ETH zürich

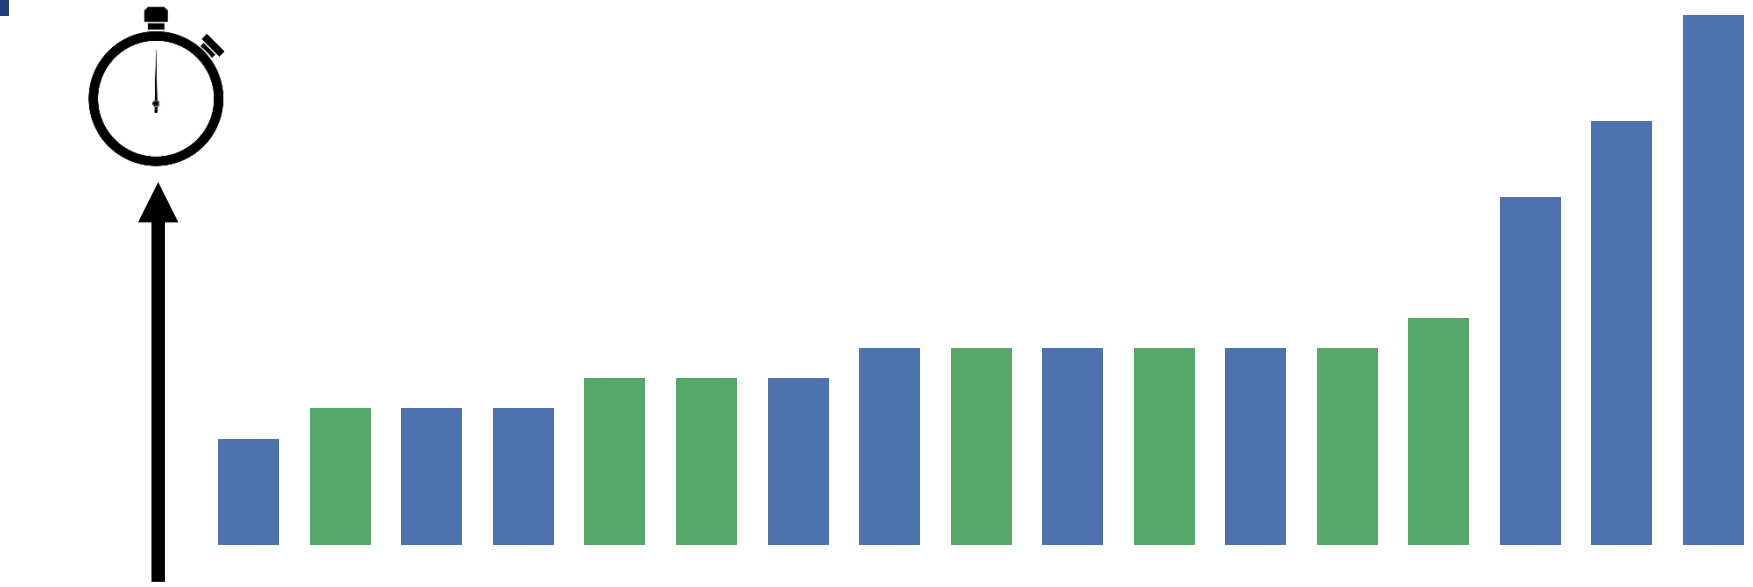

ATP TCP

Reducing the Latency-Tail of Short-Lived Flows: Adding Forward Error Correction in Data Centers

Klaus-Tycho Foerster Demian Jaeger David Stolz Roger Wattenhofer
ETH Zurich

Time is Money

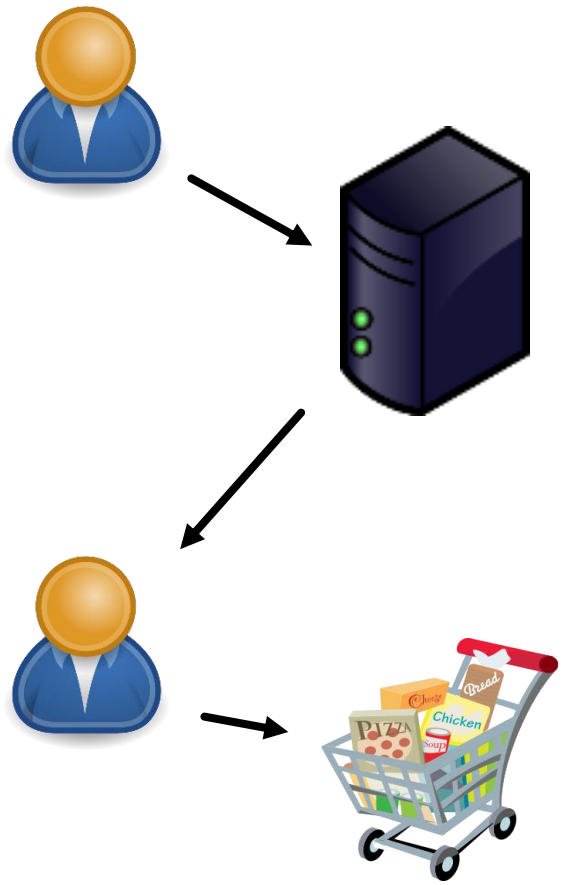

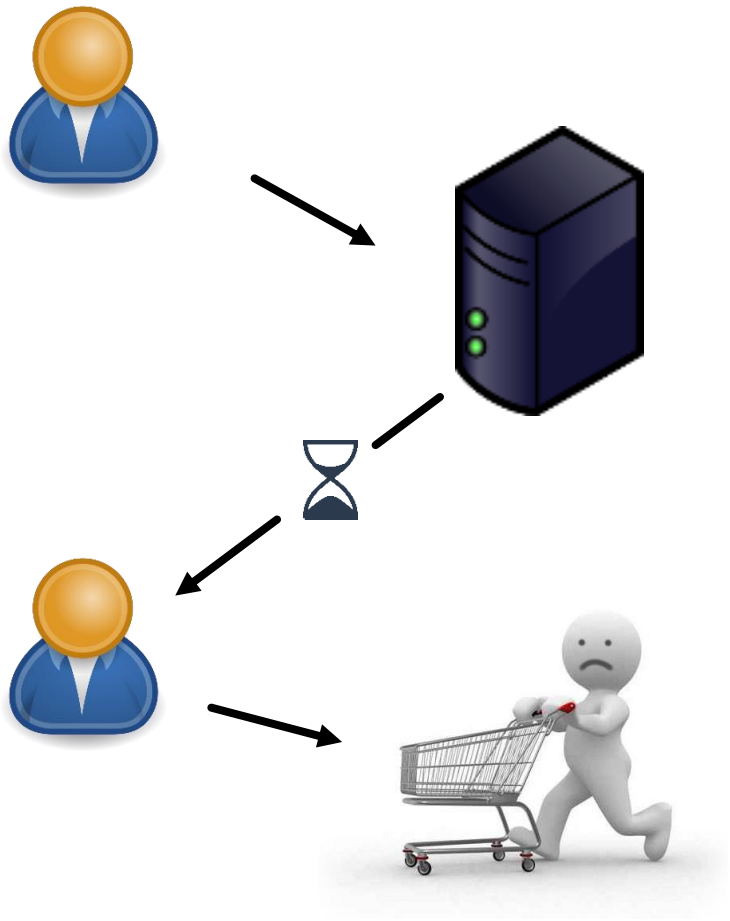

Datacenter Traffic

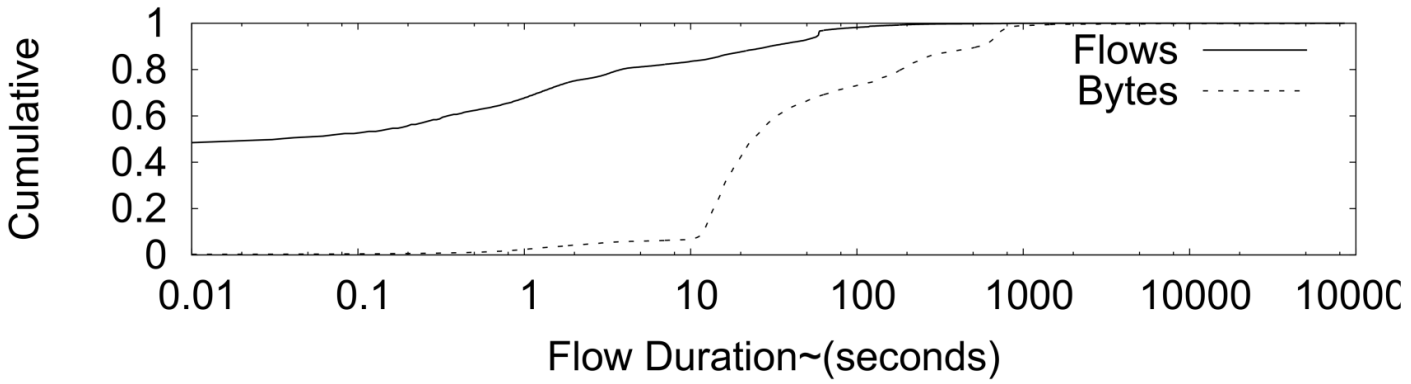

S. Kandula et al., The Nature of Datacenter Traffic. IMC 2009

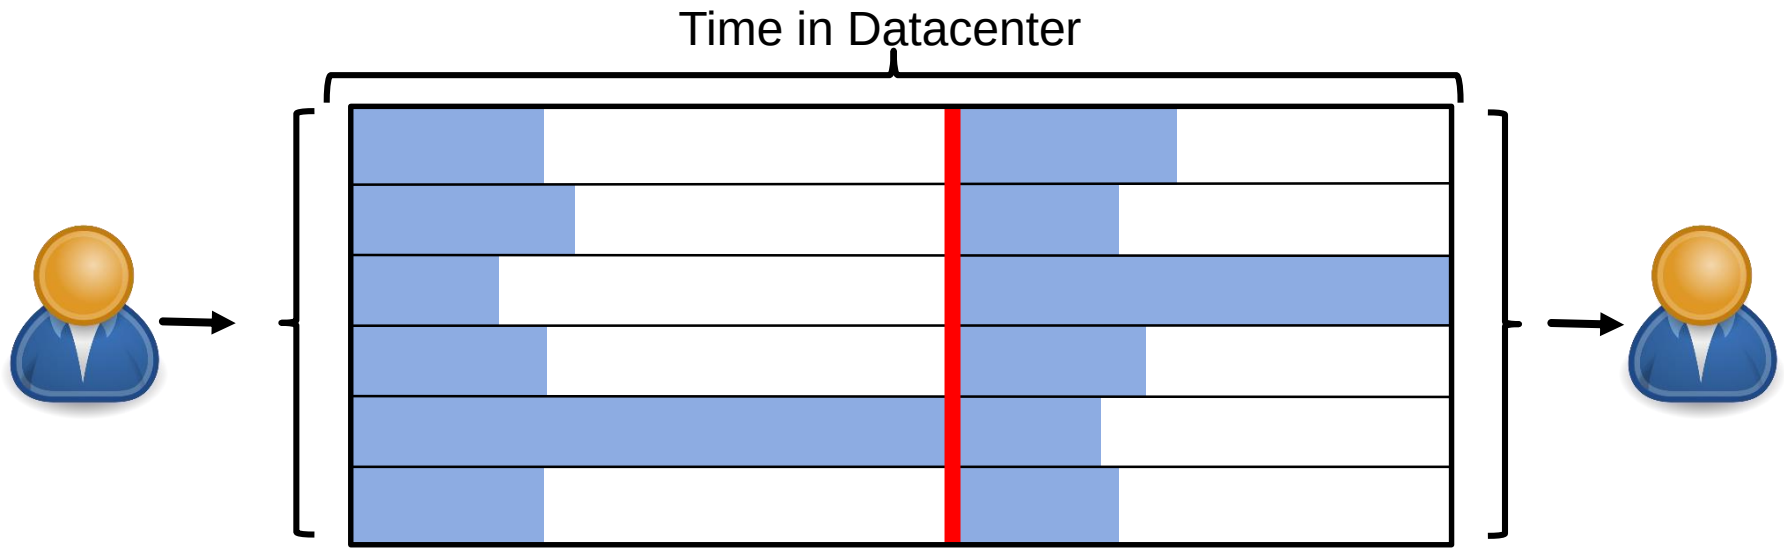

Overview

Problem

- TCP is sensitive to retransmissions
  - Induces latency-tail in congested networks

    → **Goal**

Prior Work

- Add forward error correction (FEC) at link layer
  - Wireless networks
- Add general overhead
- Reserve capacity

Our Approach

- Adaptive FEC on packet level

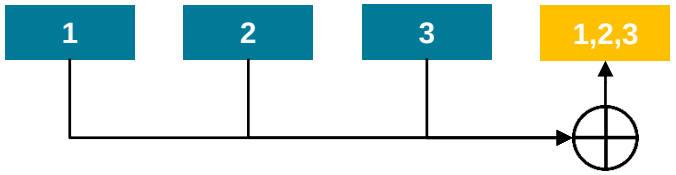

ATP: A Protocol with Error Correction

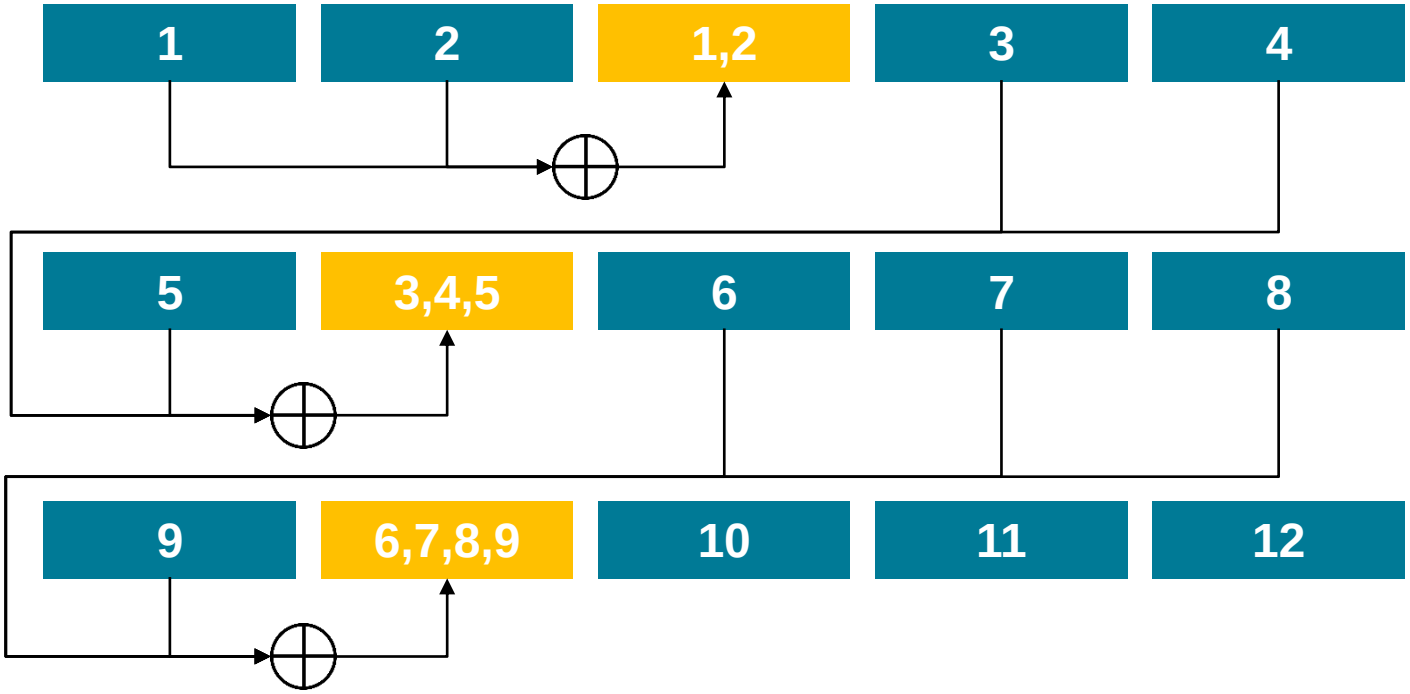

Results

Evaluation

Testbed

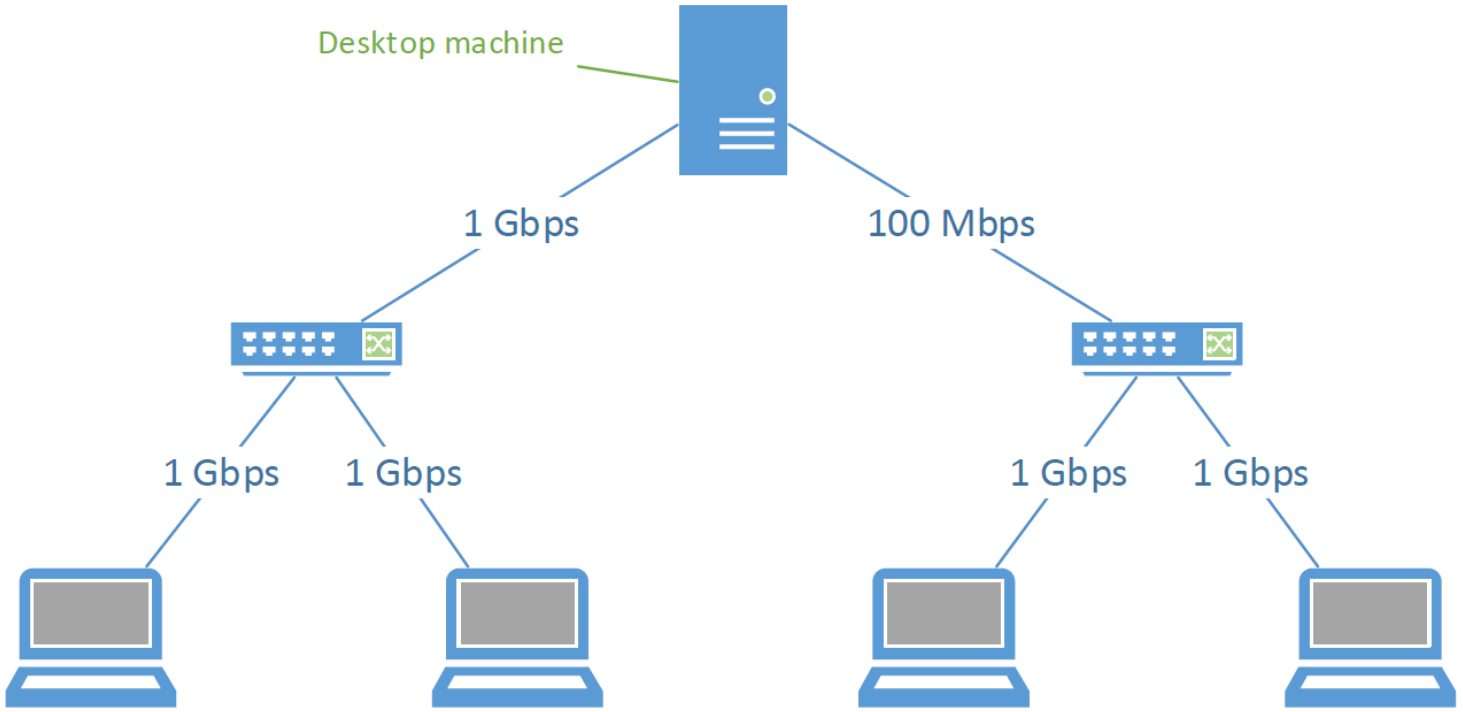

TCP - Background Traffic

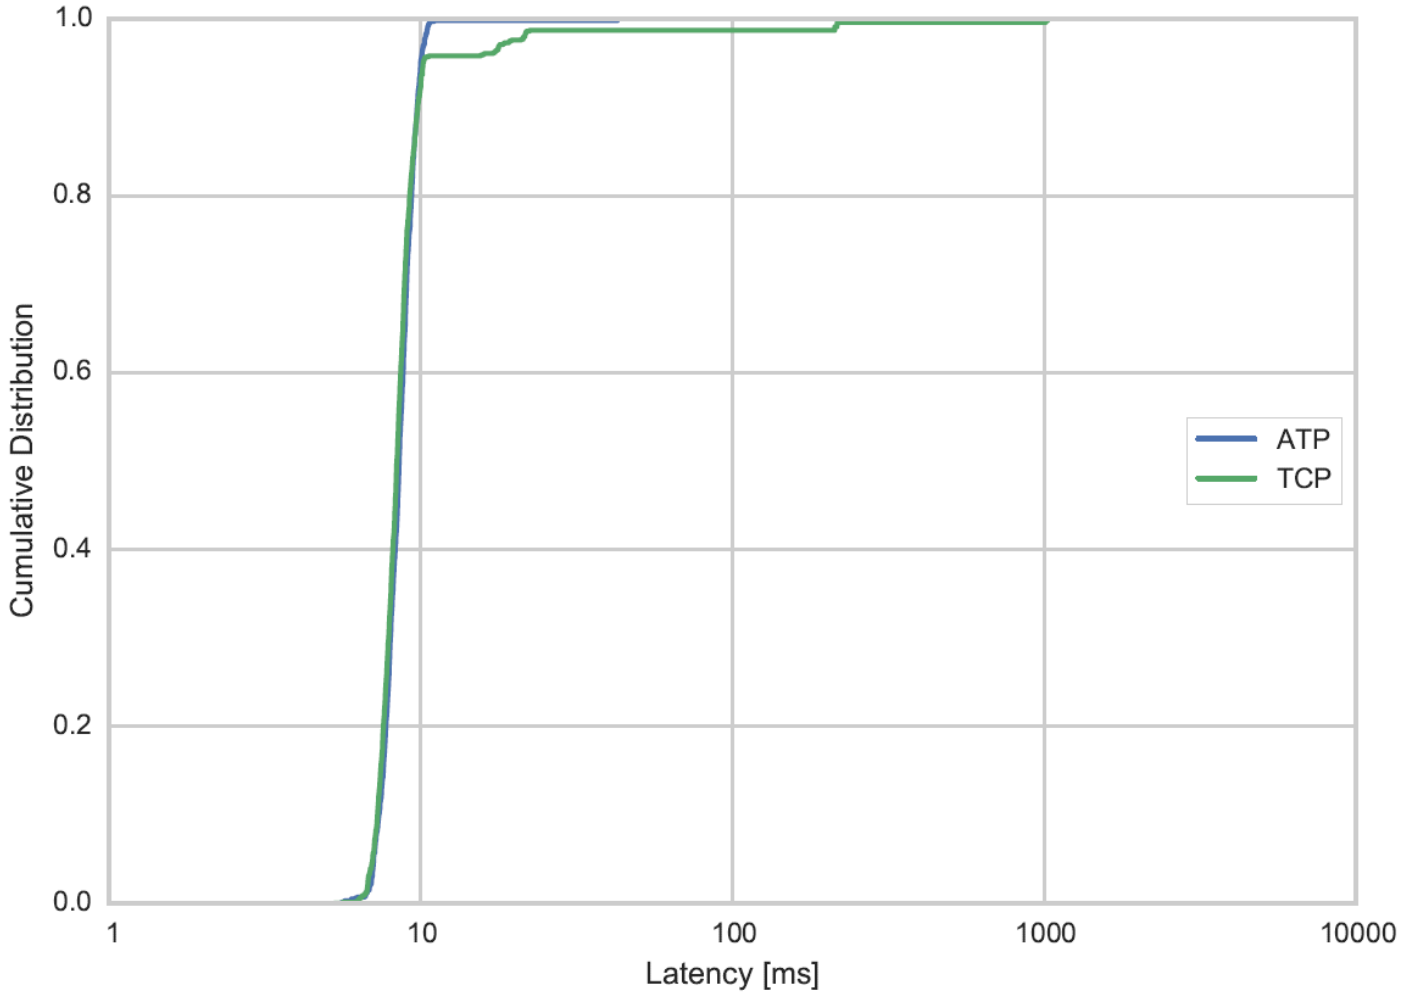

Tail Latency

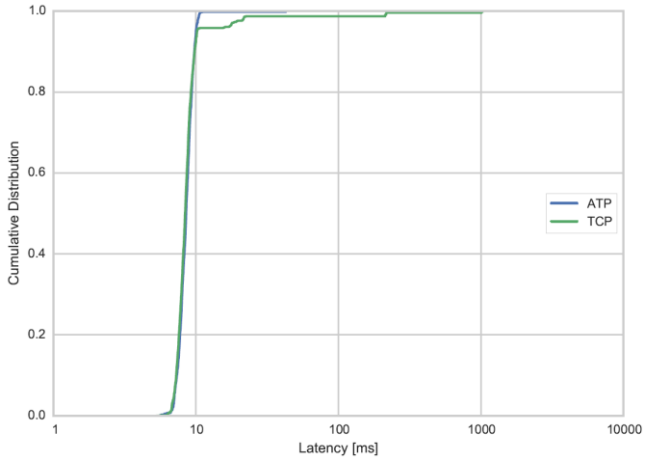

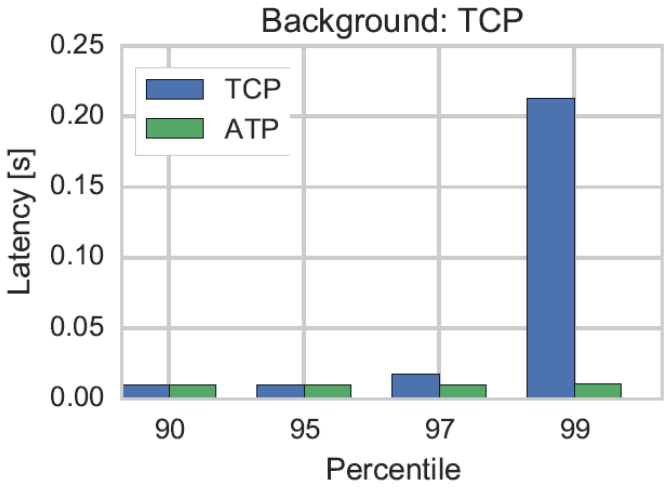

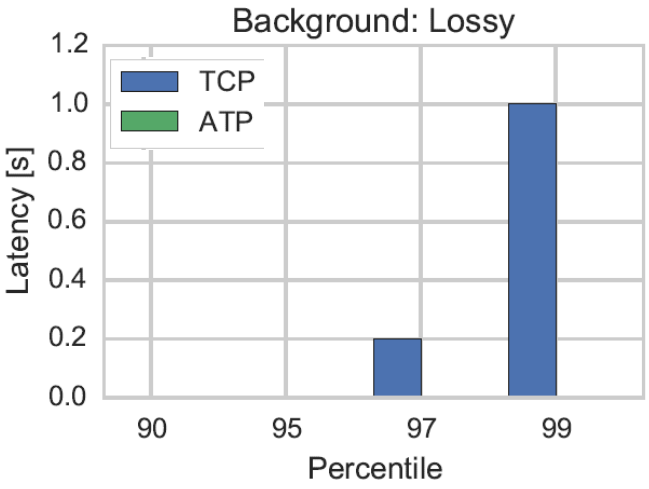

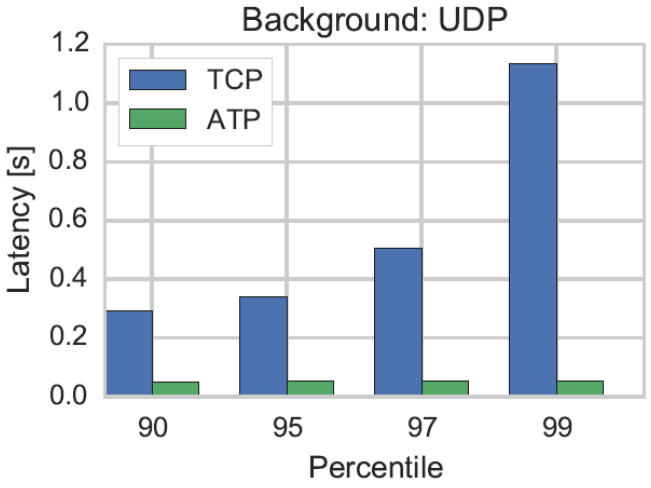

Conclusion

ATP

- Transport Layer Protocol for Datacenters
- Improve Latency of Small Flows by using variable FEC

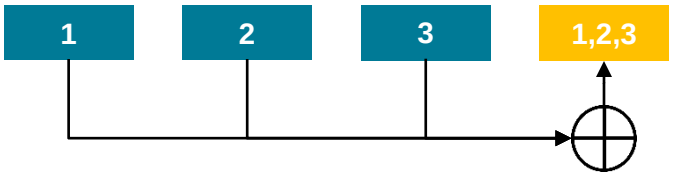

Evaluation

- Fairness to TCP and other ATP Streams
- Similar to TCP in not Congested Network
- **20 Times smaller** Tail Latency compared to TCP in busy Network