented a prototype of SafeGuard using the Ryu SDN controller
- Each node is implemented as a CpqD switch instance
- Link capacities are set to 1 Gbps

https://github.com/Meysam-Sh/SafeGuard

Network topologies:

- We considered two networks: B4 with 12 nodes and ATT with 25 nodes

Data generation:

- We used Iperf tool to generate UDP traffic
- We randomly failed a link in the network

Results

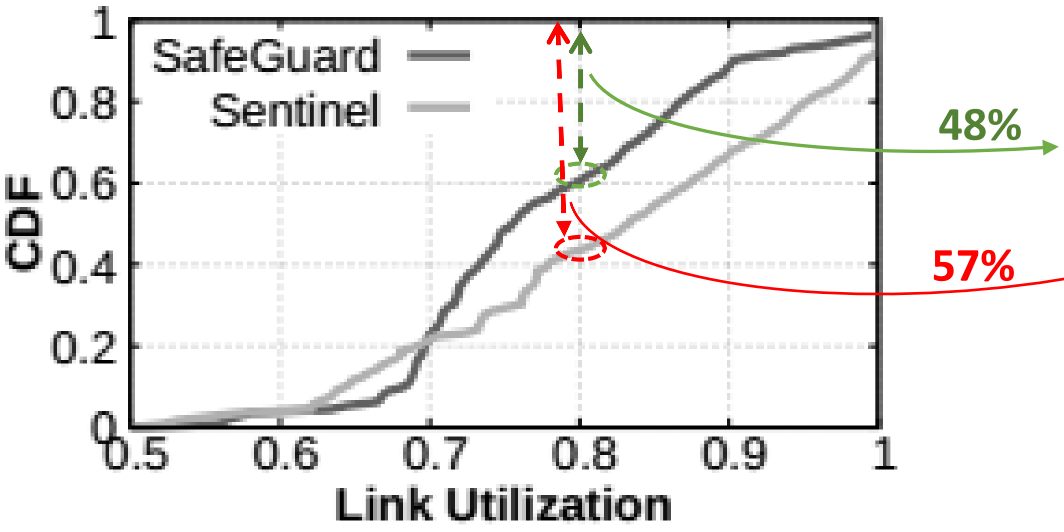

❑ SafeGuard: 48% of the links are under 80%+ load

❑ Sentinel: 57% of links are under 80% + load

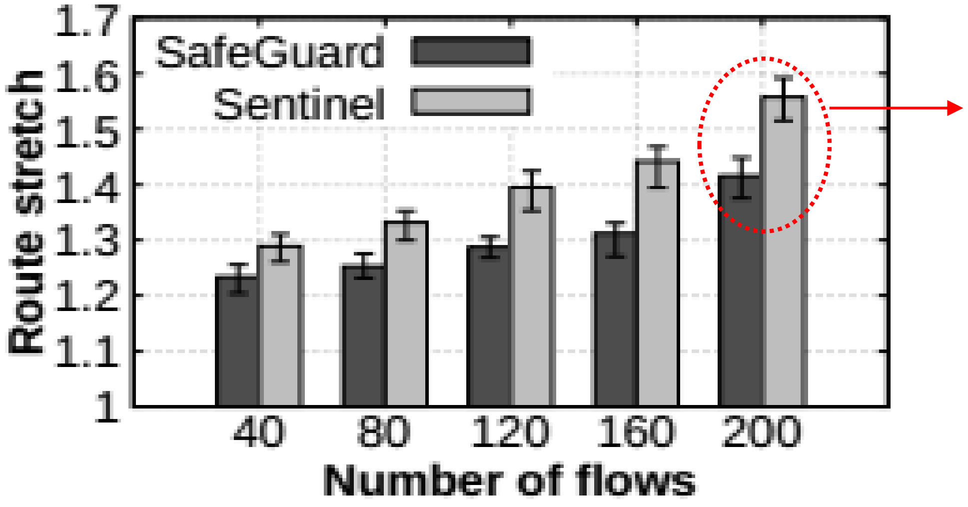

❑ SafeGuard results in backup routes around 10% shorter than Sentinel

Conclusion

- ❖ SDN is widespread in the production networks. The link failures are a common occurrence in SD-WANs.
- ❖ We formulated the failure recovery problem as a multi-objective MILP optimization problem.
- ❖ We designed and developed a heuristic as the problem was NP-hard.
- ❖ We implemented a prototype of SafeGuard using the Ryu SDN controller and evaluated it in Mininet.
- ❖ Our results show that SafeGuard can reduce the number of congested links by up to 50% compared to the state-of-the-art failure recovery scheme.

Any Question?