THE WEB CONFERENCE ACM

Research Track | Systems and Infrastructure for Web, Mobile, and WoT

ARES: Predictable Traffic Engineering under Controller Failures in SD-WANs

Songshi Dou†‡ §, Li Qi‡ §, and Zehua Guo‡ * §

†The University of Hong Kong, ‡Beijing Institute of Technology, § Zhejiang Lab,

*Beijing Institute of Technology Zhengzhou Academy of Intelligent Technology

之江实验室 ZHEJIANG LAB

① Introduction

- ***TE Meets Various QoS Requirements for Web Services through SD-WANs***
  - Cloud services bring new web applications:
    - web services, video streaming, web AR/VR …
  - Web services account for *a large share of* Wide Area Network (WAN) traffic.
  - Traffic Engineering (TE) improves the network performance of WANs and *enables differentiable QoS* for web applications.
  - SDN brings the *flexible management* of WANs since TE can be implemented at the SDN controller to update routing policies for accommodating potential traffic fluctuations.
- ***Controller Failures Pose Severe Impact on TE Performance***
  - Controller failure is *a common problem* in SD-WANs.
    - power outage, malicious attacks …
  - Switches previously controlled by the failed controller become *offline switches*, and flows traversing these offline switches *become offline and lose their path programmability*.
  - Network performance cannot be guaranteed to accommodate potential traffic variations due to the *loss of flexible network management*, leading to *significant TE performance degradation* with a large number of offline flows.
- ***State-of-the-Art Solutions and Their Limitations***
  - Reassign offline switches to other active controllers.
  - *Path programmability-centric* solutions.
  - While higher path programmability implies *a higher likelihood of accommodating unpredictable future traffic fluctuations*, it cannot guarantee a predictable TE performance.

② Observation & Motivation

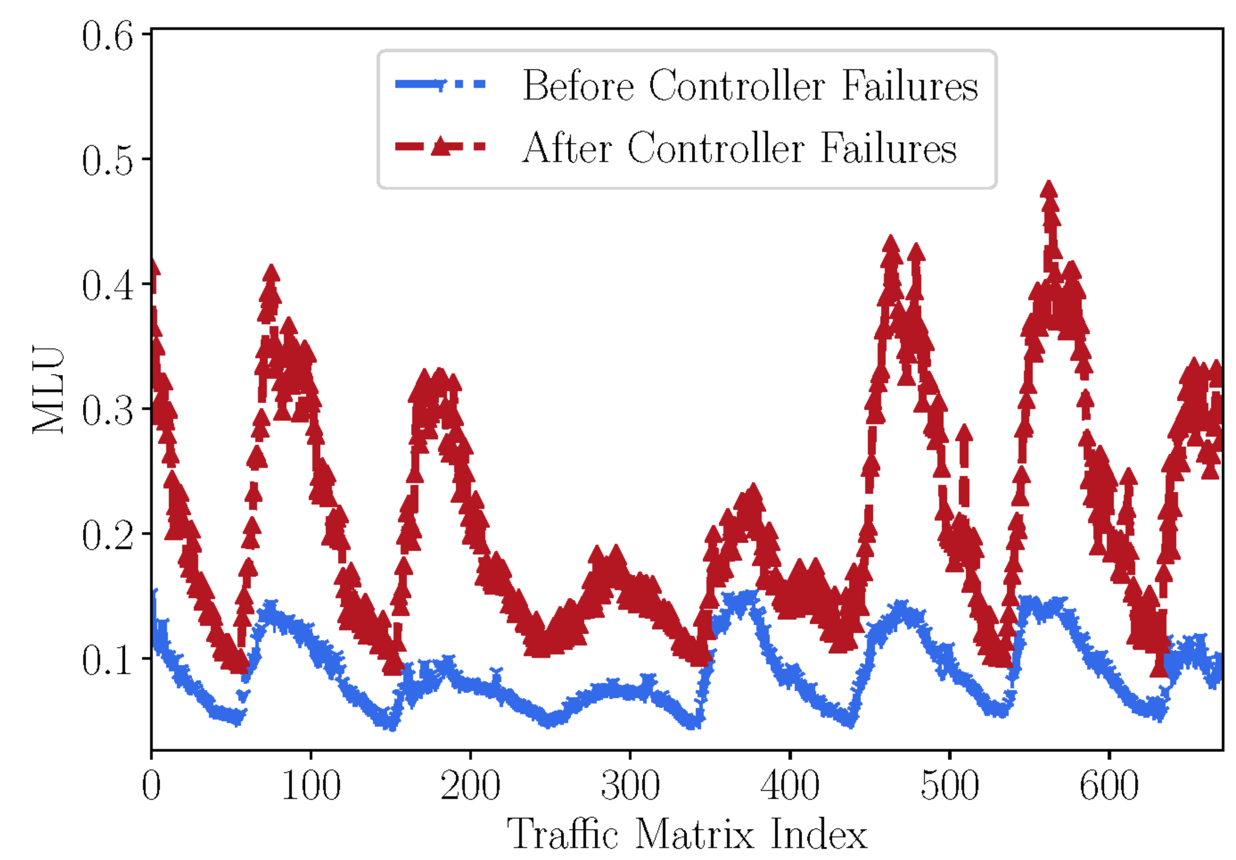

Comparison of the MLU performance before and after controller failures. The lower, the better.

- ***Observation: impact of controller failures on TE performance***
  - Doing TE operations with the objective of *minimizing the Maximum Link Utilization (MLU)*.
  - GEANT topology with real-world Traffic Matrices (TMs):
    - Topology: 23 switches, 72 links.
    - TMs: every 15 minutes, a total of 672 TMs during one week.
  - Blue line: the MLU performance *before* controller failures
  - Red line: the MLU performance *after* controller failures.
  - This observation demonstrates that controller failures *threaten TE performance*, which can *increase the MLU by up to 0.35* in the worst case.
  - The root cause lies in that *flows cannot be flexibly rerouted* to accommodate traffic fluctuations due to controller failures.

③ Solution - ARES

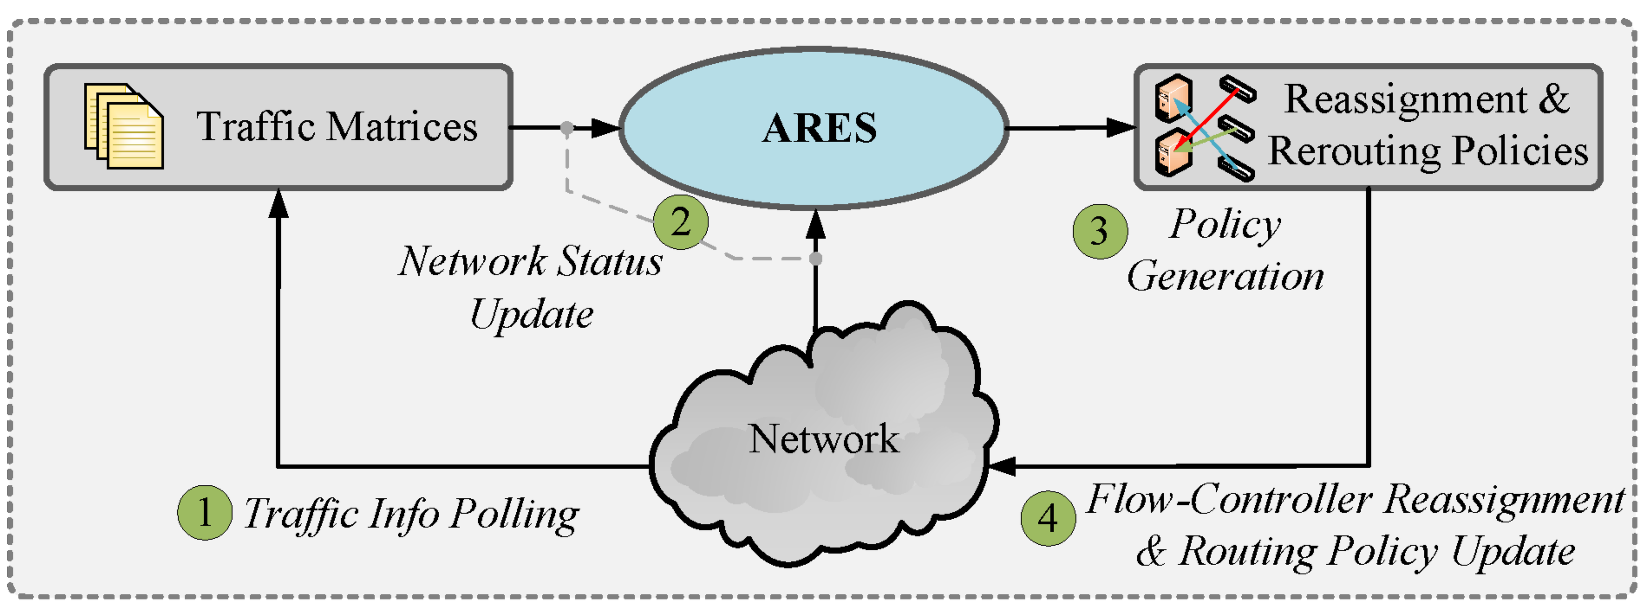

The processing logic of ARES.

- ***Opportunity: fine-grained flow-controller reassignment***
  - P4 Runtime specification is designed for P4 programmable switches and *allows multiple controllers to manage the switch*.
  - This feature *aligns with the design goal* of enabling multiple controllers to manage a single switch.
- ***Overview of ARES***
  - Firstly, it collects *real-time traffic traces* (e.g., TMs) from the network periodically.
  - Secondly, when controller failures occur, ARES *updates the current network status* with the required information (e.g., collected TMs, offline flows, and active controllers).
  - Subsequently, *flow-controller reassignment and flow rerouting policies* are determined and generated by solving the proposed optimization problem / heuristic algorithm (detailed in the paper).
  - Finally, ARES *reassigns offline flows* to corresponding active controllers and *updates routing policies*, achieving predictable load balancing performance in the whole network.

④ Evaluation Results

Scheme	One controller failure	Two controller failures
OSPF	44.46 +/- 0%	44.46 +/- 0%
ECMP	56.78 +/- 0%	56.78 +/- 0%
OPT-OSCM [23]◊	60.57 +/- 6.68%	55.68 +/- 6.97%
OPT-FRSM [24]◊	58.09 +/- 4.18%	51.40 +/- 5.44%
OPT-FMSSM [25]◊	58.50 +/- 4.44%	52.42 +/- 4.49%
OPT-TPFCRFR	**100.00 +/- 0%**	**100.00 +/- 0%**

Average PR Performance of proposed formulation and existing formulations.

Scheme	One controller failure	Two controller failures
OSPF	44.46 +/- 0%	44.46 +/- 0%
ECMP	56.78 +/- 0%	56.78 +/- 0%
RetroFlow♠	64.01 +/- 0%	63.70 +/- 0.25%
Matchmaker♠	57.46 +/- 4.05%	50.49 +/- 4.24%
PM♠	60.97 +/- 2.85%	51.44 +/- 4.18%
ARES	**87.82 +/- 1.69%**	**87.23 +/- 1.19%**

Average PR Performance of ARES and existing solutions.

- ***Evaluation results***
  - GEANT Topology: 23 switches, 72 links.
  - Real-world TMs: every 15 minutes, a total of 672 TMs.
  - *Performance Ratio* (*PR*): $PR = P_{scheme} / P_{optimal}$. A *PR* value of 1 implies that the scheme performs on par with the optimal results.
  - *Our problem formulation* exhibits comparable load balancing performance to optimal TE solution without controller failures, and *ARES* significantly improves average load balancing performance by up to 43.36% with low computation time, compared with baseline approaches.