Enabling Router Assisted Congestion Control on the Internet

Marcel Flores Alexander Wenzel
Aleksandar Kuzmanovic

Northwestern

Achieving Fair Congestion Control

- TCP achieves suboptimal performance in dynamic networks
- XCP and RCP
  - Share information in headers - not compatible with TCP
- Our proposal: Work within TCP to share explicit rate information

Kick-Ass

- Method for deploying congestion control within the TCP/IP stack
- Packet lengths implicitly communicate rates
- Kick-Ass…
  1. Has significantly improves performance compared to TCP
  2. Is effective on legacy/mixed paths
  3. Is bounded by TCP performance

Kick-Ass

- Routers use any algorithm to calculate rates
- Advertise rates in logarithmically-encoded packet lengths
- Leading fragment at receiver encodes bottleneck rate

Kick-Ass

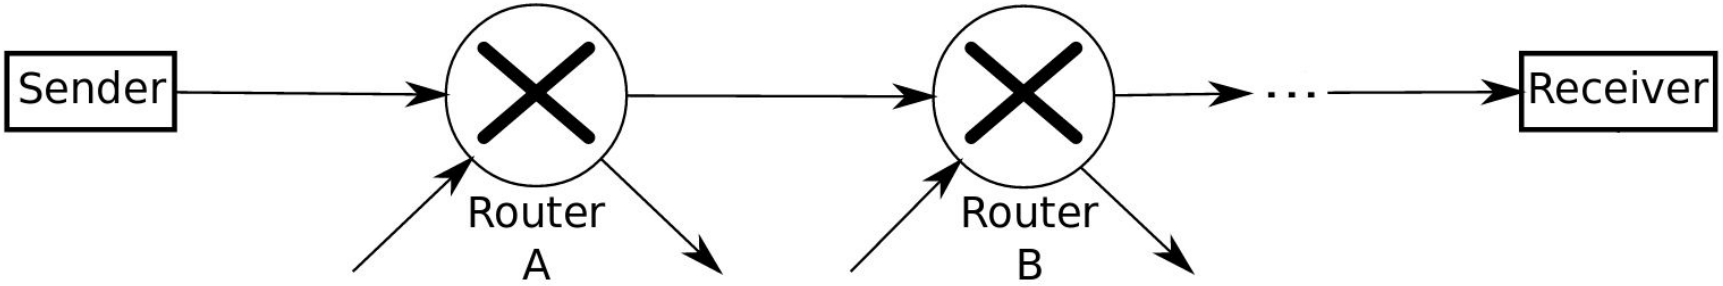

Kick-Ass

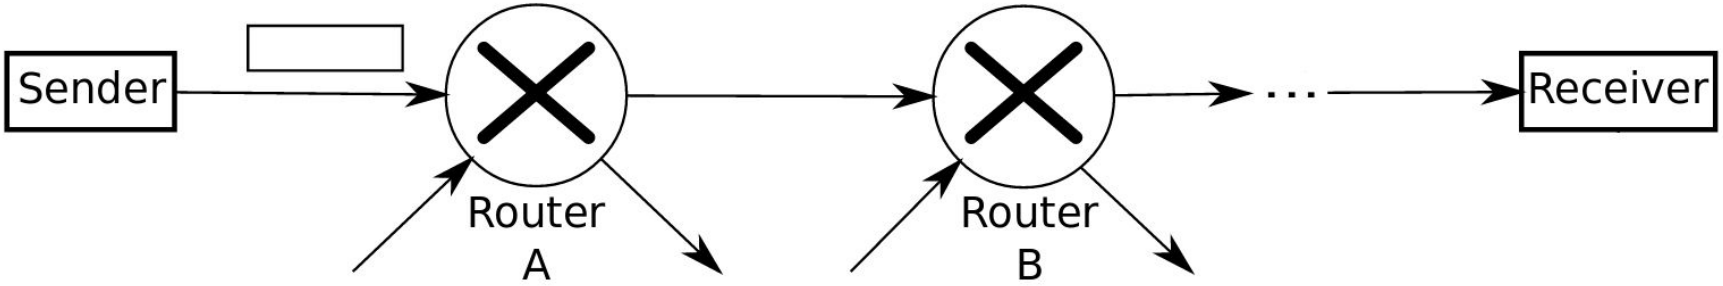

Kick-Ass

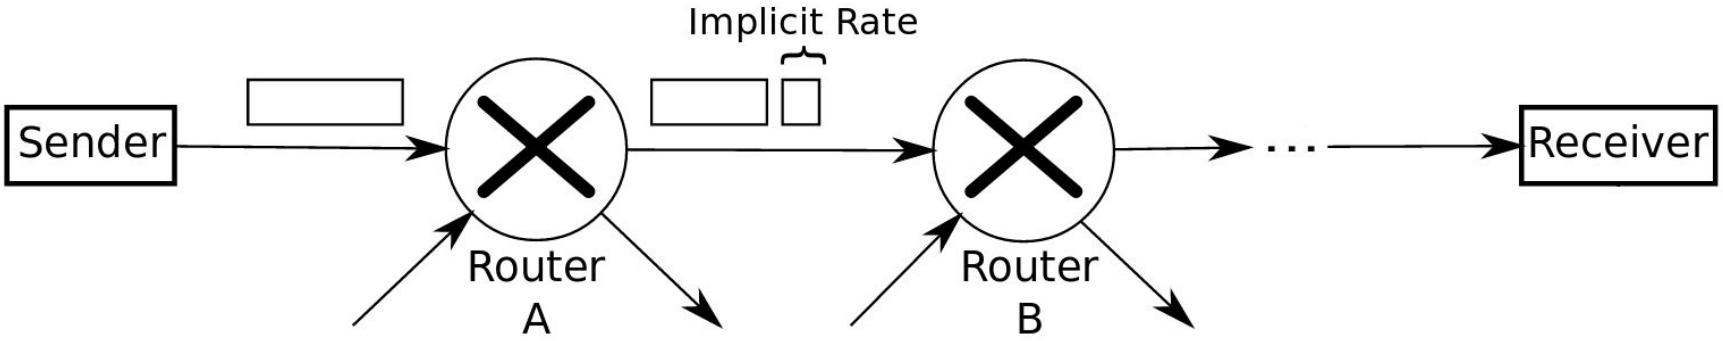

Kick-Ass

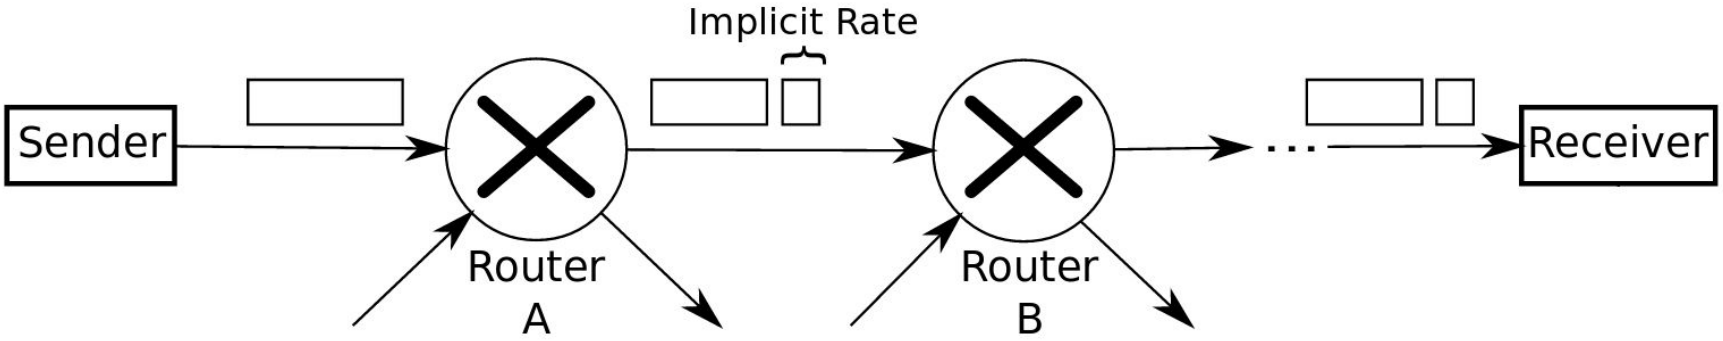

Kick-Ass

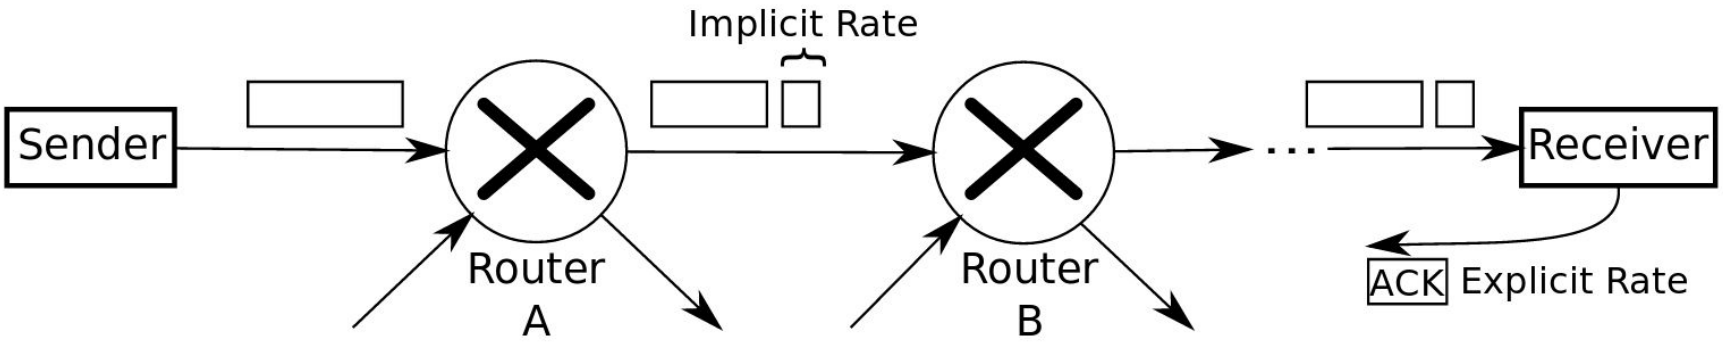
Implicit Rate
Sender
Router
A
Router
B
Receiver
ACK
Explicit Rate

Simulations

- Kick-Ass, TCP, and RCP evaluated in NS-3 simulations
- Evaluate light/heavy congestion in full and mixed deployment
- Parameters
  - Link rate: 150 Mbps
  - Queue size: 1000 packets
  - RTT: 100 ms

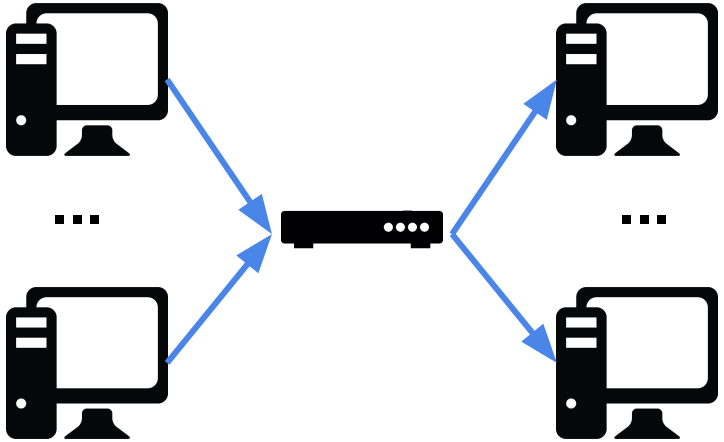

Full Deployment Kick-Ass vs RCP

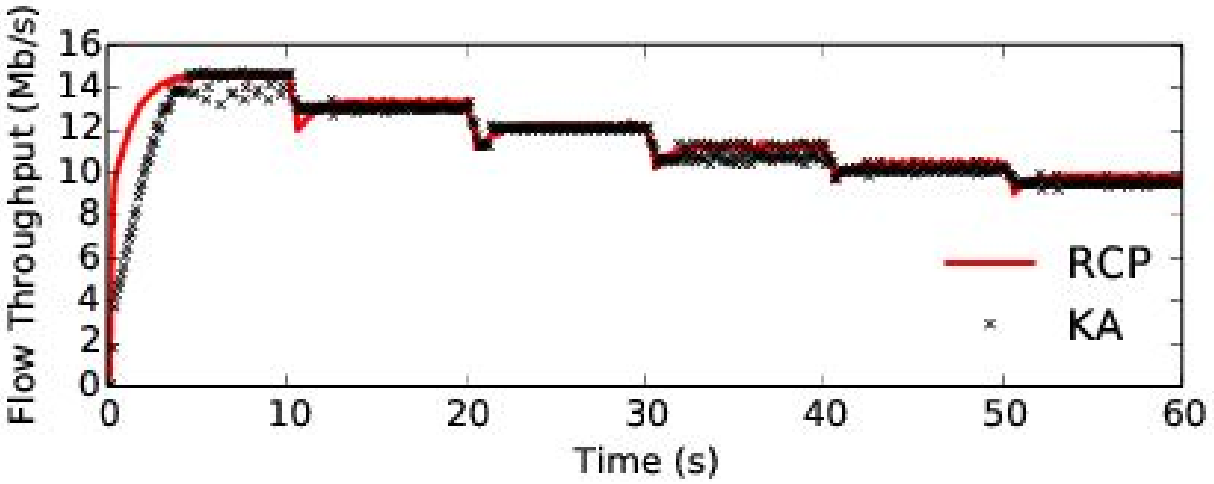

Kick-Ass achieves comparable fair-sharing to RCP

Full Deployment - Light Congestion

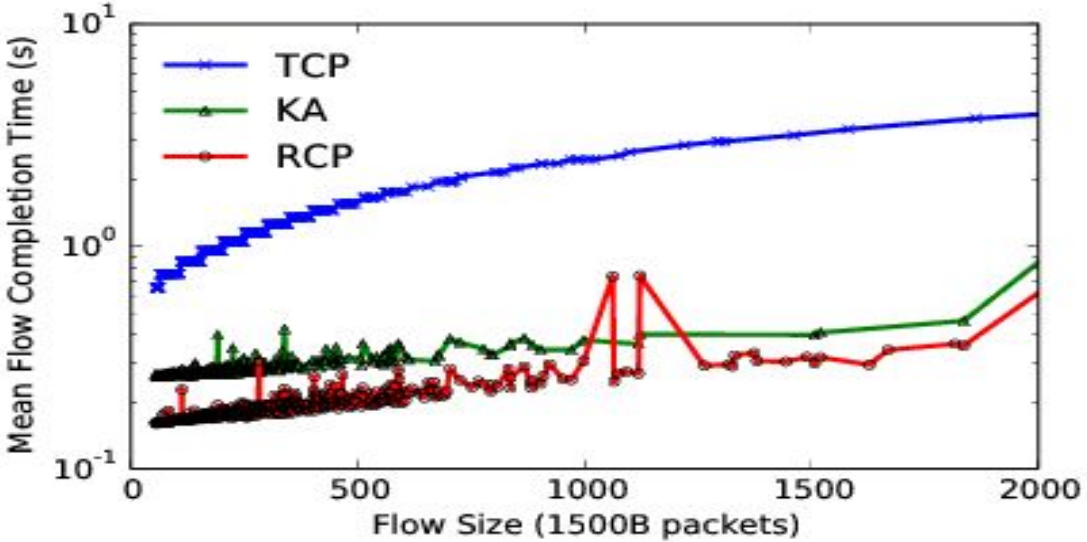

Kick-Ass achieves performance benefits over TCP

Light Congestion - Mixed Traffic

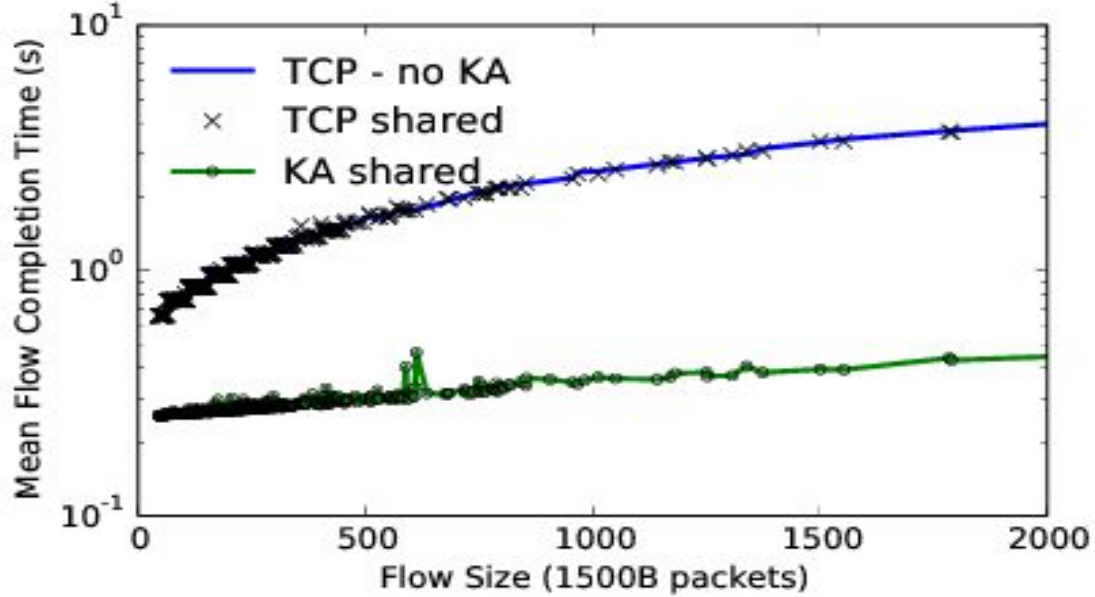

Kick-Ass shows improvements in mixed environments, does not harm TCP

Kick-Ass in Heavy Congestion

- Problem: Heavy TCP will starve fair-share protocols
- Solution: TCP Mode - Kick-Ass detects heavy congestion and switches to TCP to compete
- Kick-Ass worst case performance is TCP

Kick-Ass in Heavy Congestion

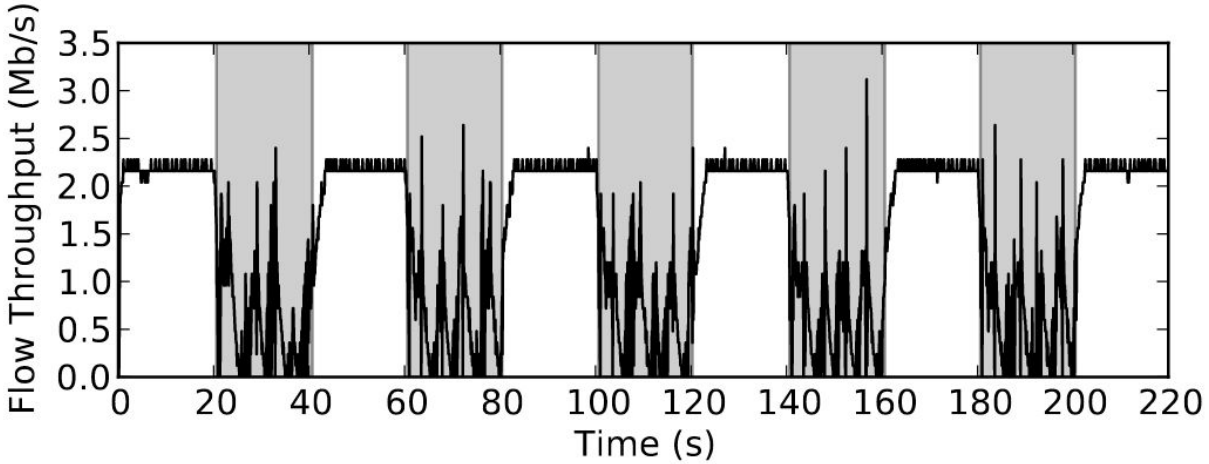

Kick-Ass quickly detects heavy congestion and quickly returns to optimal performance afterwards

Testbed Experiments

- Protocols Evaluated:
  - Kick-Ass
  - TCP (CUBIC)
  - PCC
  - FQ-CoDel

- Parameters
  - Kick-Ass implemented as Linux queueing discipline
  - Router+Endpoints: 3.3 Ghz Intel i5 running Arch Linux 3.12.3
  - Netem creates 100 ms+ RTT

Light Congestion in Testbed

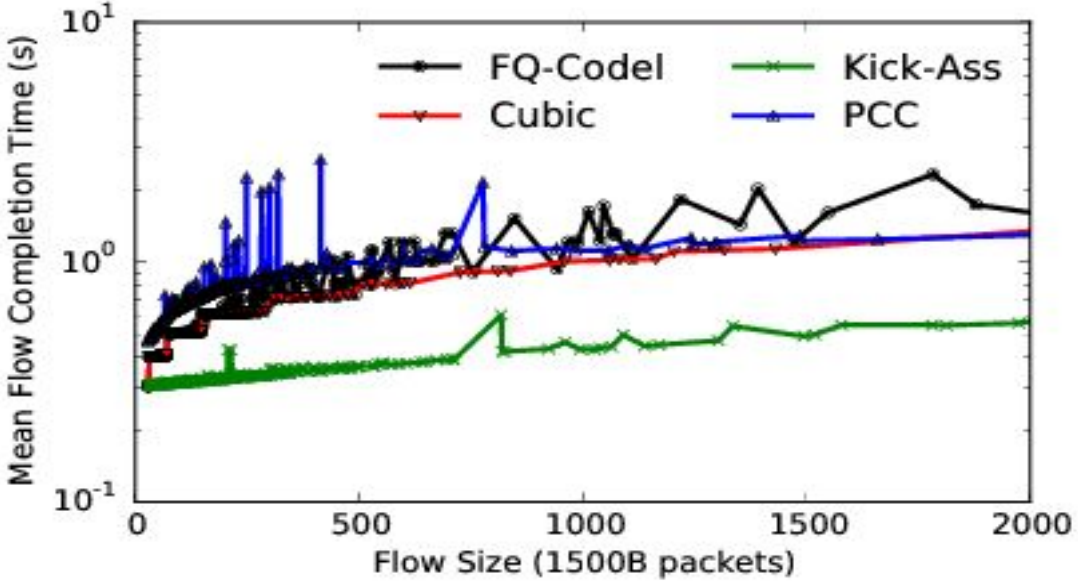

Kick-Ass outperforms all protocols

Heavy Congestion in Testbed

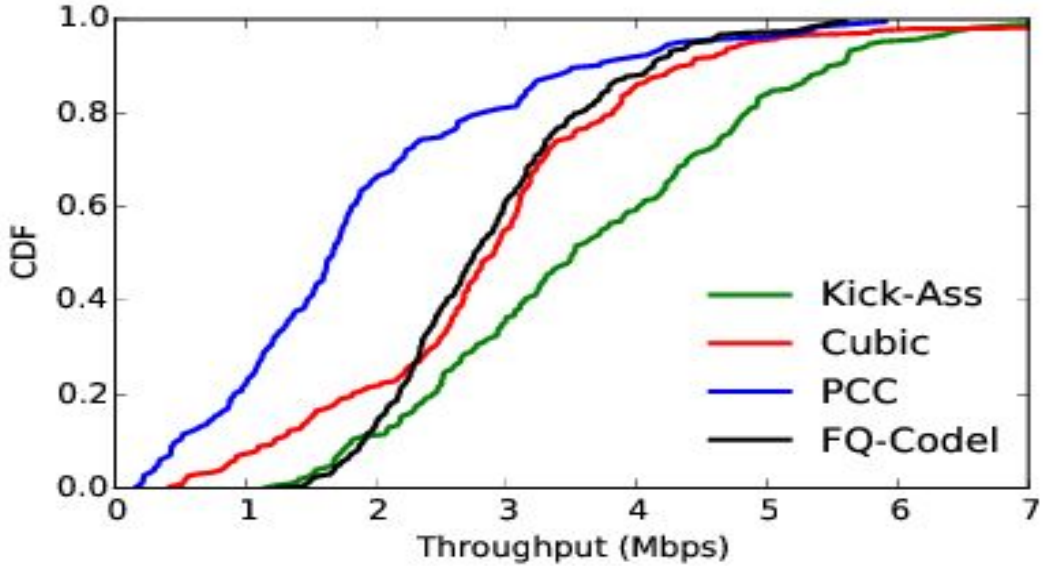

Kick-Ass achieves higher throughput in > 80% of cases

Kick-Ass in the Internet

- Flows routed through local Kick-Ass router to destinations in the Internet
- Kick-Ass outperforms TCP 1.2x-4x depending on path

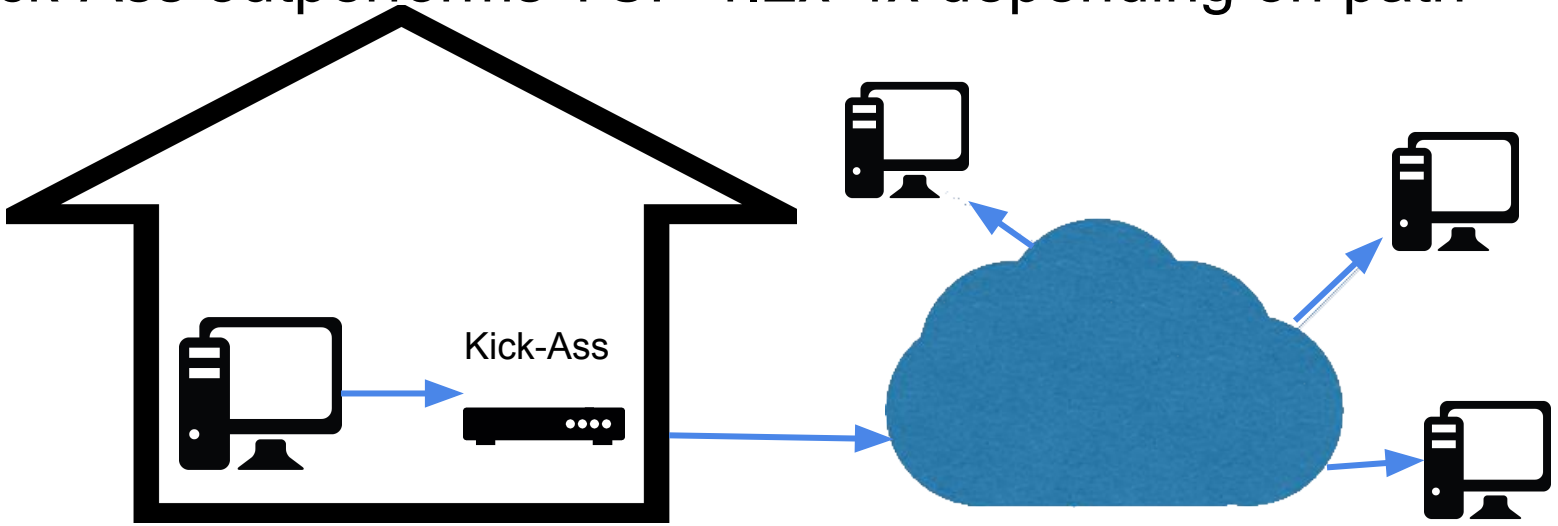

Kick-Ass

- Congestion control is possible within the TCP/IP stack. Packet fragmentation can be used to communicate explicit rate.

Kick-Ass

- Congestion control is possible within the TCP/IP stack. Packet fragmentation can be used to communicate explicit rate.

- Kick-Ass outperforms other protocols in simulations, testbed experiments, and on the Internet

Kick-Ass

- Congestion control is possible within the TCP/IP stack. Packet fragmentation can be used to communicate explicit rate.
- Kick-Ass outperforms other protocols in simulations, testbed experiments, and on the Internet
- Kick-Ass is compatible with legacy flows. Its performance is lower-bounded by TCP