FavorQueue: a Stateless Active Queue Management to Improve TCP Traffic Performance

Emmanuel Lochin ISAE

Latency matters

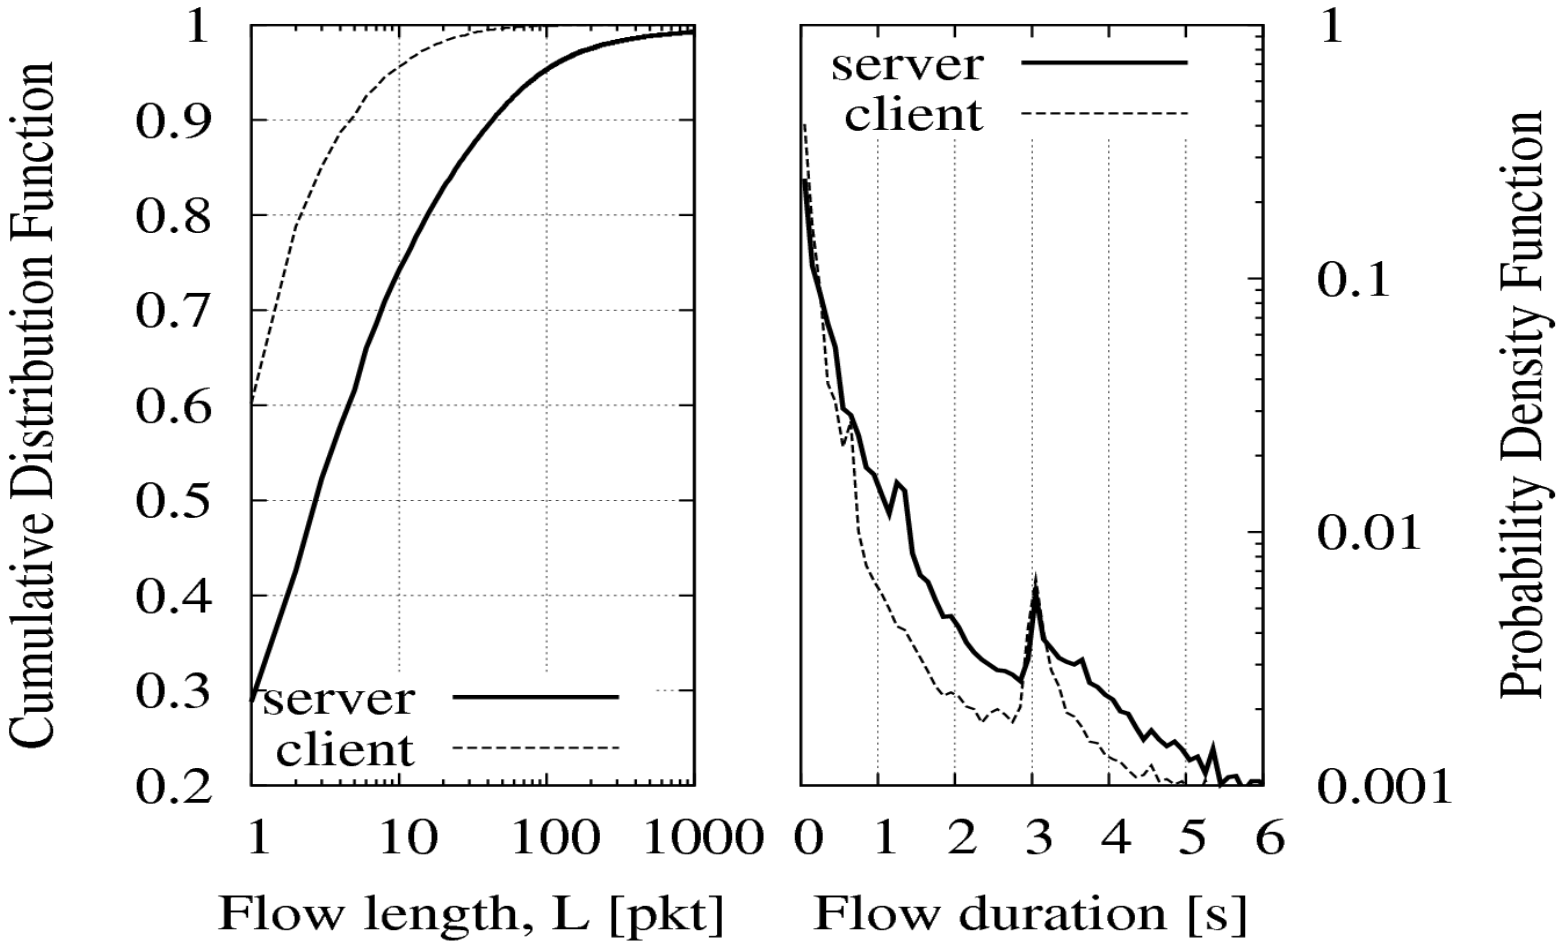

By courstesy of D. Ciullo, M. Mellia, and M. Meo. Two schemes to reduce latency in short lived TCP flows. Communications Letters, 13(10), October 2009.

What is the problem?

- More than 95% (70%) of the client (server) flows are shorter than 10 full-size segments
- As a result, their delay performance is mainly driven by:
  - The end-to-end transfer delay
    - This delay can be reduced if the queueing delay of each router is low
  - The potential losses at the beginning connection
    - The first packets lost at the beginning of a TCP connection (i.e. in the slow-start phase) are mainly recovered by the RTO mechanism
    - At the beginning the initila value is(was) set to 1(3) second(s)

ns-2 experiment

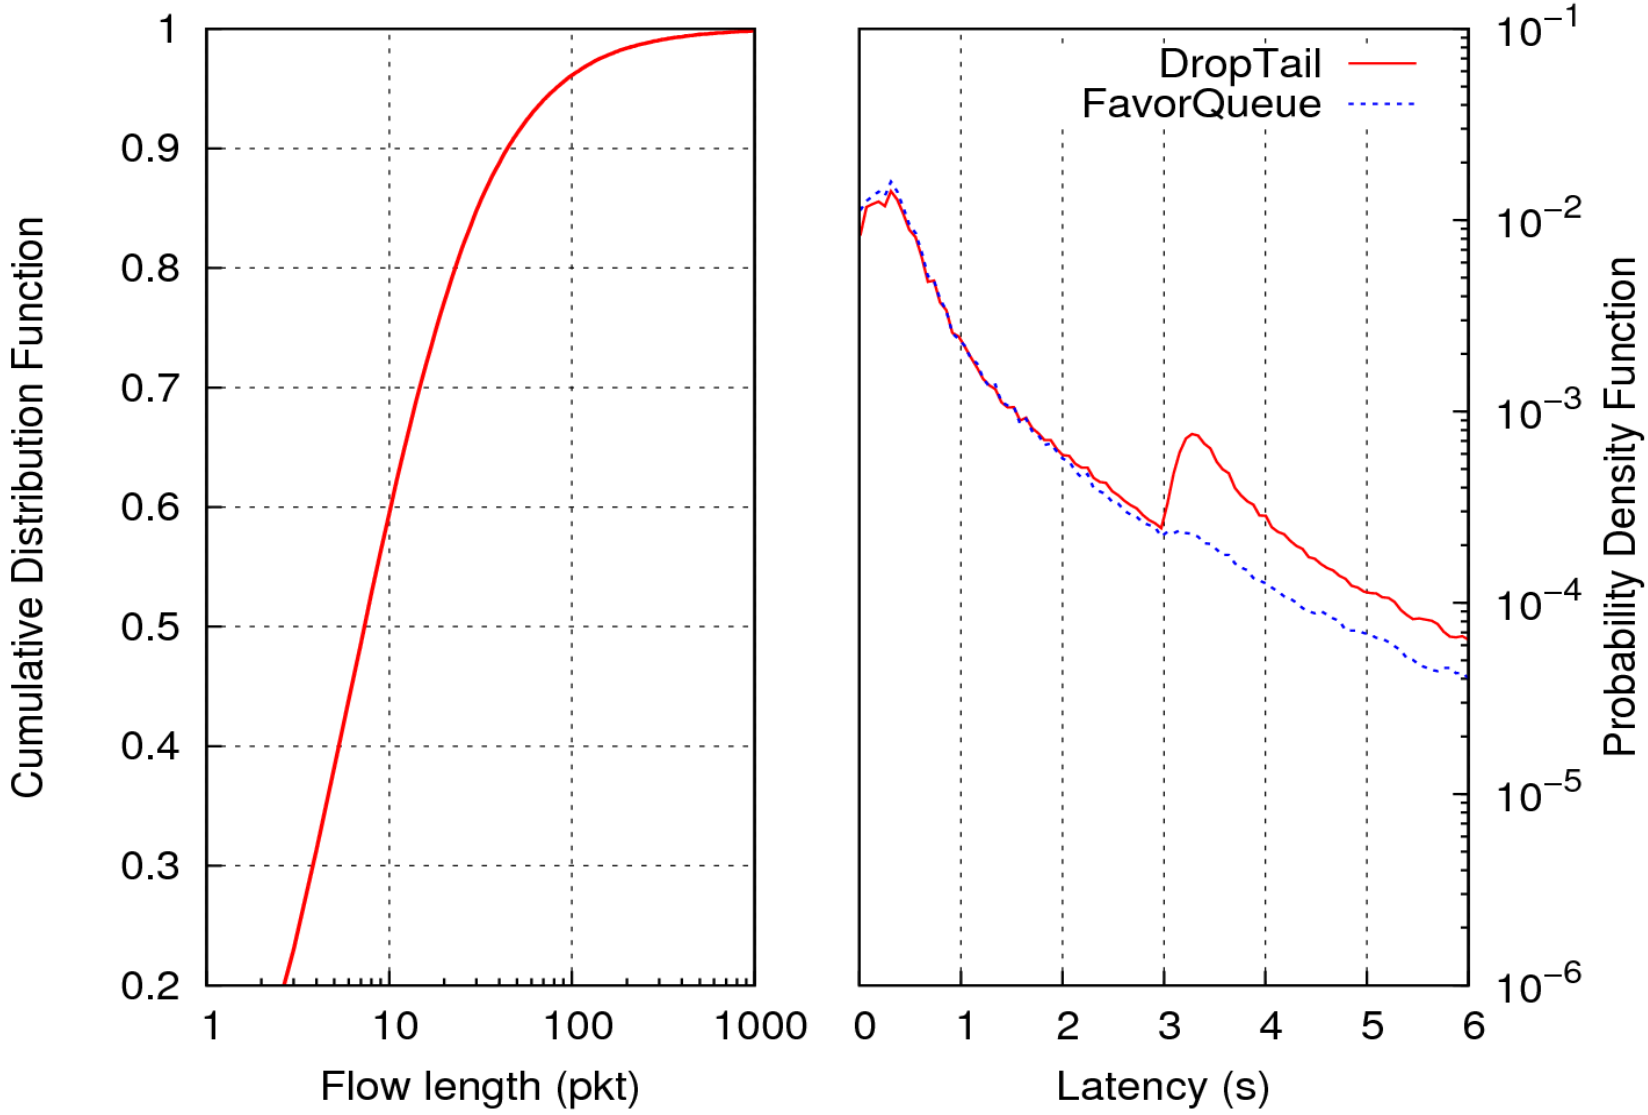

56% of dropped packets are recovered by RTO expires (versus 70% in the previous one)

Our proposal FavorQueue (FaQ)

- Follow a well-known queueing theory result ►the overall mean latency is reduced when shortest job is served first (L. Kleinrock Queuing Systems)
- FavorQueue favors certain packets in order to speed up delay transfer by giving preferential access and to protect them from drop
- Basic principle:
  - When a packet is enqueued, a check is done on the whole queue to seek another packet from the same flow
  - If no other packet is found, it becomes a favored packet and is served first
  - Push-out when queue is full
- FaQ stateless, no parameters to set

Simulations setup

- We follow « Common TCP evaluation suite »
  - The traffic demand, expressed as a bit rate, is the product of the flow arrival rate λ and the average flow size $E[\sigma]$
  - The link load ρ is given by $\rho=\lambda E[\sigma]/C$
- Results correspond to ten averaged experiments (a dataset of nearly 17 million of packets)

Latency results

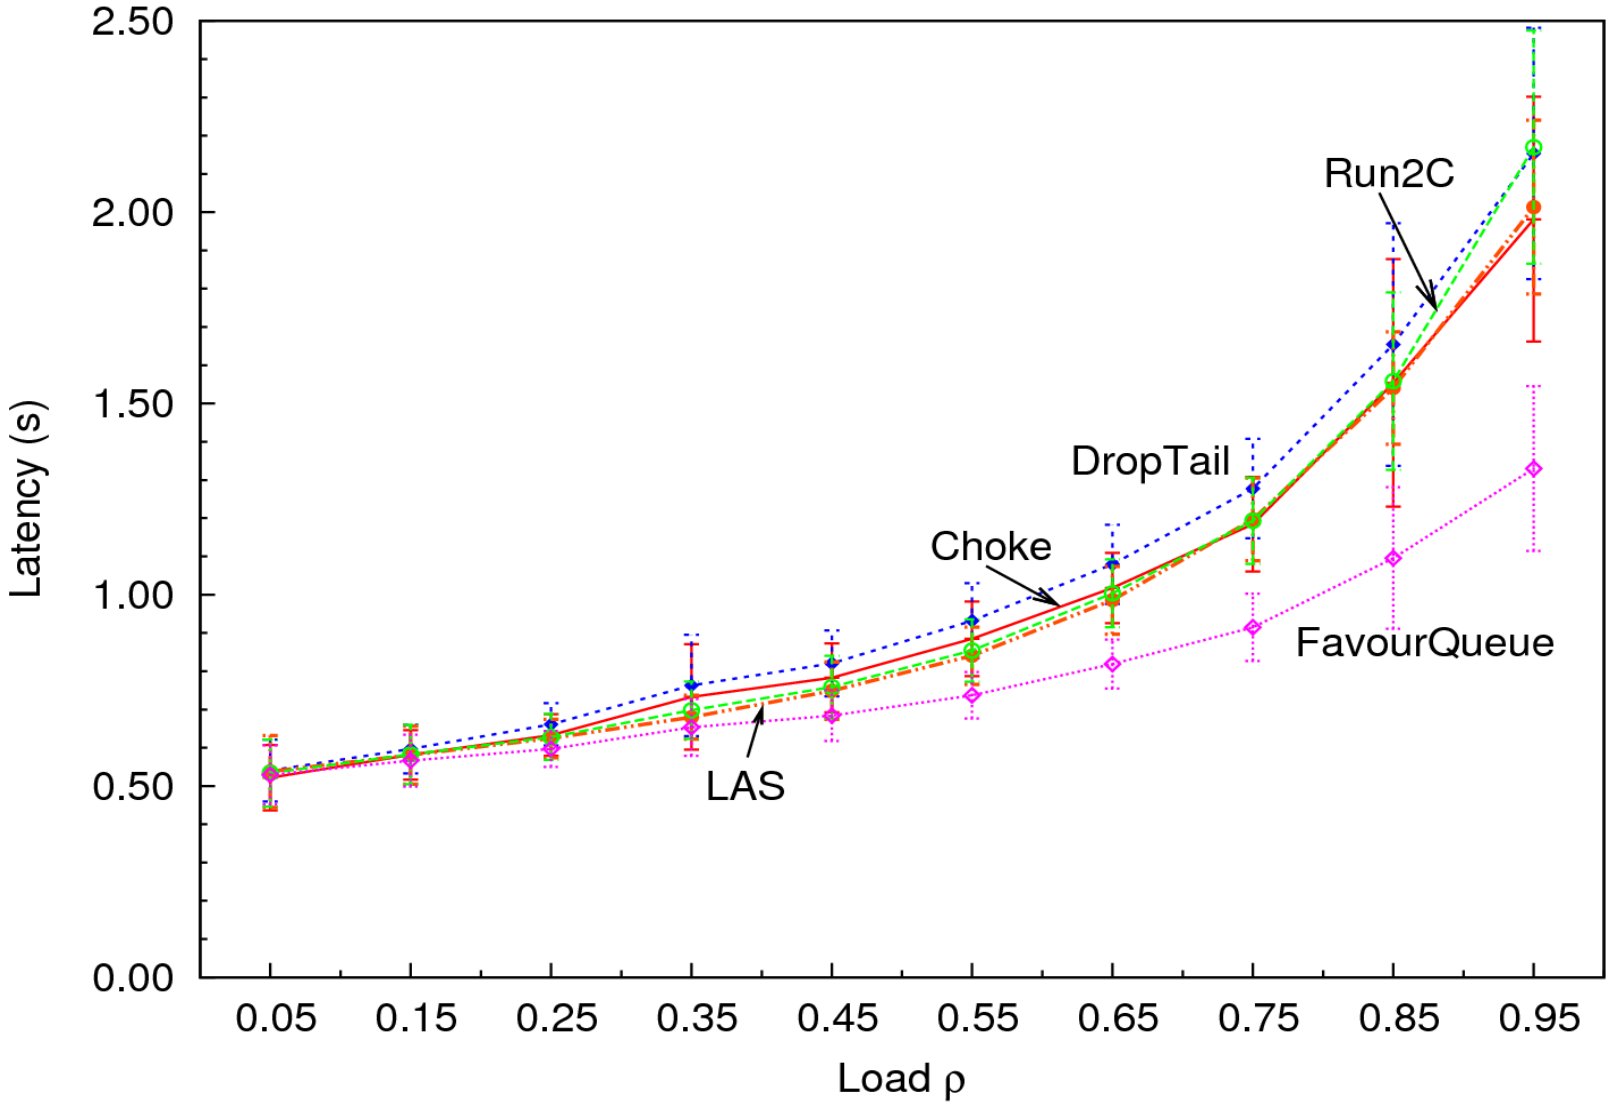

Drop ratio results

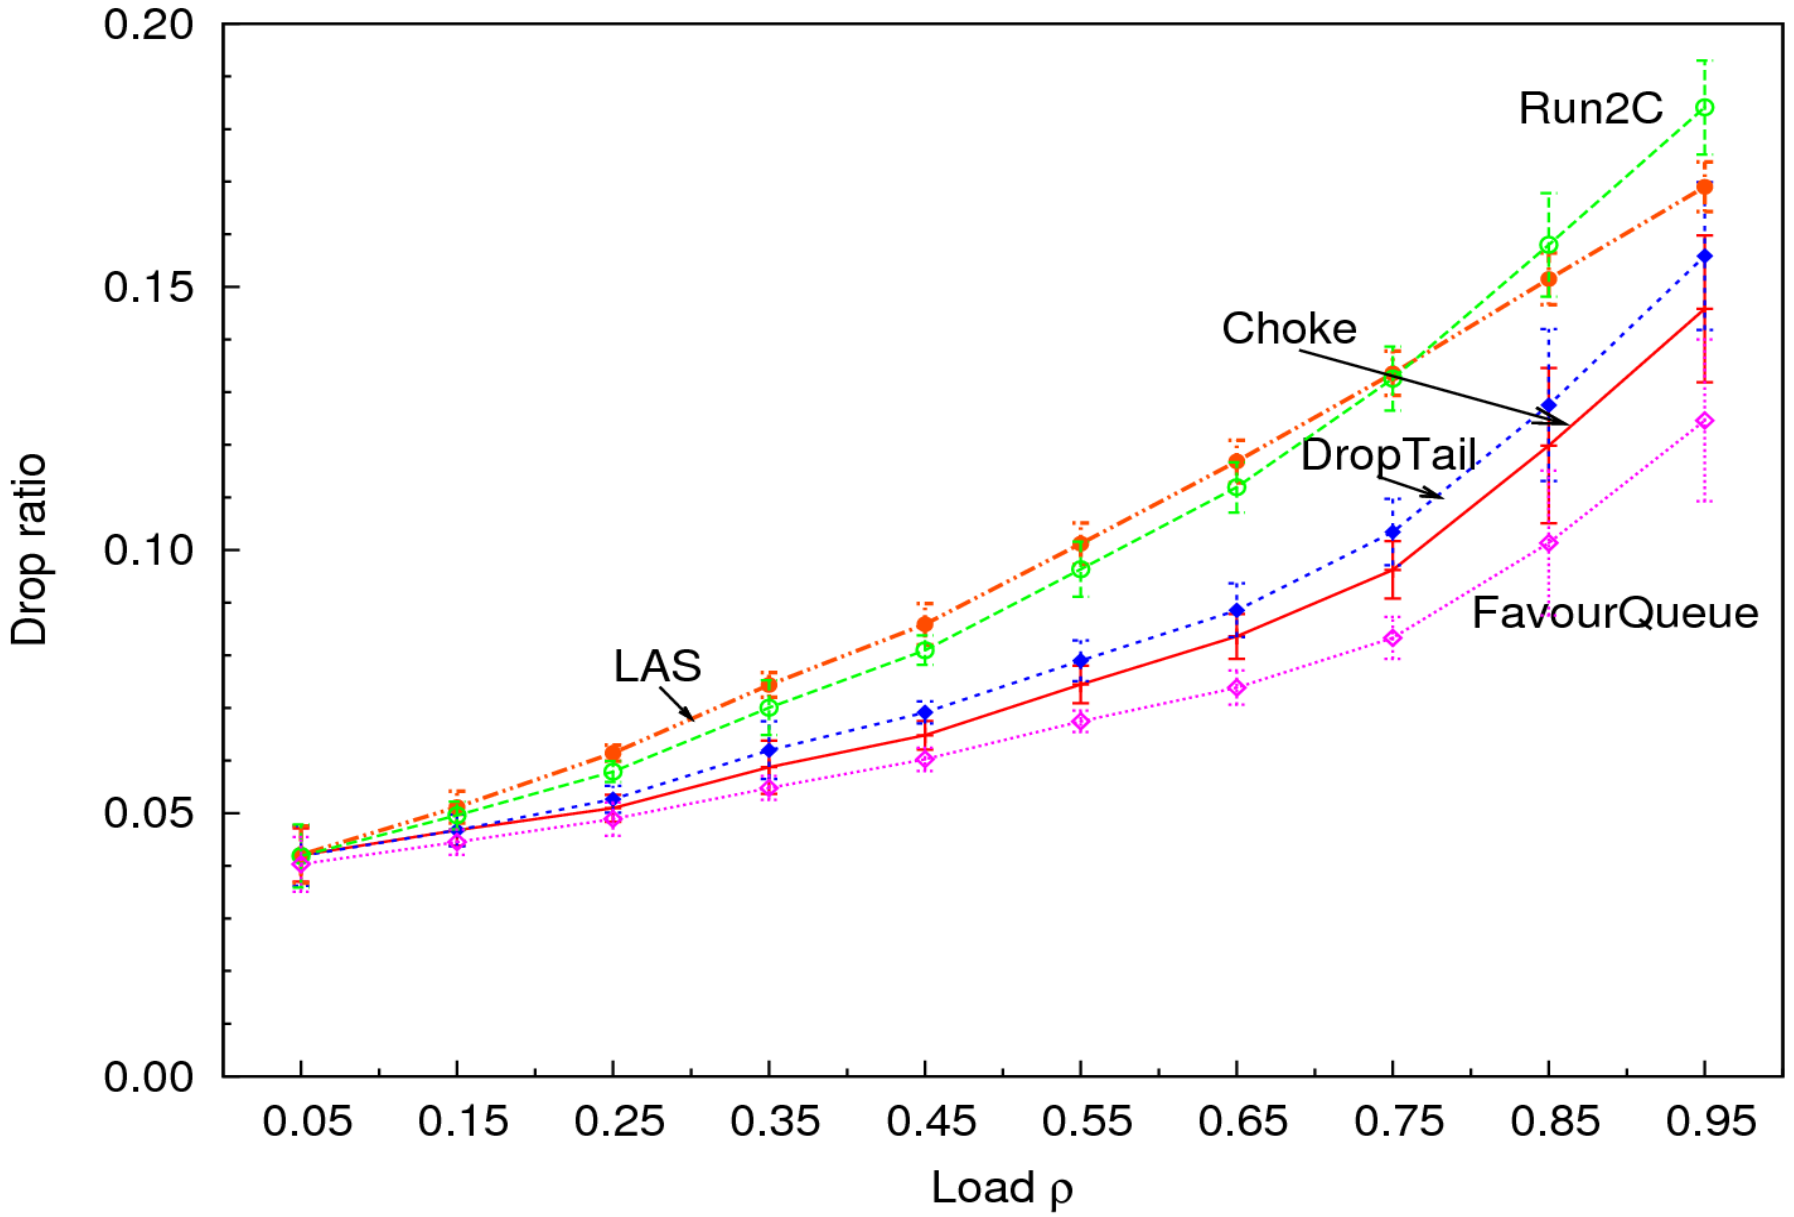

RTO recovery ratio

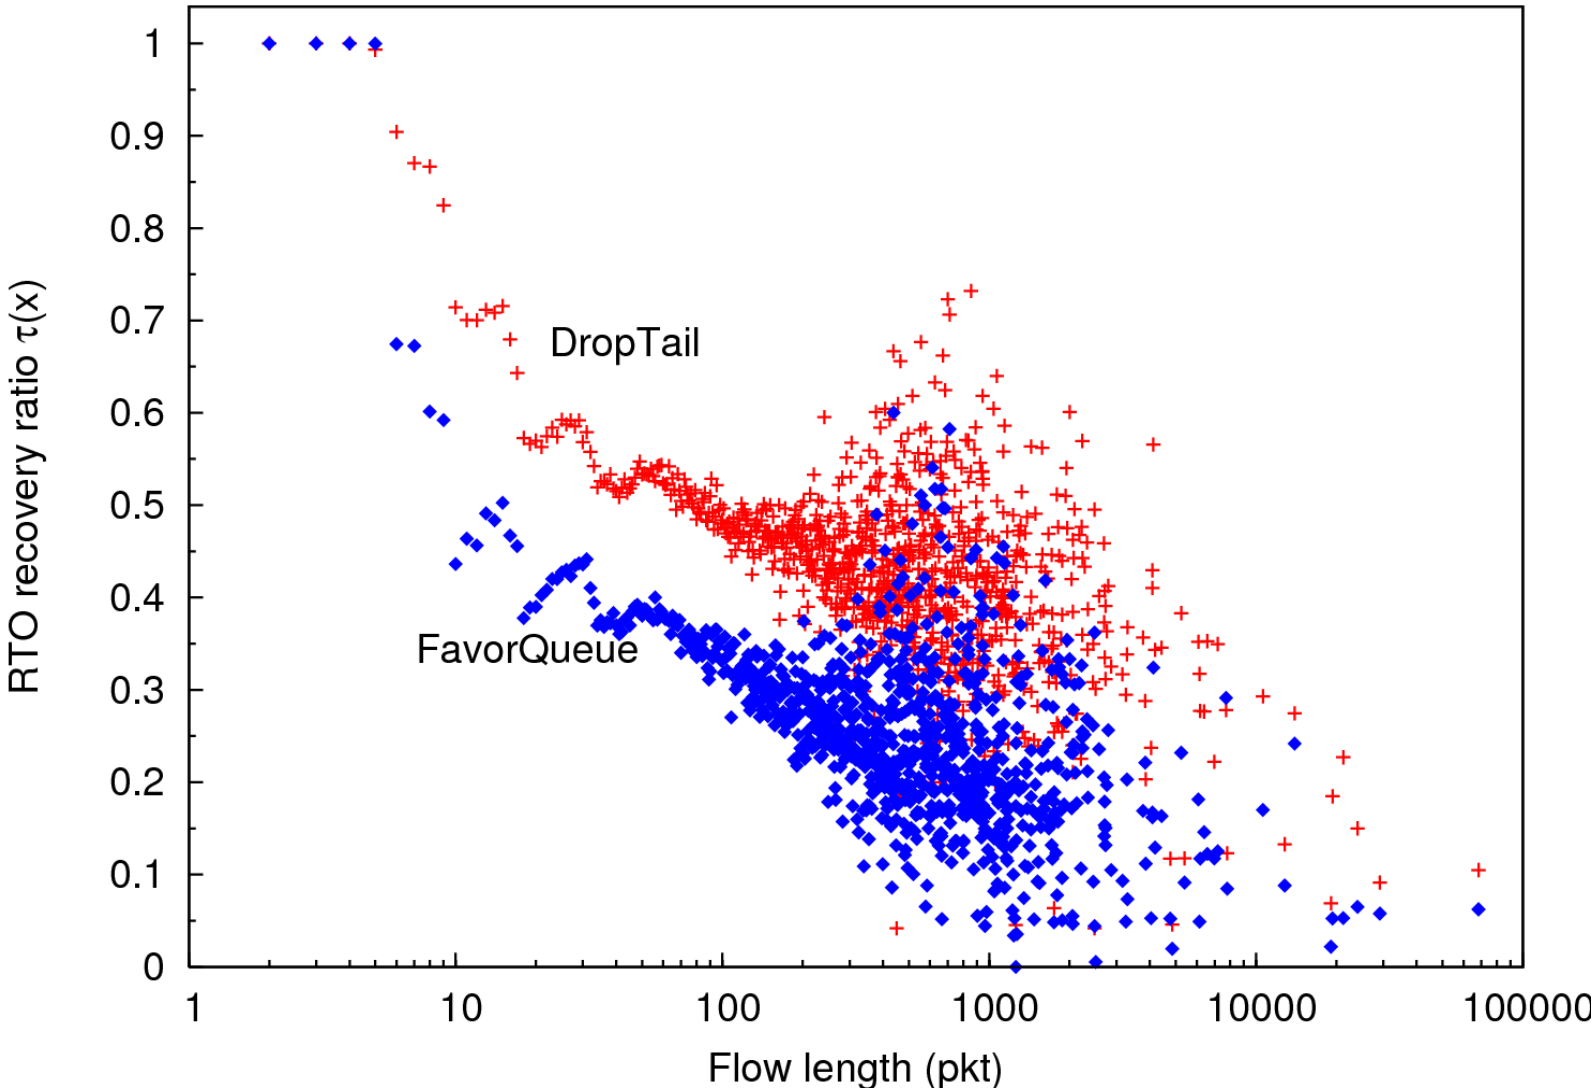

What happens?

- Flows with a size of two packets are always favored
- Middle sized flows that mainly remain in a slow-start phase are less favored
- Long TCP flows get a favoring ratio around 70%

Partial deployment impact

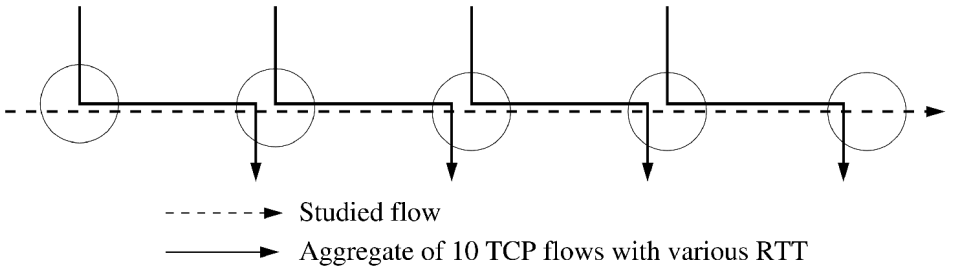

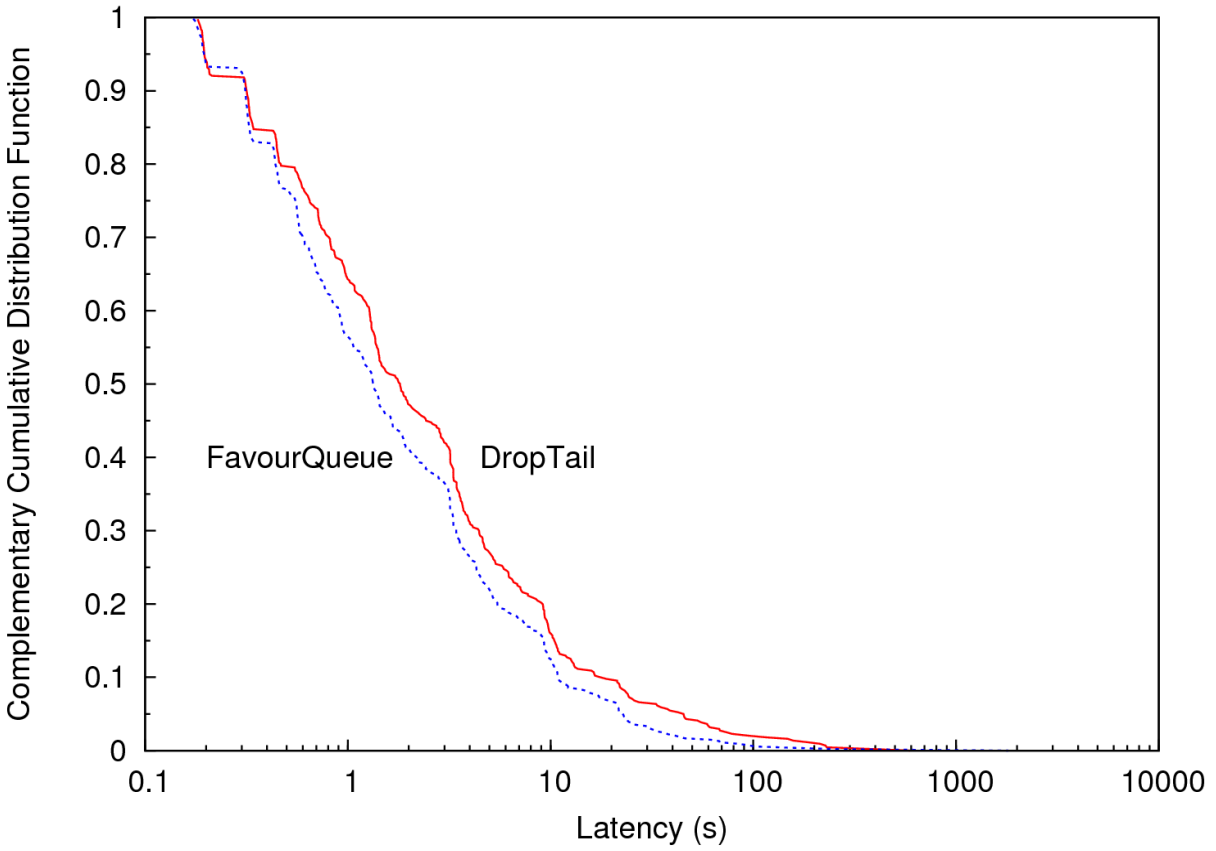

Conclusions

- All results and stochastic analysis available on arXiv
- In all experiments FaQ improves the latency up to 30%
- Stateless, no parameters
- Beneficial to all flows
- ns-2 code available (on-going TC qdisc code)