Measuring Latency Variation in the Internet

Toke Høiland-Jørgensen[⋆], Bengt Ahlgren[†], Per Hurtig[⋆] and Anna Brunstrom[⋆]

[⋆]Karlstad University, [†]SICS

{firstname.lastname}@kau.se, bengta@sics.se

Research question

How much bufferbloat exists in the internet?

The idea

- Combine large-scale active measurements with passive captures

The idea

- Combine large-scale active measurements with passive captures
- Use *latency span* as metric

The idea

- Combine large-scale active measurements with passive captures
- Use *latency span* as metric
- Estimate queueing latency by:

The idea

- Combine large-scale active measurements with passive captures
- Use *latency span* as metric
- Estimate queueing latency by:
  - Looking at latency drop after TCP congestion event

The idea

- Combine large-scale active measurements with passive captures
- Use *latency span* as metric
- Estimate queueing latency by:
  - Looking at latency drop after TCP congestion event
  - Correlating latency with link load

The datasets

We combine two datasets:

1. M-lab NDT dataset

The datasets

We combine two datasets:

1. M-lab NDT dataset
   - User-initiated active measurements (10s download)

The datasets

We combine two datasets:

1. M-lab NDT dataset
   - User-initiated active measurements (10s download)
   - Total 265.8 M test runs, spanning 2010–2015 (incl)

The datasets

We combine two datasets:

1. M-lab NDT dataset
   - User-initiated active measurements (10s download)
   - Total 265.8 M test runs, spanning 2010–2015 (incl)
   - Data source: TCP state machine RTT samples (span per flow)

The datasets

We combine two datasets:

1. M-lab NDT dataset
   - User-initiated active measurements (10s download)
   - Total 265.8 M test runs, spanning 2010–2015 (incl)
   - Data source: TCP state machine RTT samples (span per flow)
2. Passive capture from ISP access network

The datasets

We combine two datasets:

1. M-lab NDT dataset
   - User-initiated active measurements (10s download)
   - Total 265.8 M test runs, spanning 2010–2015 (incl)
   - Data source: TCP state machine RTT samples (span per flow)
2. Passive capture from ISP access network
   - 1 Gbps aggregation links serving 50 and 400 customers (respectively)

The datasets

We combine two datasets:

1. M-lab NDT dataset
   - User-initiated active measurements (10s download)
   - Total 265.8 M test runs, spanning 2010–2015 (incl)
   - Data source: TCP state machine RTT samples (span per flow)
2. Passive capture from ISP access network
   - 1 Gbps aggregation links serving 50 and 400 customers (respectively)
   - Collected over a period of 8 months in 2014

The datasets

We combine two datasets:

1. M-lab NDT dataset
   - User-initiated active measurements (10s download)
   - Total 265.8 M test runs, spanning 2010–2015 (incl)
   - Data source: TCP state machine RTT samples (span per flow)
2. Passive capture from ISP access network
   - 1 Gbps aggregation links serving 50 and 400 customers (respectively)
   - Collected over a period of 8 months in 2014
   - Data source: Delay between SYN+ACK and ACK for outgoing flows (span per user)

Latency span (NDT dataset)

All flows

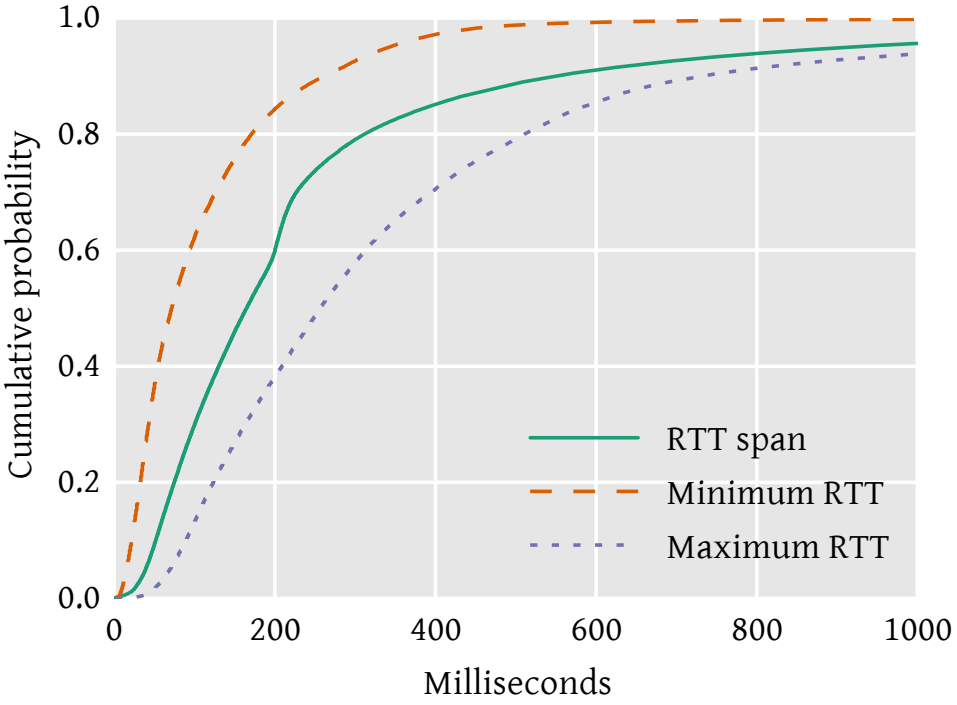

Per continent (2015)

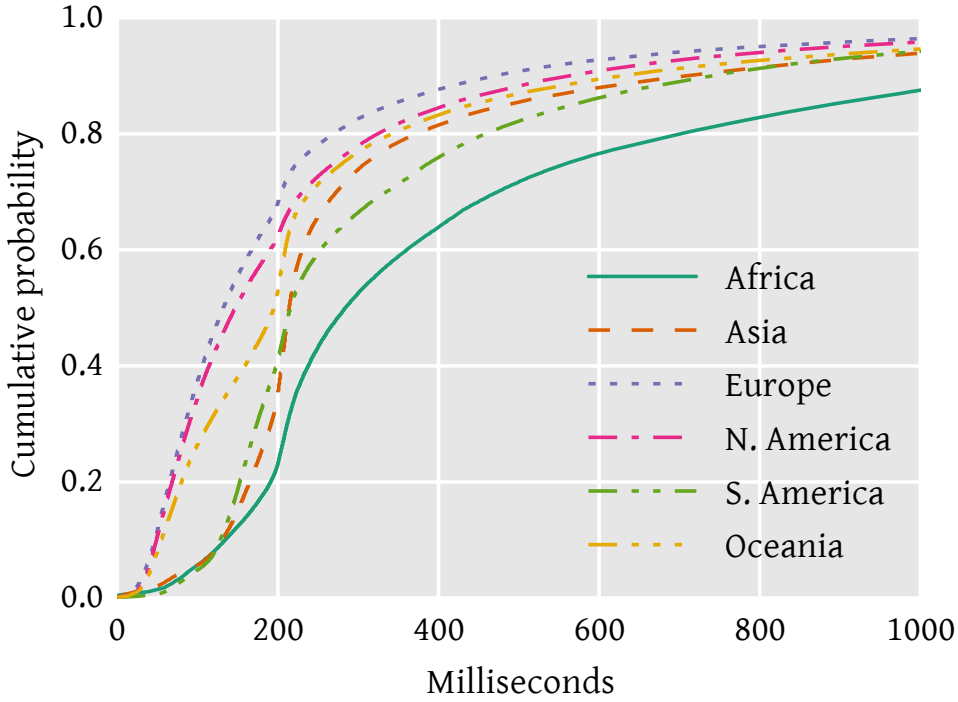

Over time (NDT dataset)

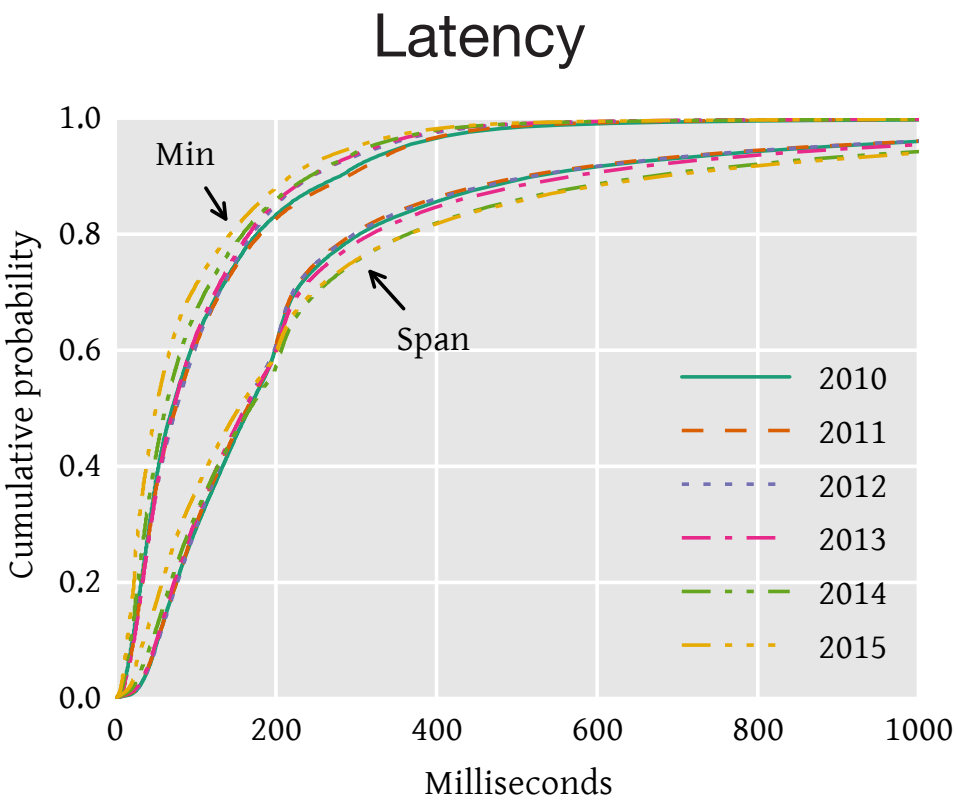

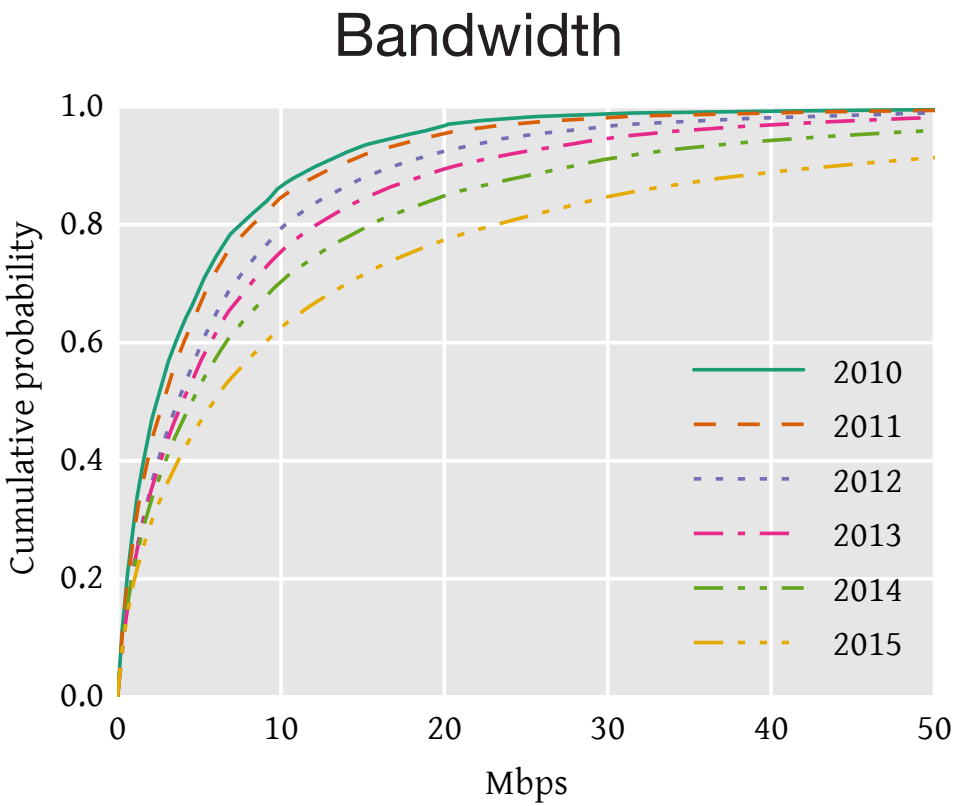

Access network data

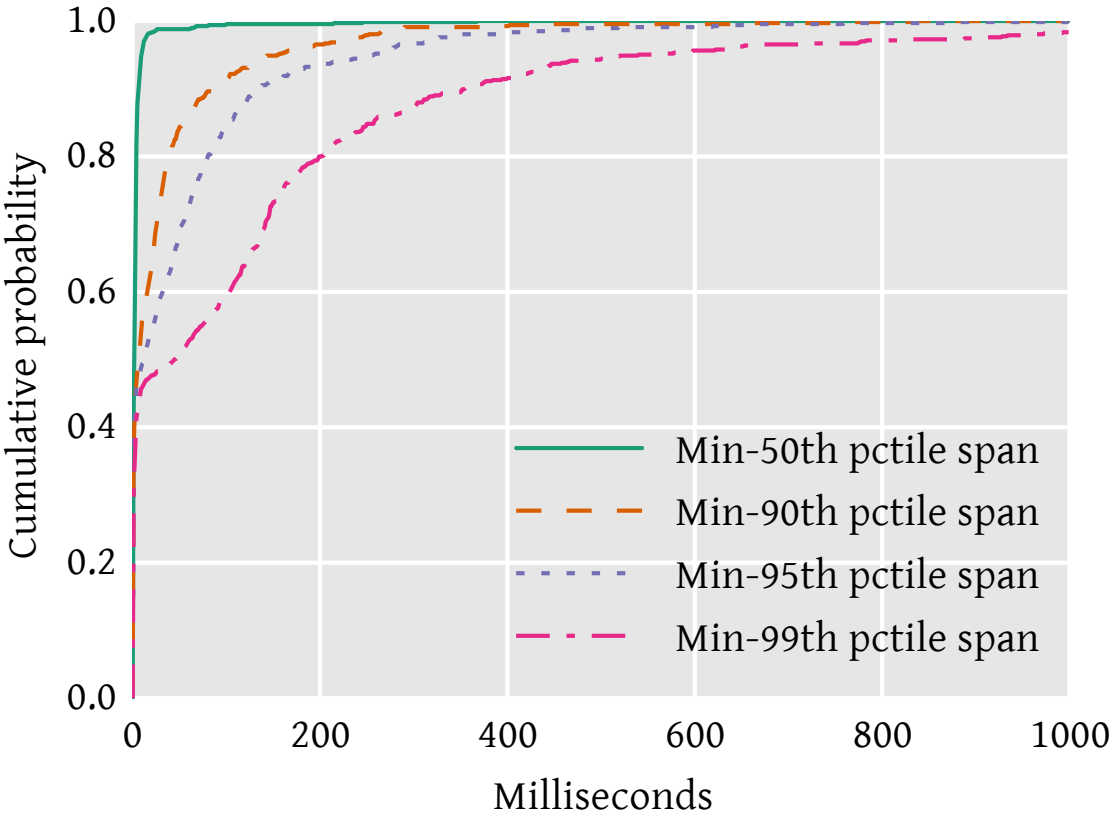

A full day at the first aggregation link.

Queueing latency (NDT data; 5.7 M flows)

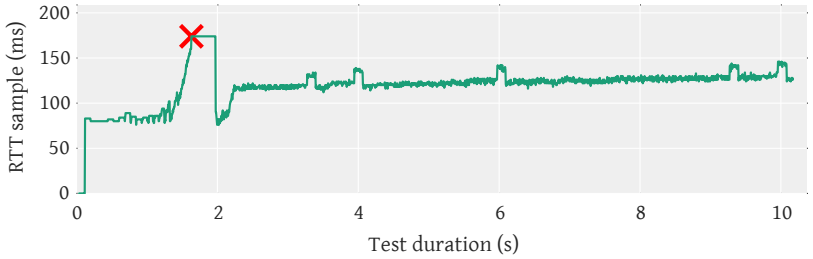

Queueing latency (NDT data; 5.7 M flows)

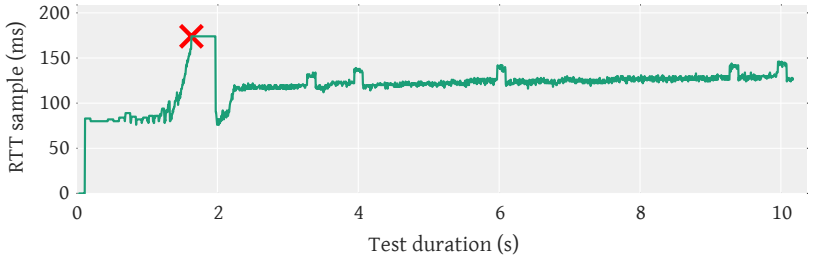

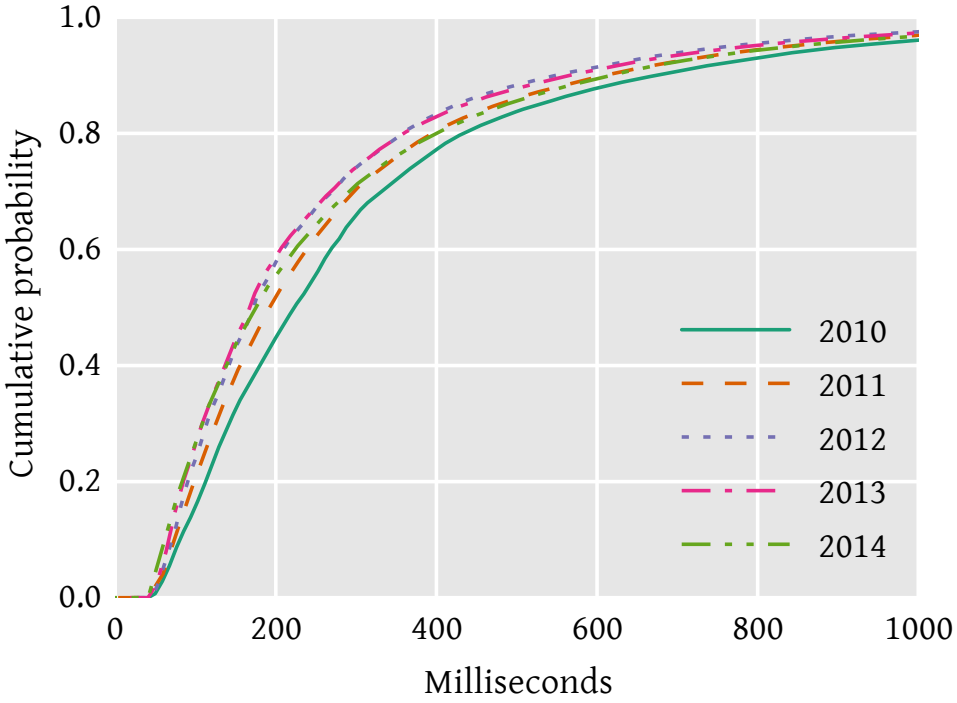

Queueing latency (ISP capture; single flow)

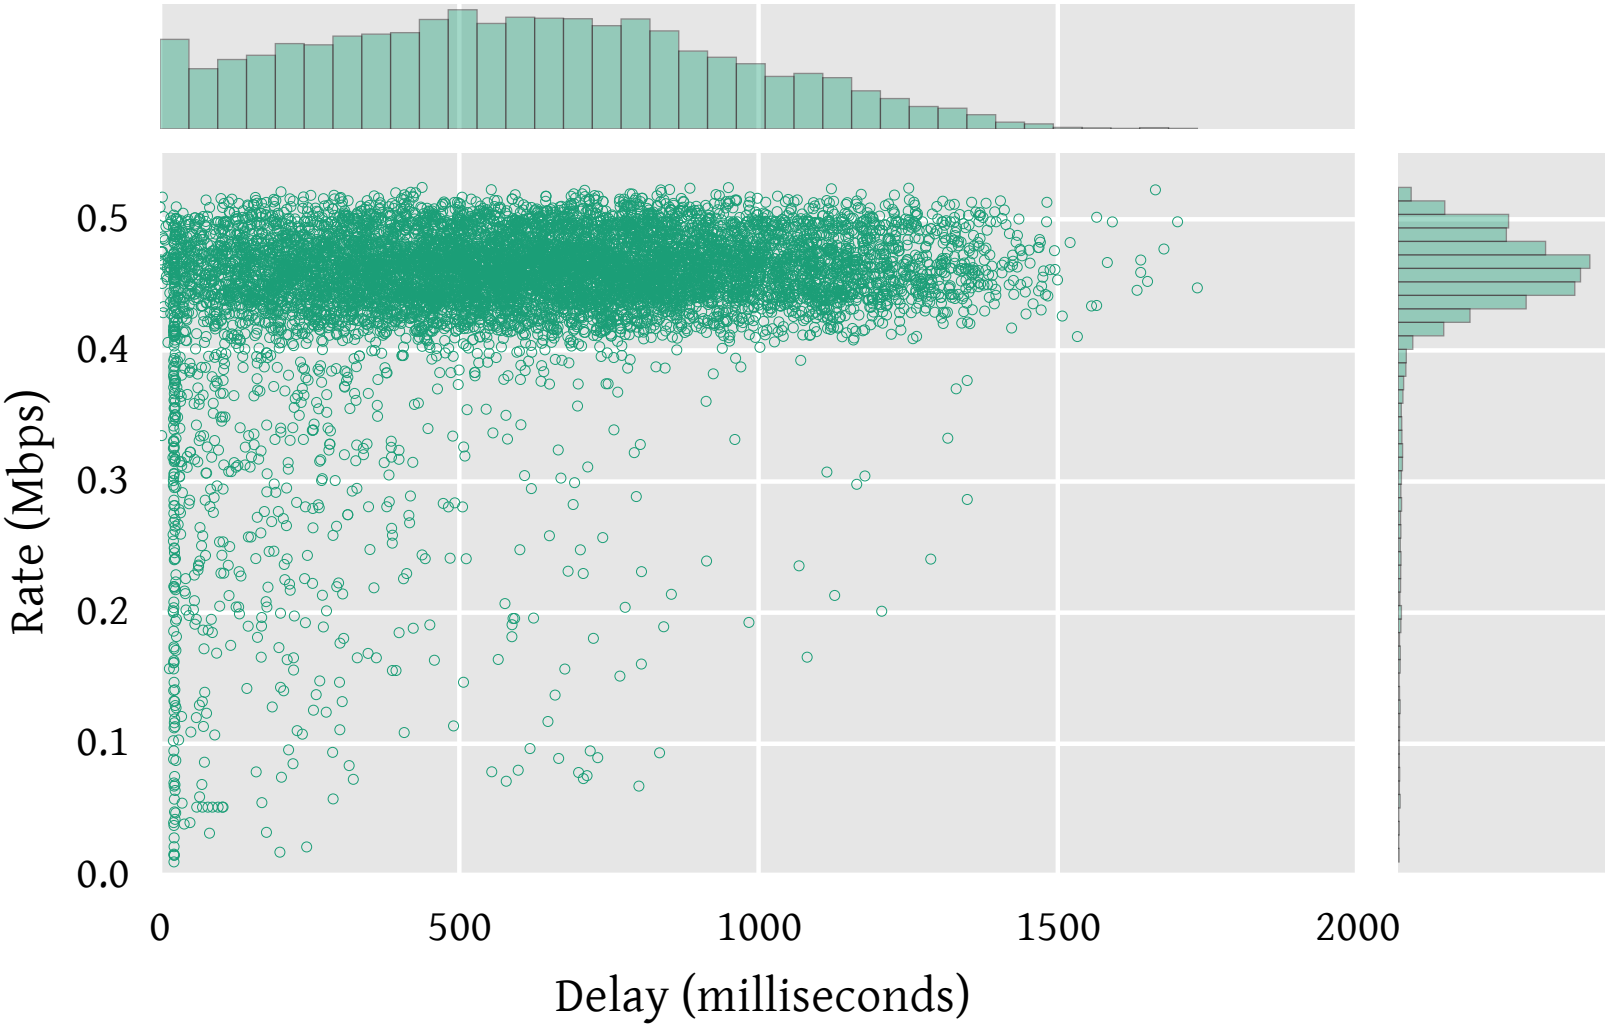

Conclusions

- Latency variation in the internet is significant
  - It has not improved over time
  - There are significant regional differences

Conclusions

- Latency variation in the internet is significant
  - It has not improved over time
  - There are significant regional differences
- At least some of it can be attributed to queueing
  - Where queueing occurs, its magnitude is significant

We need to pay more attention to latency

- Higher bandwidth $\neq$ a better connection

We need to pay more attention to latency

- Higher bandwidth $\neq$ a better connection
- Deploy better queue management today!

We need to pay more attention to latency

- Higher bandwidth $\neq$ a better connection
- Deploy better queue management today!
- Better congestion control? E.g., BBR.