



Using BPF to Measure Latency at High Link Speeds

Simon Sundberg, Anna Brunstrom

Department of Computer Science, Karlstad University, Sweden

Email: simon.sundberg@kau.se

Network latency monitoring

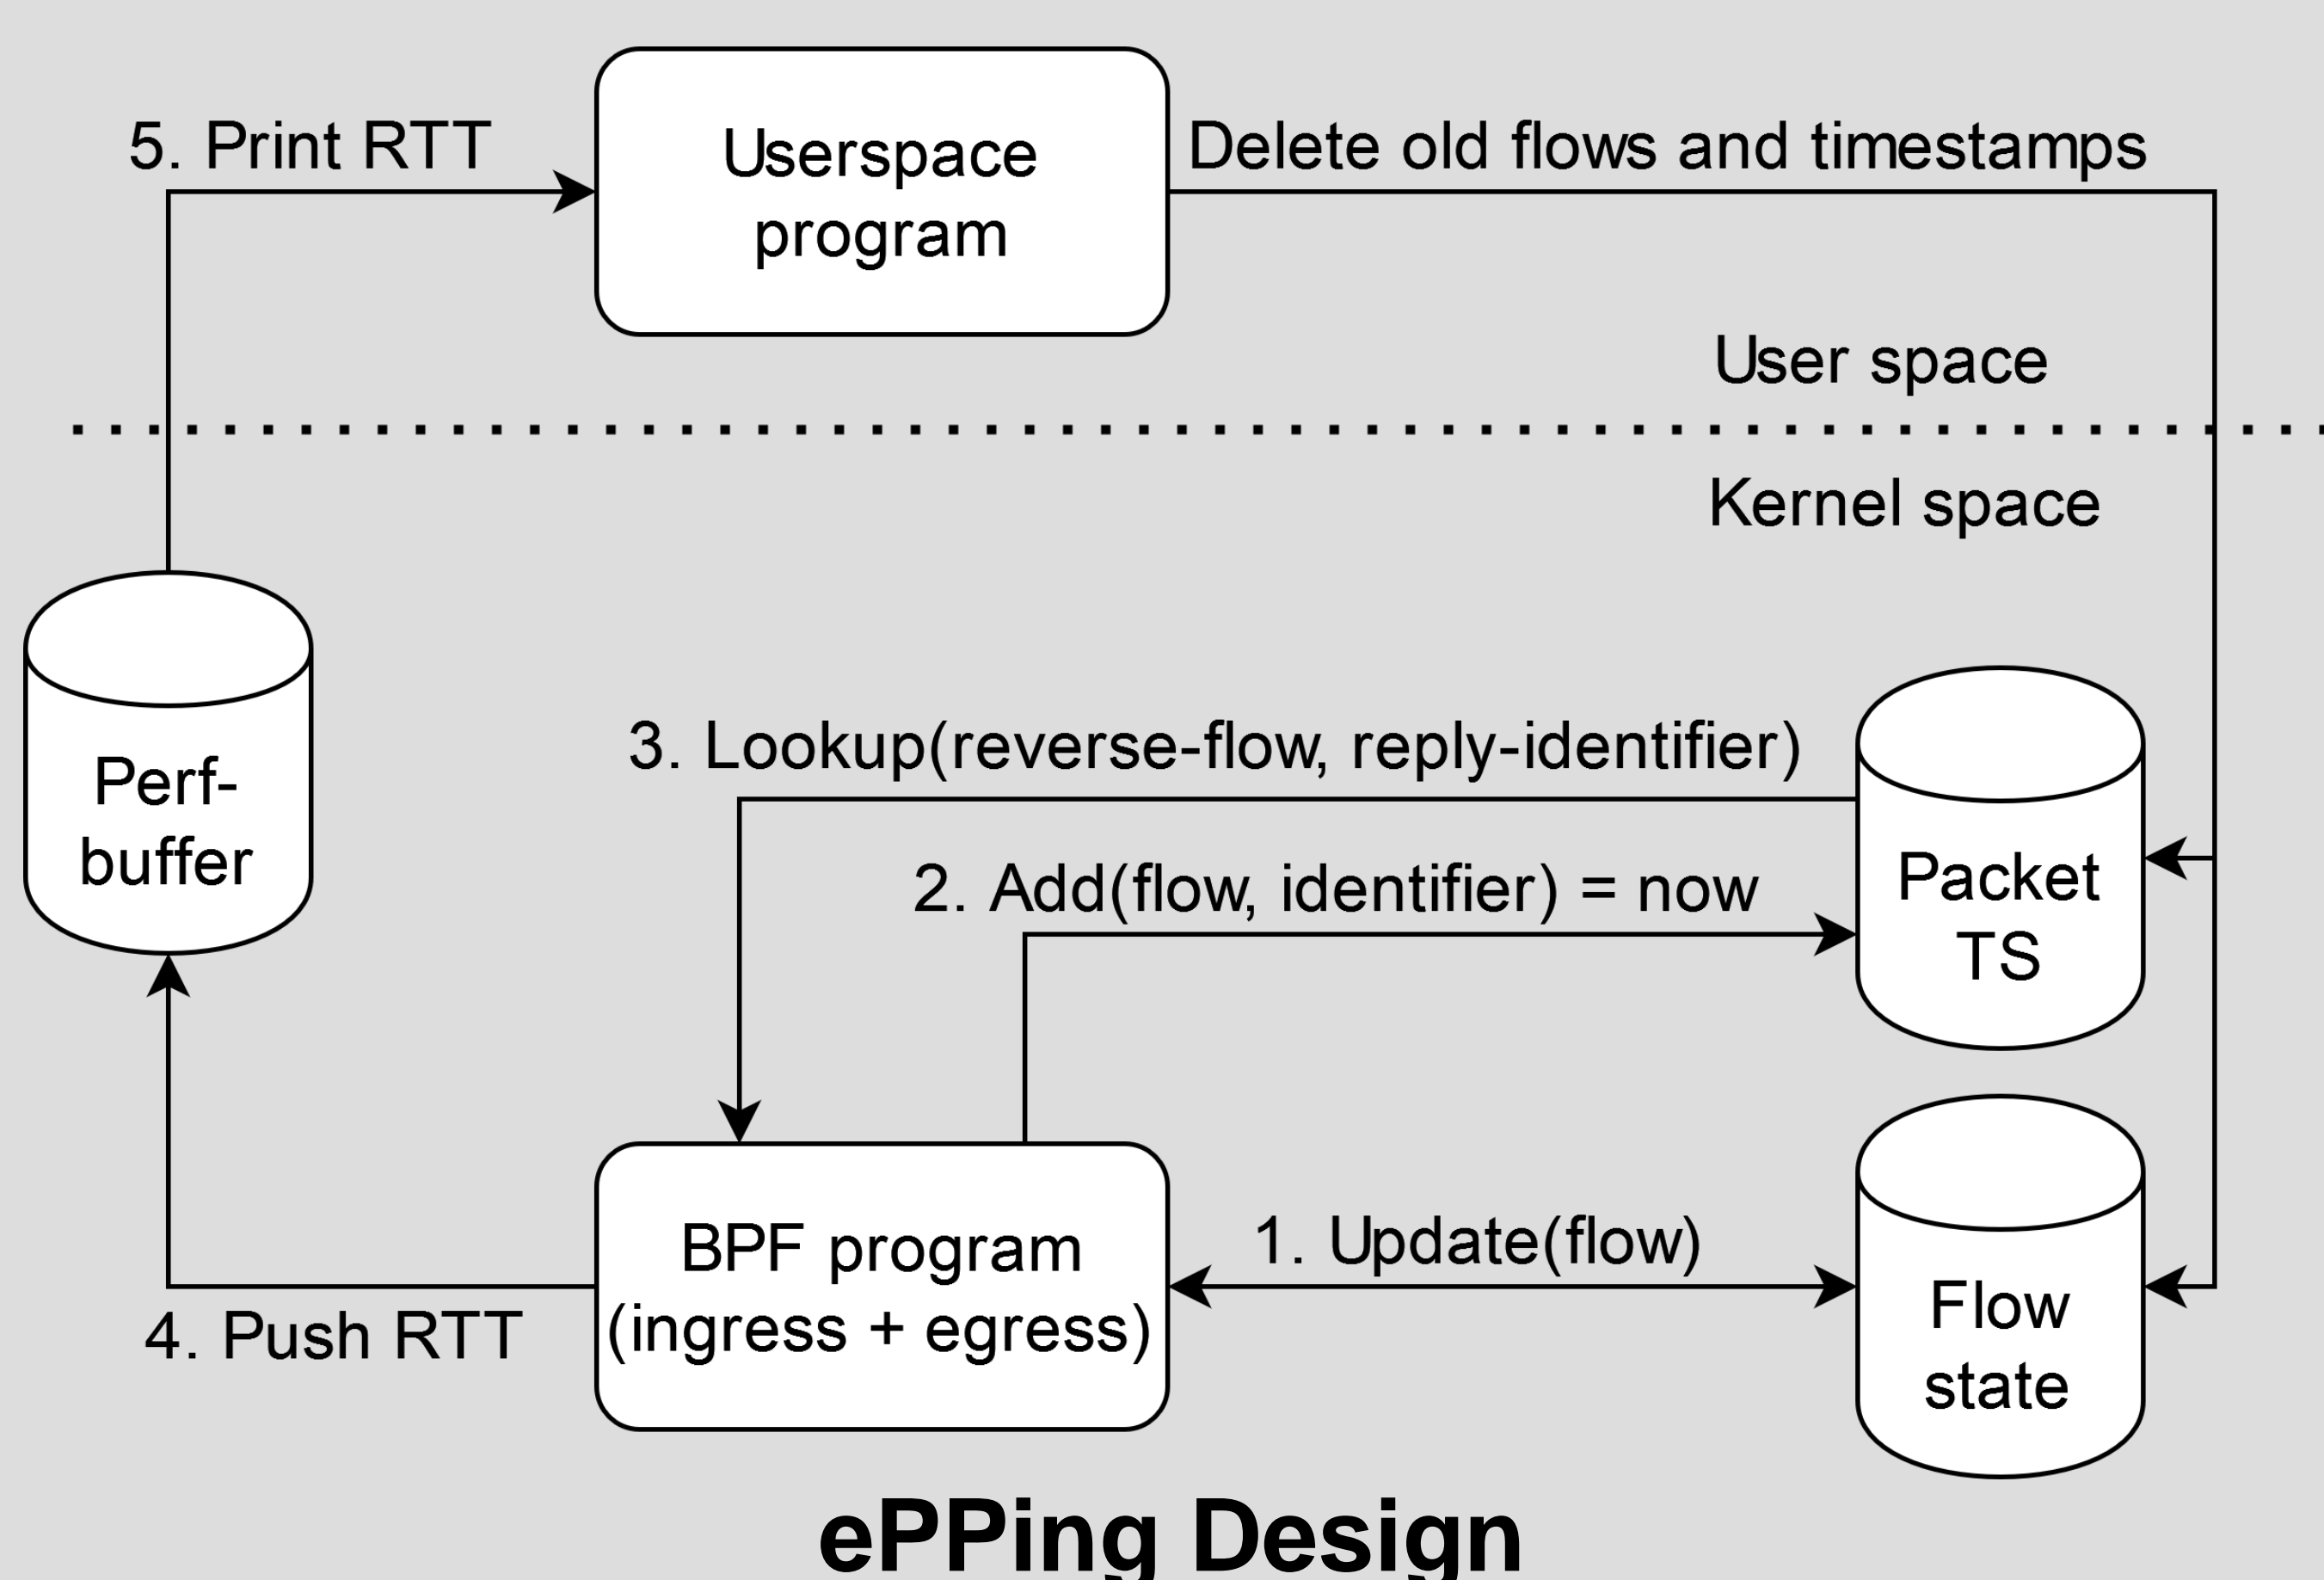
- Network latency monitoring is useful for determining QoE, troubleshooting, network performance analysis etc.
- ePPing is a general passive monitoring tool for measuring RTT designed to work at high line rates

Issues with existing monitoring tools

- Active tools like ping:
 - Add additional network traffic
 - Probes may be treated specially
- Passive userspace tools like PPing¹
 - + Work on any computer
 - High overhead, fail to keep up with link speed
- Passive P4 tools like Dapper² and P4RTT³
 - + Good performance at high link speeds
 - Require hardware support for P4

Our proposed solution - ePPing

- Extend PPing by using BPF to parse packets
 - + Runs on any Linux machine
 - + Low overhead at high line rates

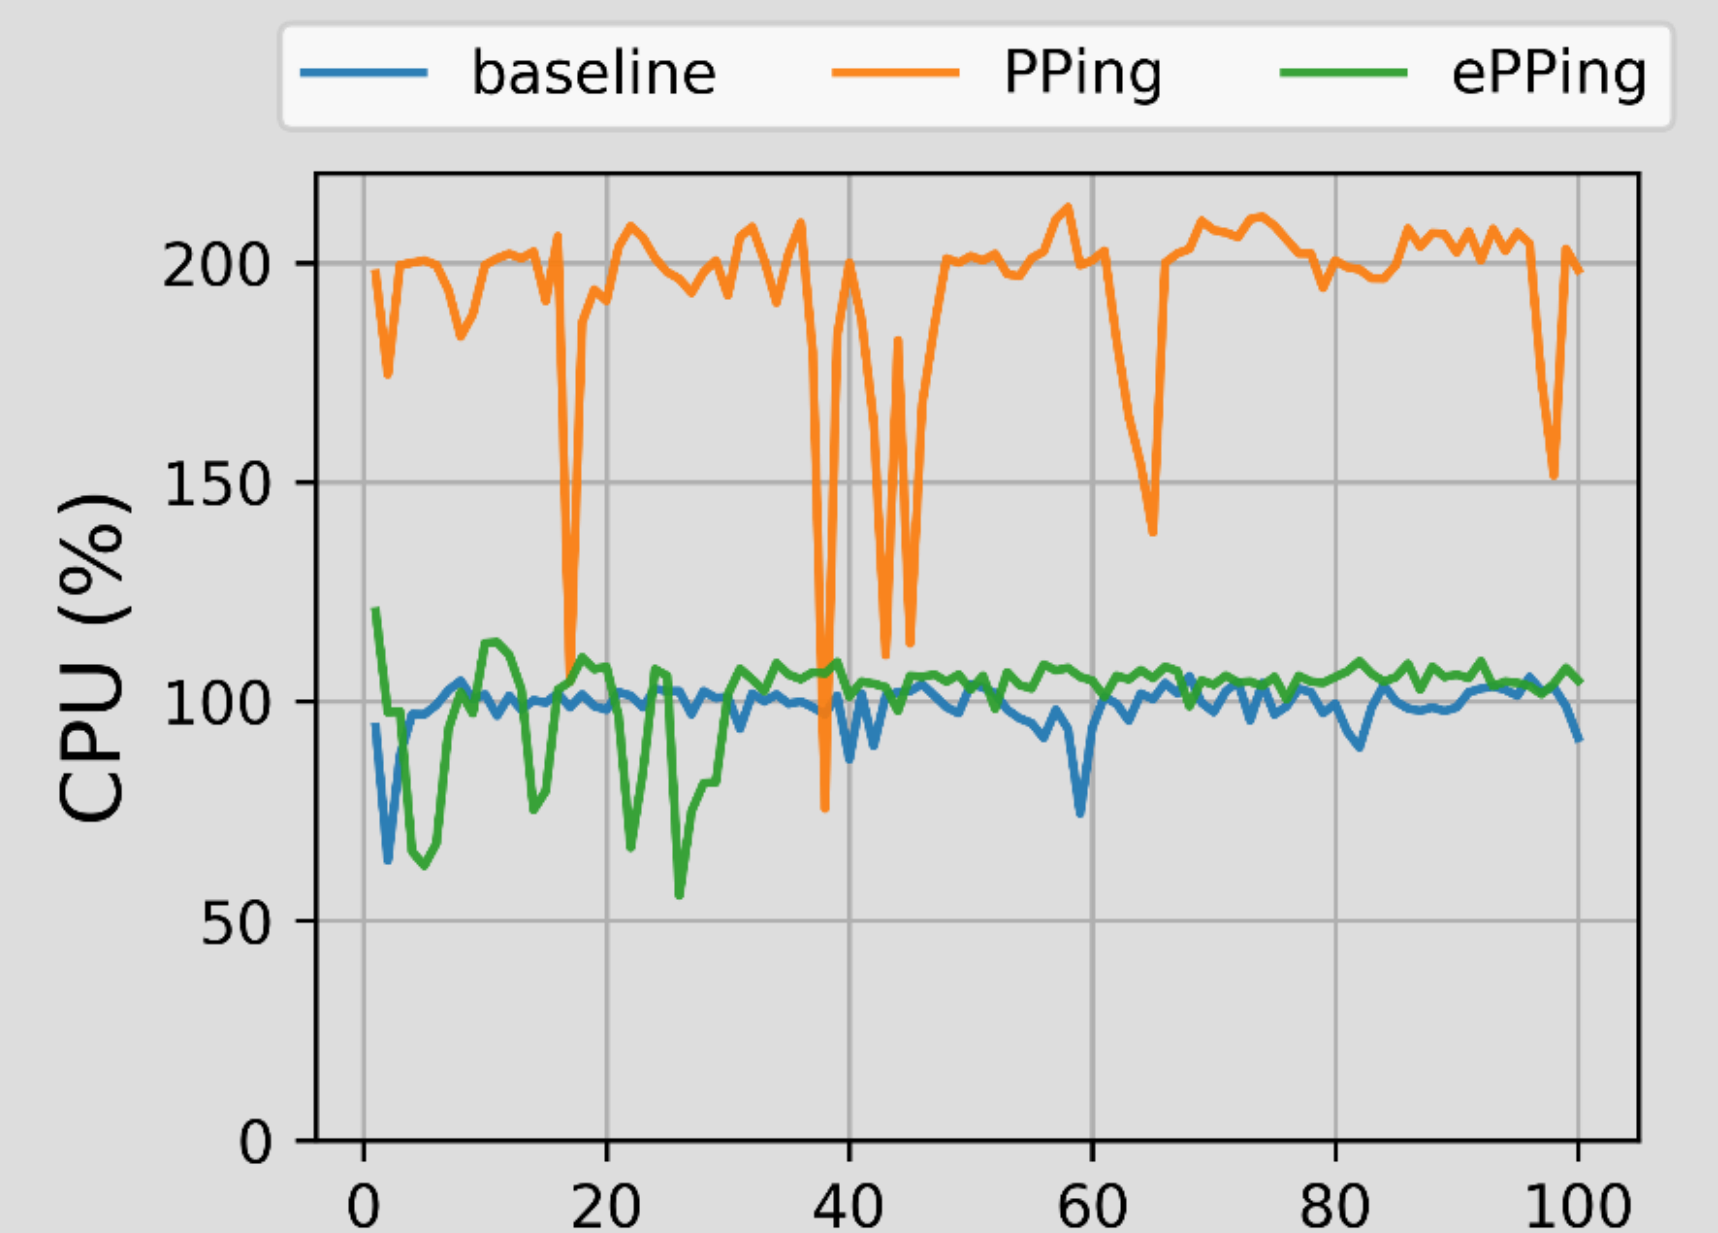


ePPing available at: <https://github.com/xdp-project/bpf-examples/tree/master/pping>

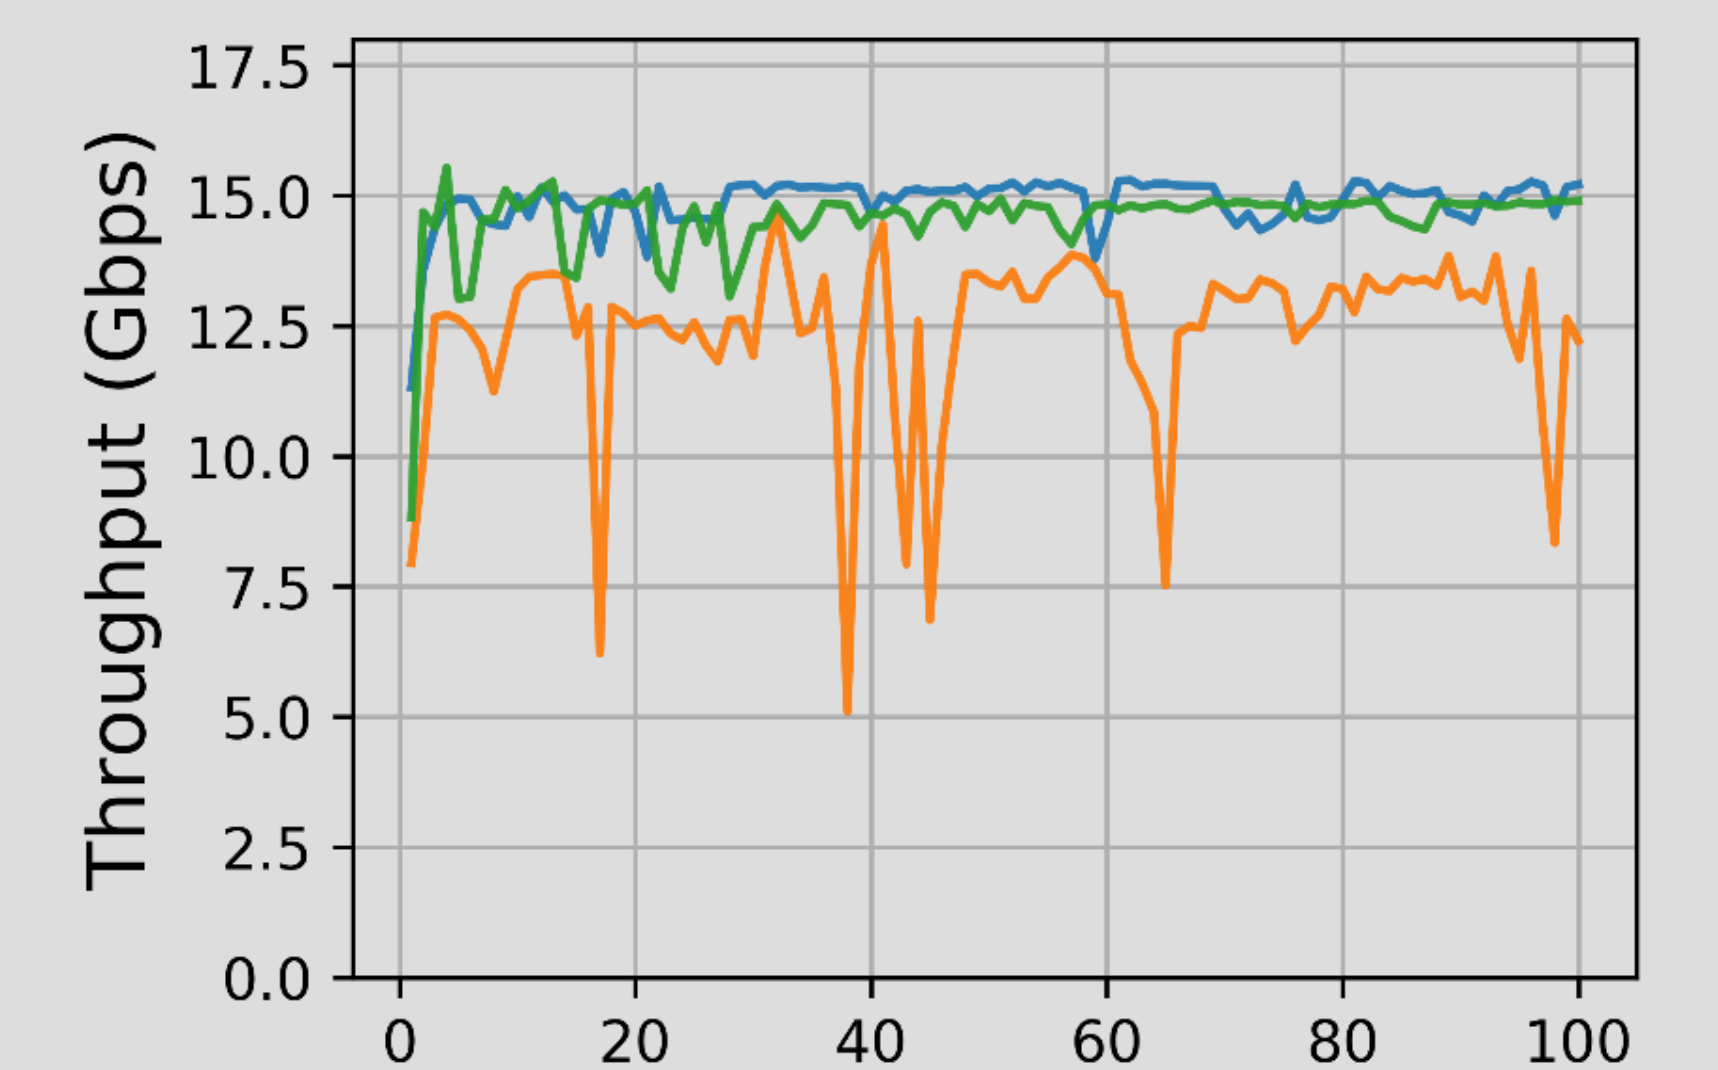


Initial results


- Low overhead



- Low impact on throughput



- Does not miss packets



Conclusions and ongoing work

- New BPF-based ePPing tool designed and implemented
- Initial results show promising improvements for ePPing compared to PPing
- TODO
 - Finalize ePPing tool
 - Perform extensive performance evaluation

Selected References

- [1] K. Nichols. 2018. <https://github.com/pollere/pping>
- [2] M. Ghasemi et al. 2017. Dapper: Data plane performance diagnosis of TCP, in *SOSR '17*. DOI: 10.1145/3050220.3050228
- [3] X. Chen et al. 2020. Measuring TCP round-trip time in the data plane, in *SPIN '20*. DOI: 10.1145/3405669.3405823