TEAVAR: Striking the Right Utilization-Availability Balance in WAN Traffic Engineering

Jeremy Bogle, Nikhil Bhatia, Manya Ghobadi

Background

Utilization vs Availability

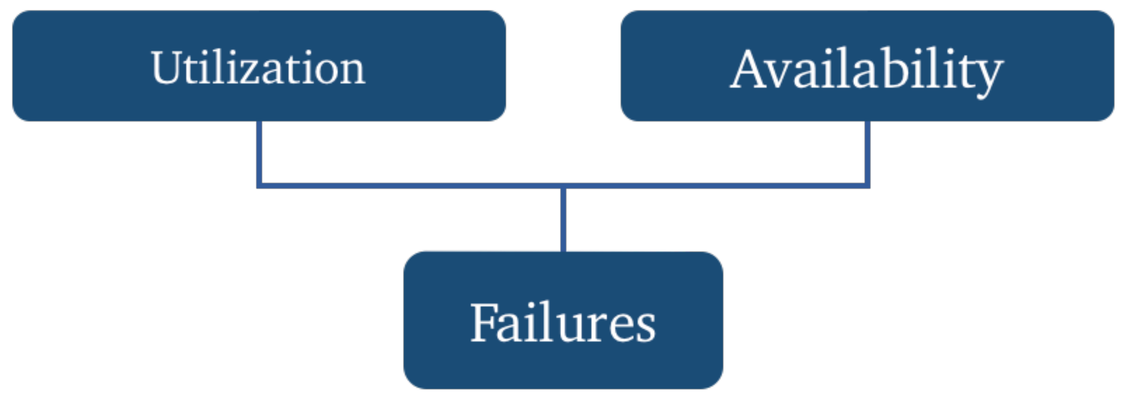

- Goal: Optimize WANs for high-performance cloud infrastructure.
- Challenge: striking a good balance between network *utilization* and *availability.*
- If we want to enable network operators to strike this balance, we must consider alternative ways of understanding risk management in traffic engineering (TE). We propose TEAVAR (Traffic Engineering Applying Value at Risk), a novel approach to TE that leverages empirical data to maximize bandwidth allocation subject to a probabilistic model of network failures.

Motivation

State of the Art

- State of the art maximizes bandwidth allocations to be robust against up to k concurrent link failures, i.e. support the maximum amount of traffic from *s* to *d* that is resilient to at most two *(k = 2)* link failures.

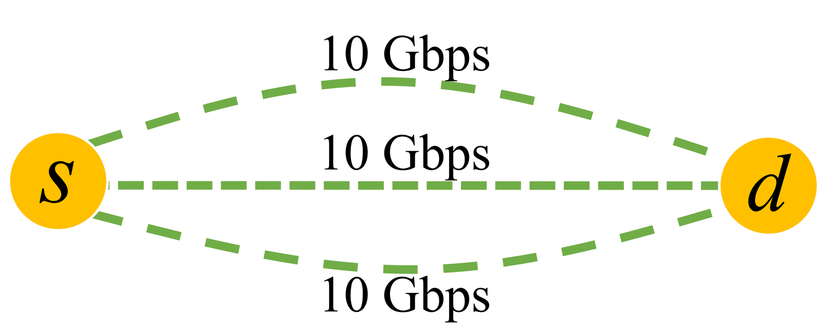

- Below is the optimal solution under FFC: rate-limit the (s,d)-flow to 10Gbps and always split traffic equally between all intact links. This solution guarantees a bandwidth of 10Gbps under up to two link failure scenarios.

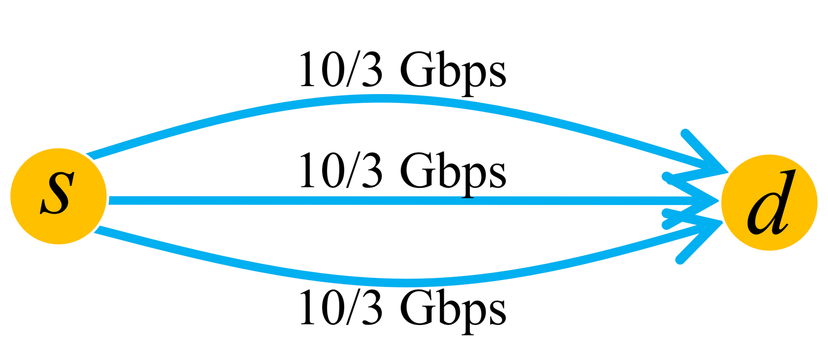

- However, this guarantee comes at the cost of keeping each link at one third utilization when no failures occur. Can we do better?

A Better Alternative

- We ask whether high availability can be achieved without such drastic over-provisioning. Suppose that the link failure probabilities are described in the figure below.

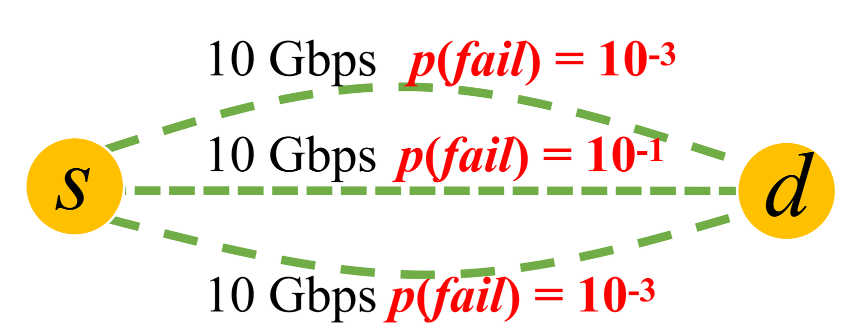

- Under these failure probabilities, the bandwidth assignment described below guarantees a total throughput of 20Gps at a level of availability of nearly 3 nines (99.8%)

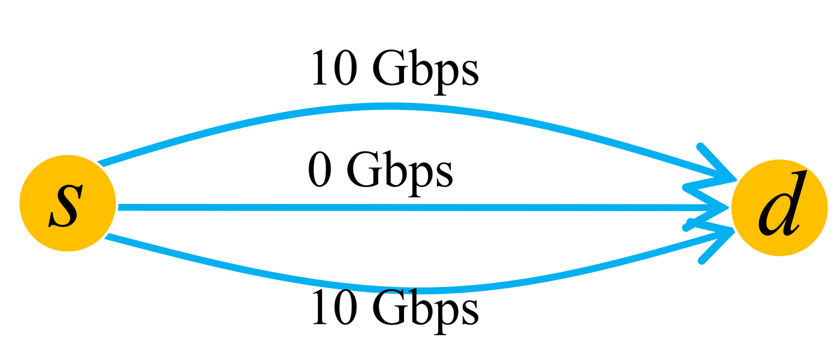

- This novel approach allows network administrators to find the right balance between network utilization and availability. But, how do we accomplish this?

Approach

Portfolio Optimization

- In portfolio optimization, we have a monetary resource pool that can be allocated across a number of stocks. Our goal as risk-averse grad students might be to minimize portfolio loss over daily market outcomes.
- Suppose we want to give a bound on the loss we might incur, such as "the loss will be less that $100 with probability 0.95." Value at Risk (VaR), a tool from financial risk theory, captures precisely this.

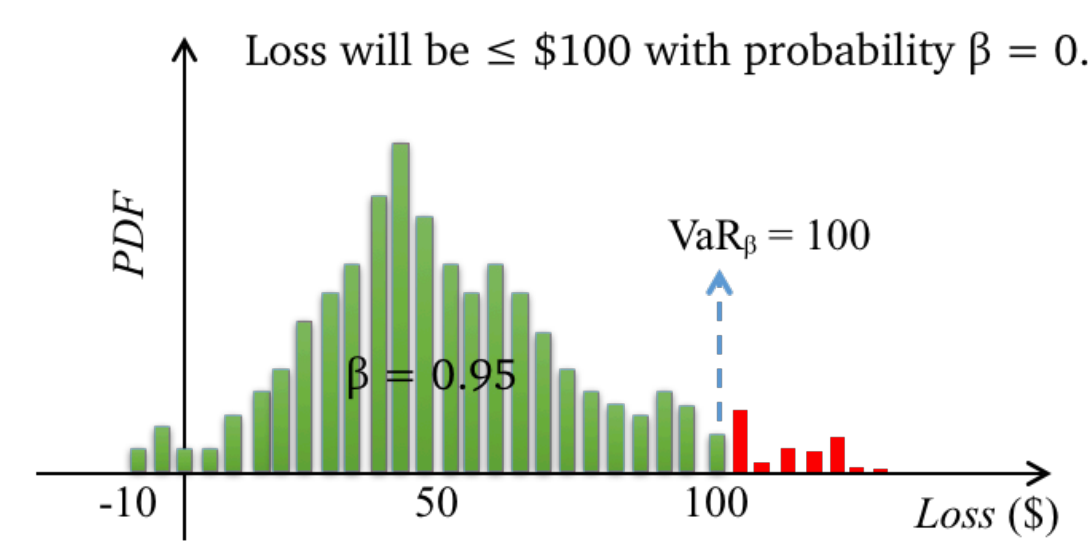

- Given a probability threshold β, VaR_β provides a probabilistic upper bound on the loss: "the loss is less than VaR_β with probability β.

TE Optimization

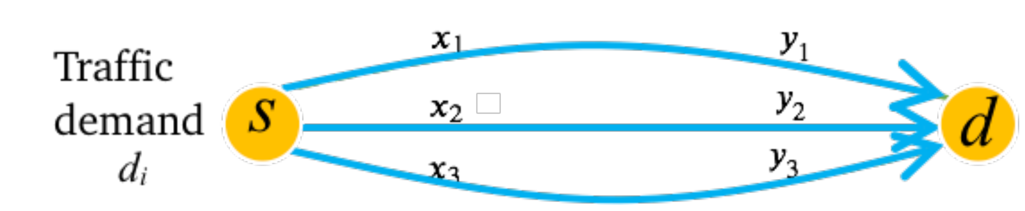

- In network optimization, we are allocating our total resource, demand, across tunnels in our network. We define loss in terms of the fraction of demand *not* satisfied, so as to maximize concurrent flow.

Scenario-level loss function:

$$Loss(x, y) = \max_i \left[1 - \frac{\sum_{r \in R_i} x_r y_r}{d_i}\right]$$

Applying conditional VaR to TE:

- **We define a probabilistic failure model using link failure rates, calculating possible failure scenarios and their corresponding probabilities of occurring.**
- **We then apply VaR by minimizing β-Conditional Value at Risk ($CVaR_\beta$) over these scenarios.**

Challenges of applying CVaR:

1. Achieving fairness across network users.
2. Capturing fast rerouting of traffic in data plane.
3. Achieving computational tractability.
4. Capturing correlated failures.

Results

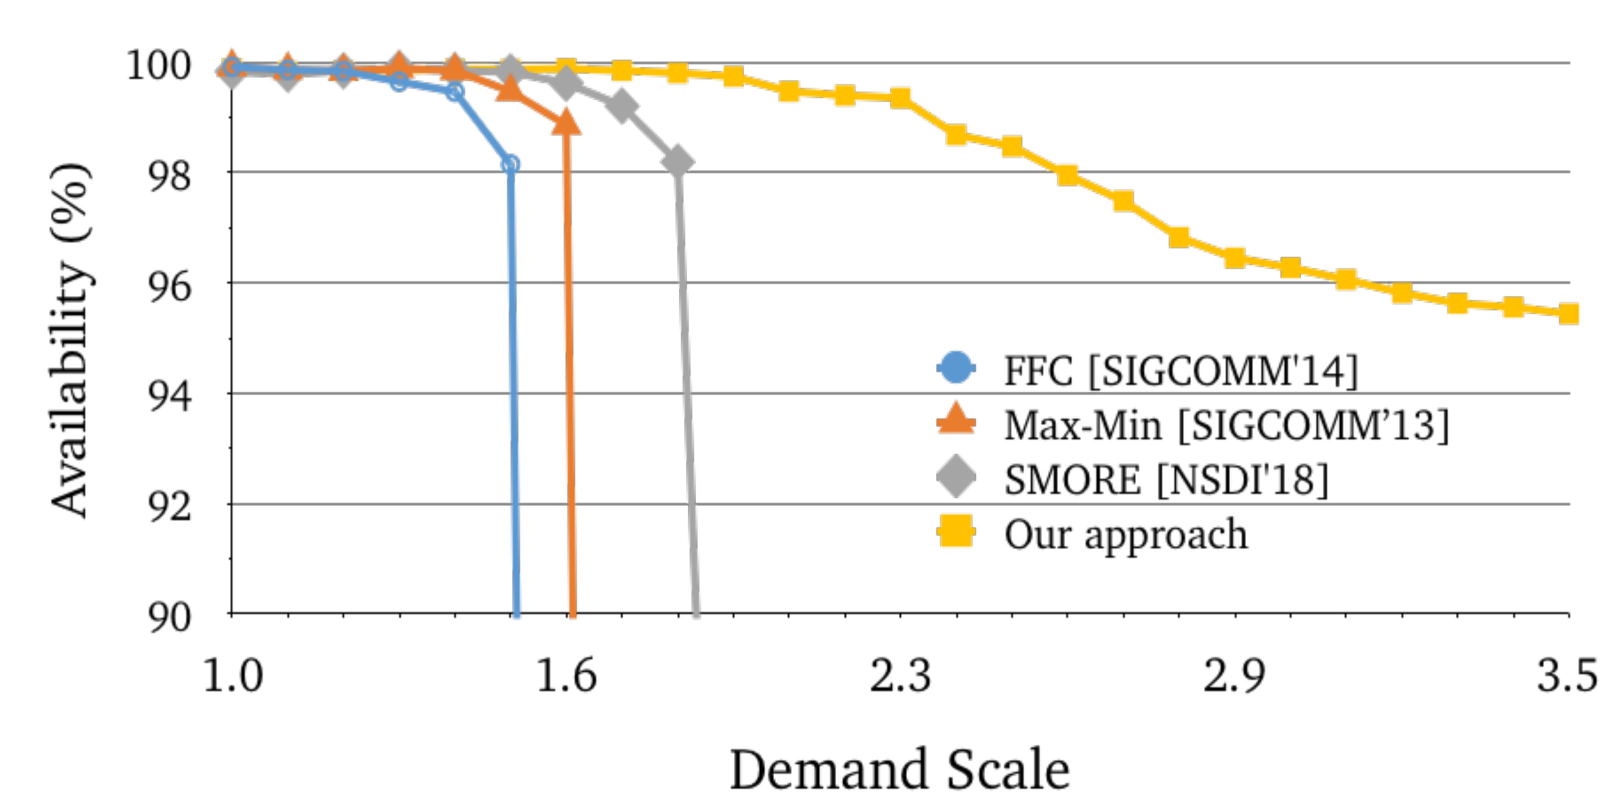

- We define the availability of a TE scheme as the percentage of simulated scenarios for which the scheme can *fully satisfy* demand.
- TEAVAR is capable of sustaining consistently higher availability for a given demand scale when compared to state-of-the-art TE schemes.