

# Boosting the BGP Convergence in SDXes in SWIFT

Philipp Mao, Rüdiger Birkner, Thomas Holterbach, Laurent Vanbever  
ETH Zürich

## 1. Motivation

- ◆ BGP converges slowly [1], for 2 reasons:
  - Learning about a failure is slow, especially for remote failures
  - Updating thousands of forwarding entries takes time
- ◆ Existing fast reroute solutions only partially solve the problem
  - using hardware-based signals
  - using a hierarchical forwarding table

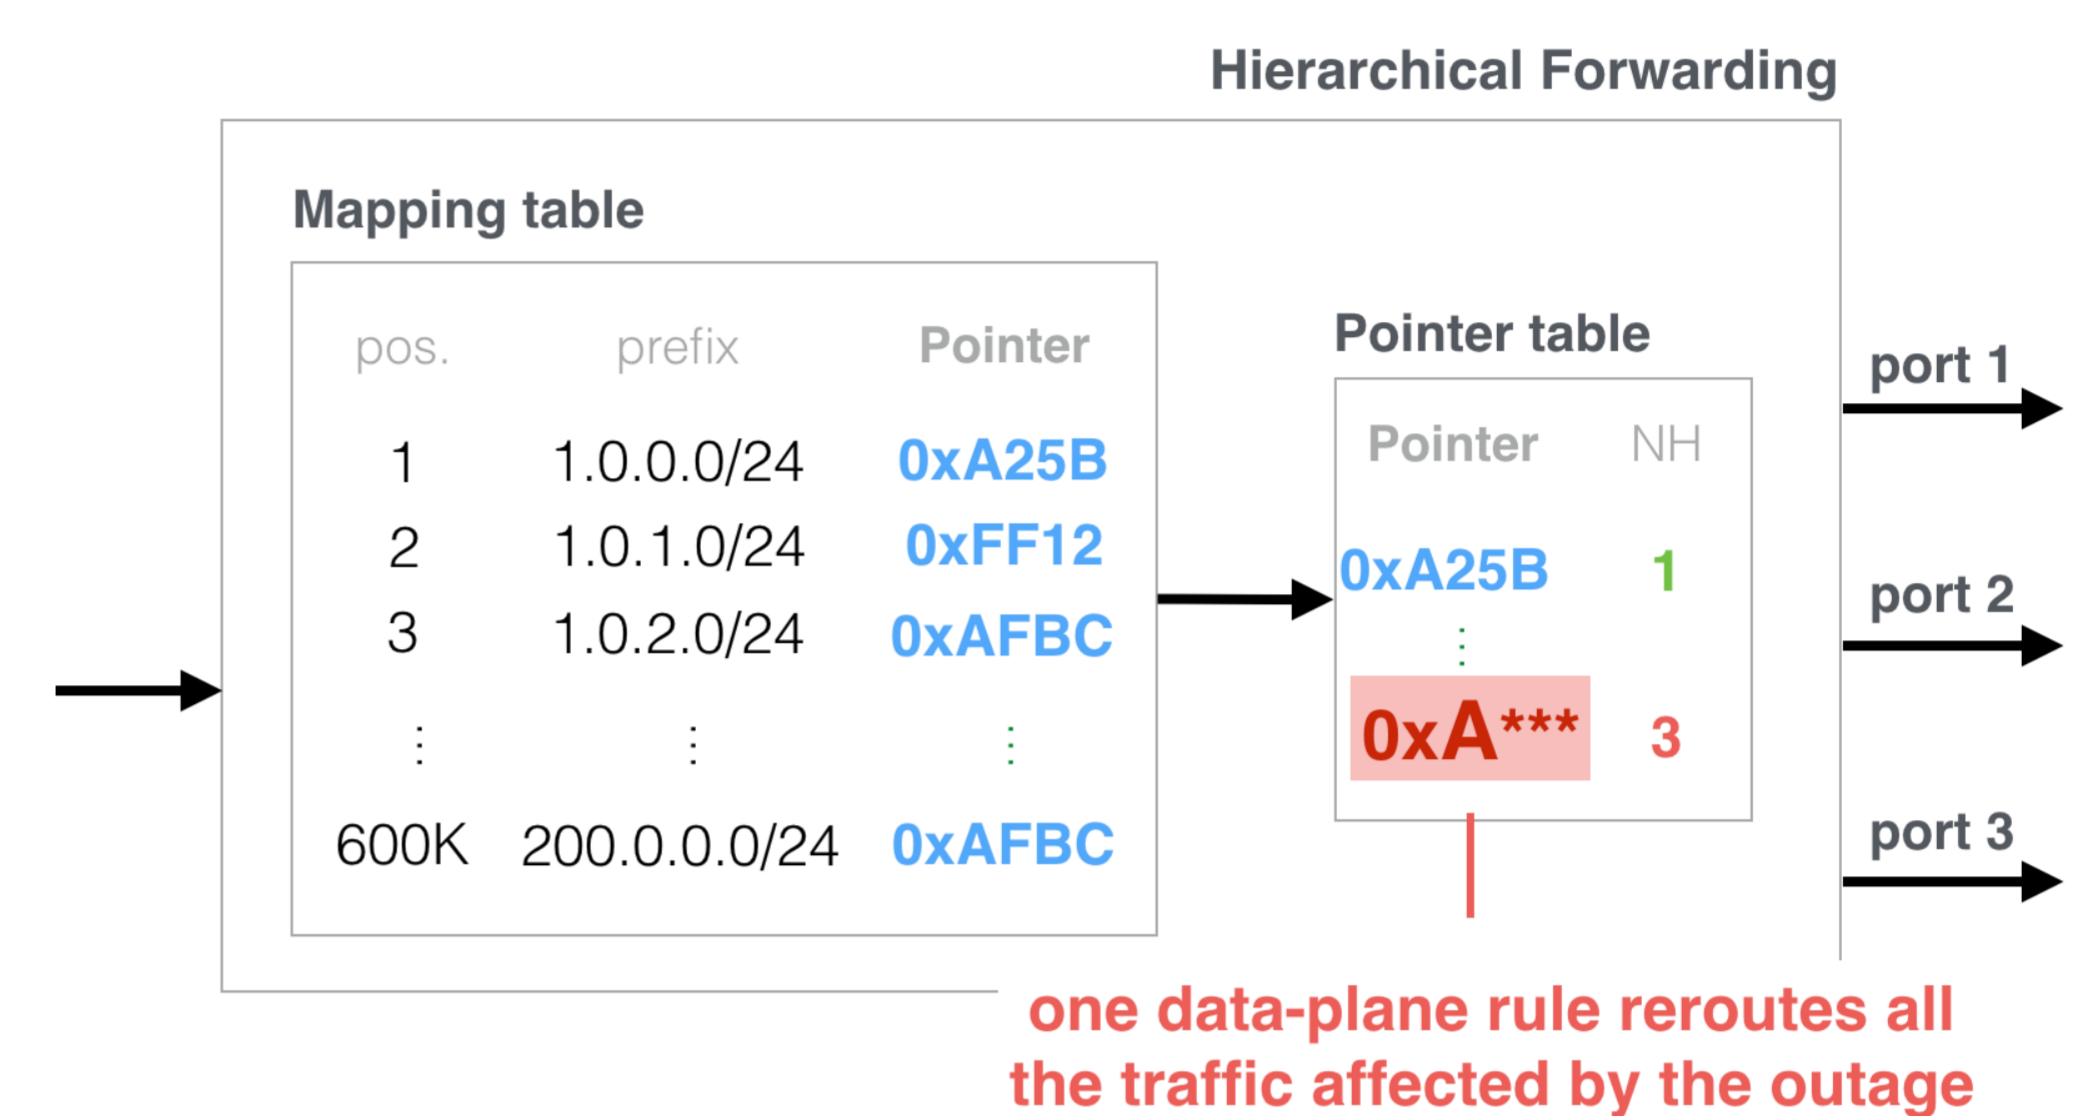
**But they only work upon local outages**

◆ SWIFT [2] is a new fast reroute framework that works upon local *and* remote outages

- ◆ We show that deploying SWIFT at SDXes can boost the convergence time of hundreds of networks
  - SDX and SWIFT architectures are compatible with each other

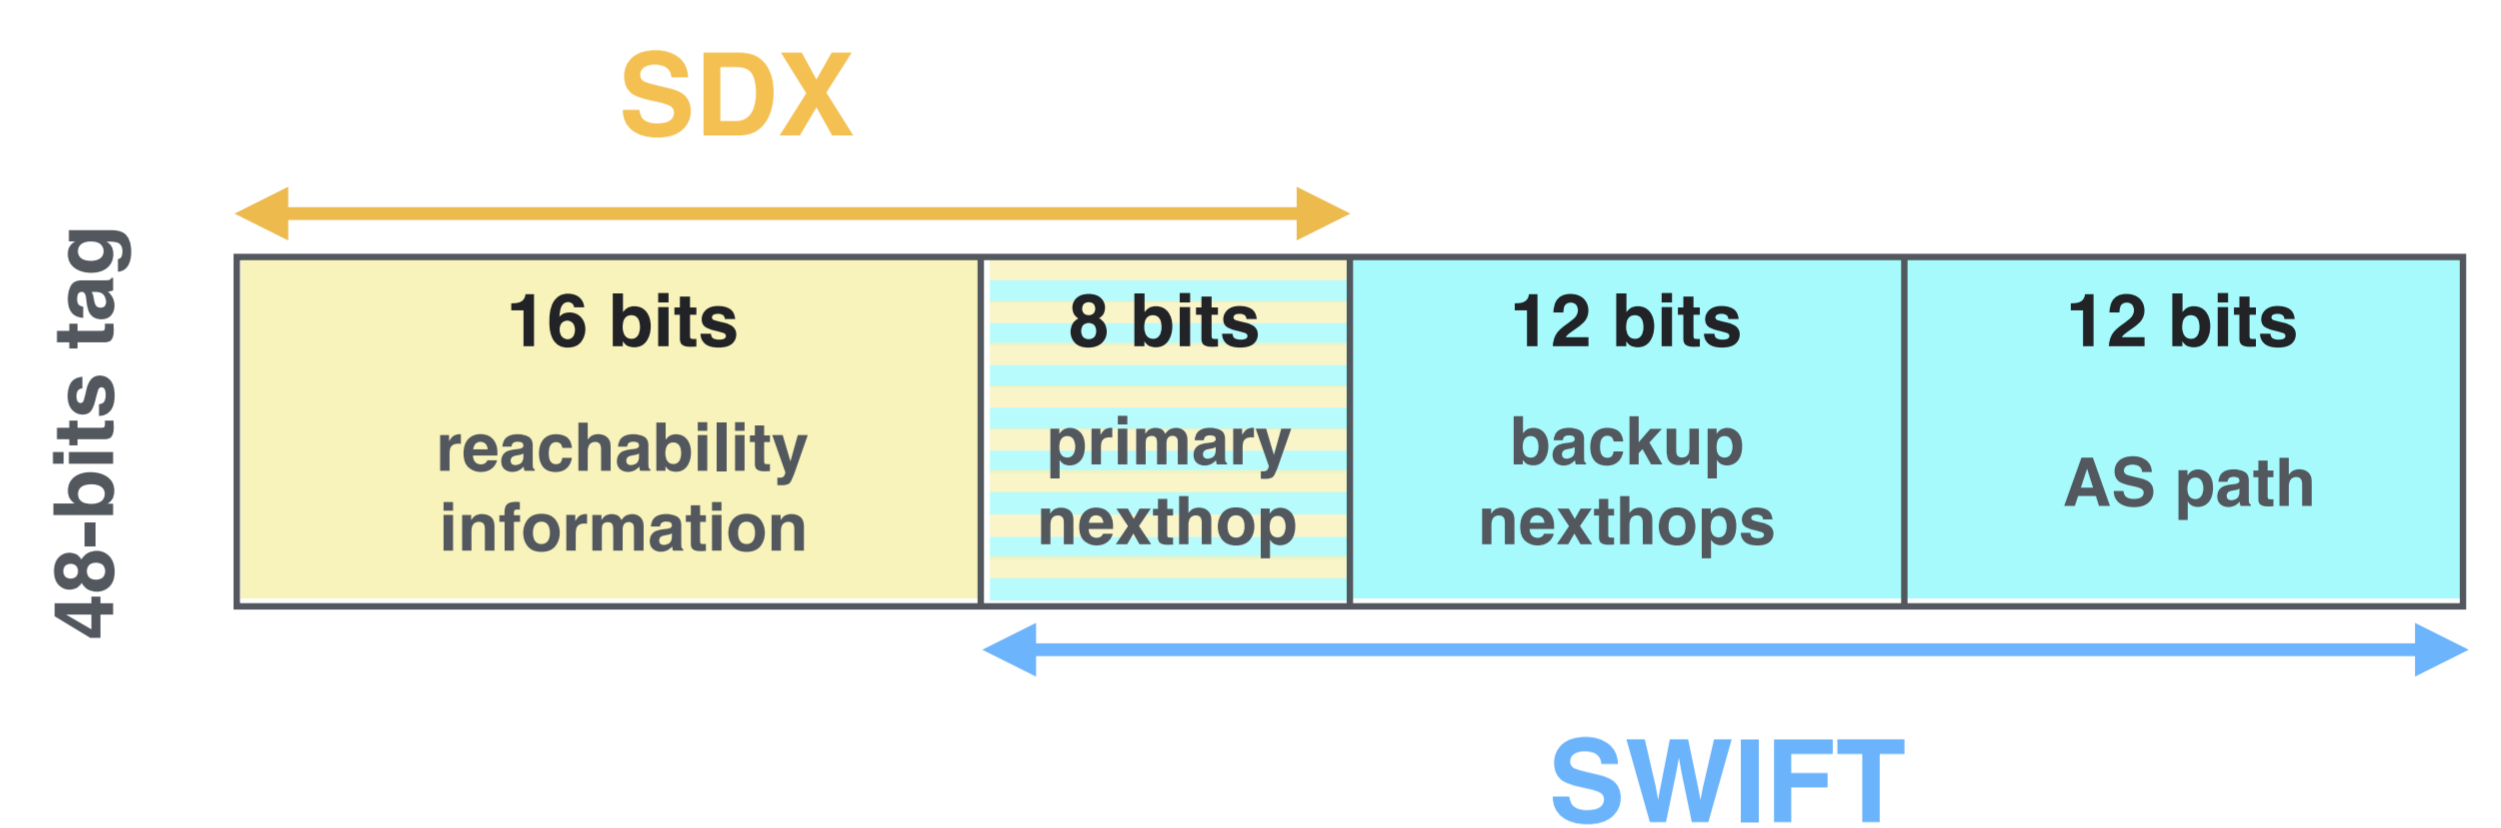
## 2. SWIFT in a nutshell

- ◆ SWIFT uses inference techniques to quickly **locate a failure**
  - using a subset of the control-plane messages
- ◆ SWIFT matches on preprovisioned dataplane tags to **quickly reroute traffic**
  - Using a 2-stage forwarding table



## 3. Integrating SWIFT into SDXes

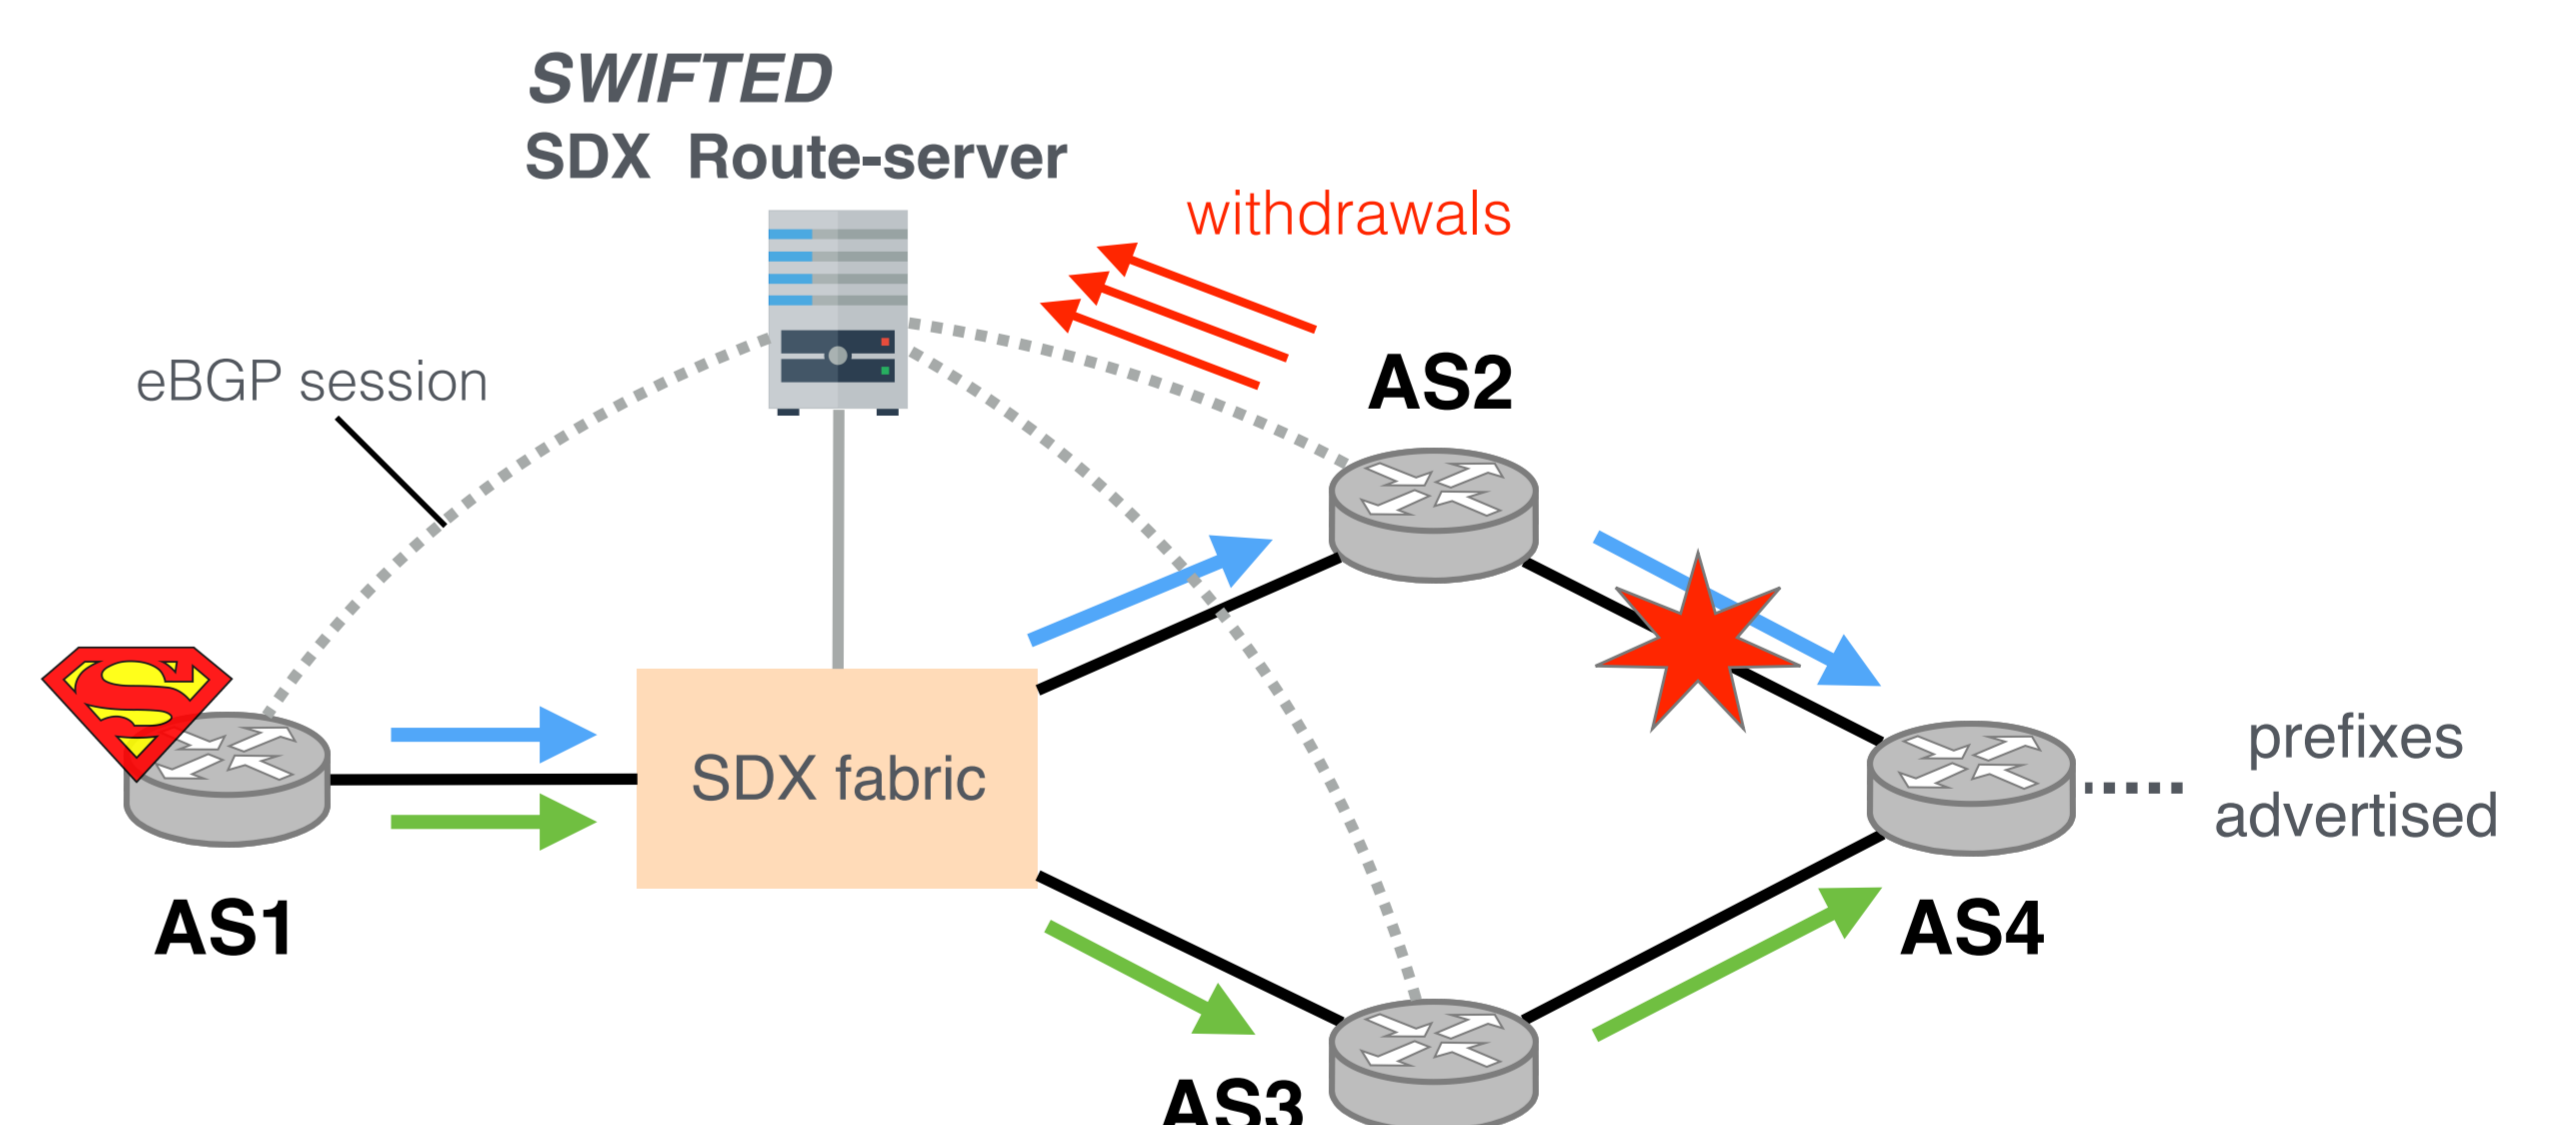

- ◆ Both SWIFT and the SDX rely on a 2-stage forwarding table
  - The first stage groups packets using a tag
  - The second stage forwards packets according to their tags' values
- ◆ Yet, SWIFT and the SDX group packets differently
  - SWIFT groups packets based on the resources they are sharing (AS path)
  - SDX groups packets based on their forwarding equivalence class
- ◆ We show that SWIFT and the SDX can share the same data-plane tag



We reserve 16 bits for the SDX, 24 bits for SWIFT, 8 bits are used for the primary next-hop (used by both)

## 4. Demonstration

- ◆ We modified the current implementation of iSDX to support SWIFT
  - Source code available at <https://github.com/nsg-ethz/iSDX>
- ◆ We built the network depicted on the right with Mininet [4]
  - We configured AS4 to advertise 5k, 10k, 50k, 100k, 250 and 500k prefixes
  - and made sure AS1 prefers AS2 to reach AS4
- ◆ We simulate a **remote failure** on the link between AS2 and AS4



- ◆ Without SWIFT, BGP takes 90 seconds to converge with 500k prefixes
- ◆ With SWIFT, BGP always converges within 1.4 seconds
  - Irrespective of the number of prefixes affected by the failure

## 5. Conclusion

1. SWIFT is a fast reroute framework that works upon remote outages
2. SDX platforms converge slowly upon remote outages
3. SWIFT can be deployed at SDX with a simple software update
4. We deployed SWIFT in an SDX and showed that the convergence time is always within 1.4s

## 6. References

- [1] Craig Labovitz et al. "Delayed Internet routing convergence" ACM SIGCOMM CCR, 2000
- [2] Thomas Holterbach et al. "SWIFT: Predictive Fast Reroute" ACM SIGCOMM, 2017
- [3] Arpit Gupta et al. "An industrial-scale software defined internet exchange point" USENIX NSDI 2016
- [4] Bob Lantz et al. "A Network in a Laptop: Rapid Prototyping for Software-defined Networks. SIGCOMM Hotnets, 2010