

Understanding the Performance of Different Packet Reception and Timestamping Methods in Linux

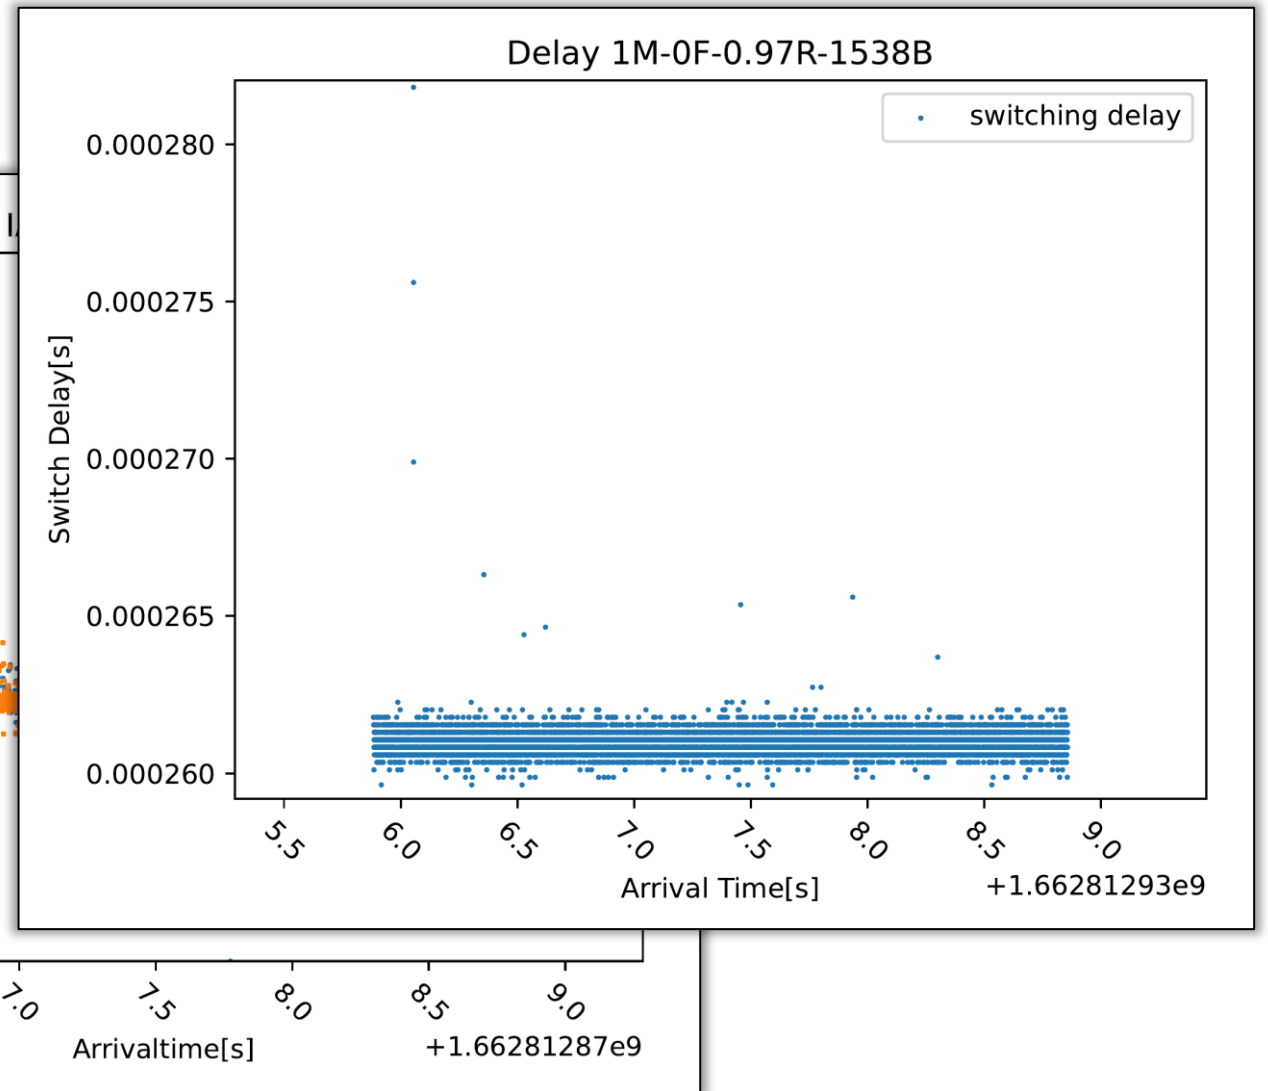
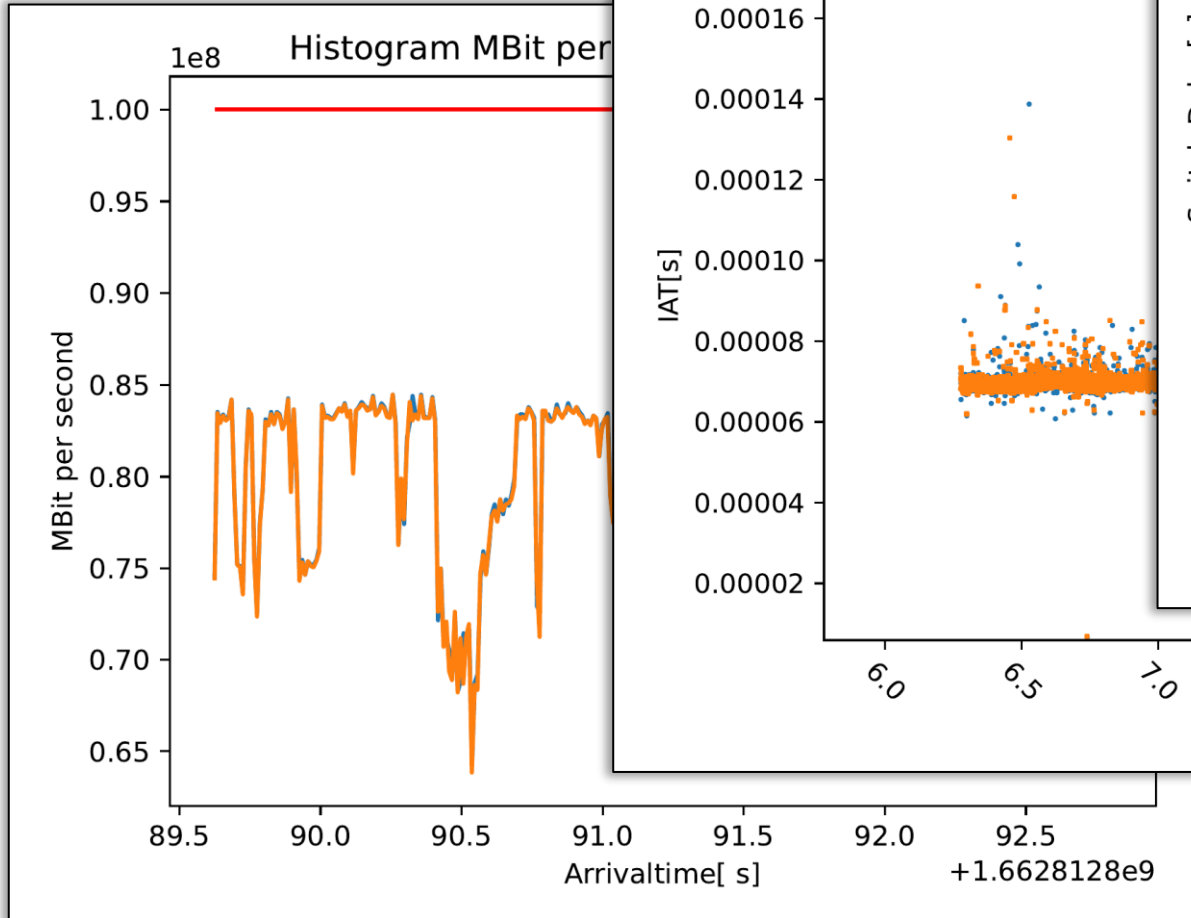
KuVS Fachgespräch – WueWoWAS'23 – Würzburg



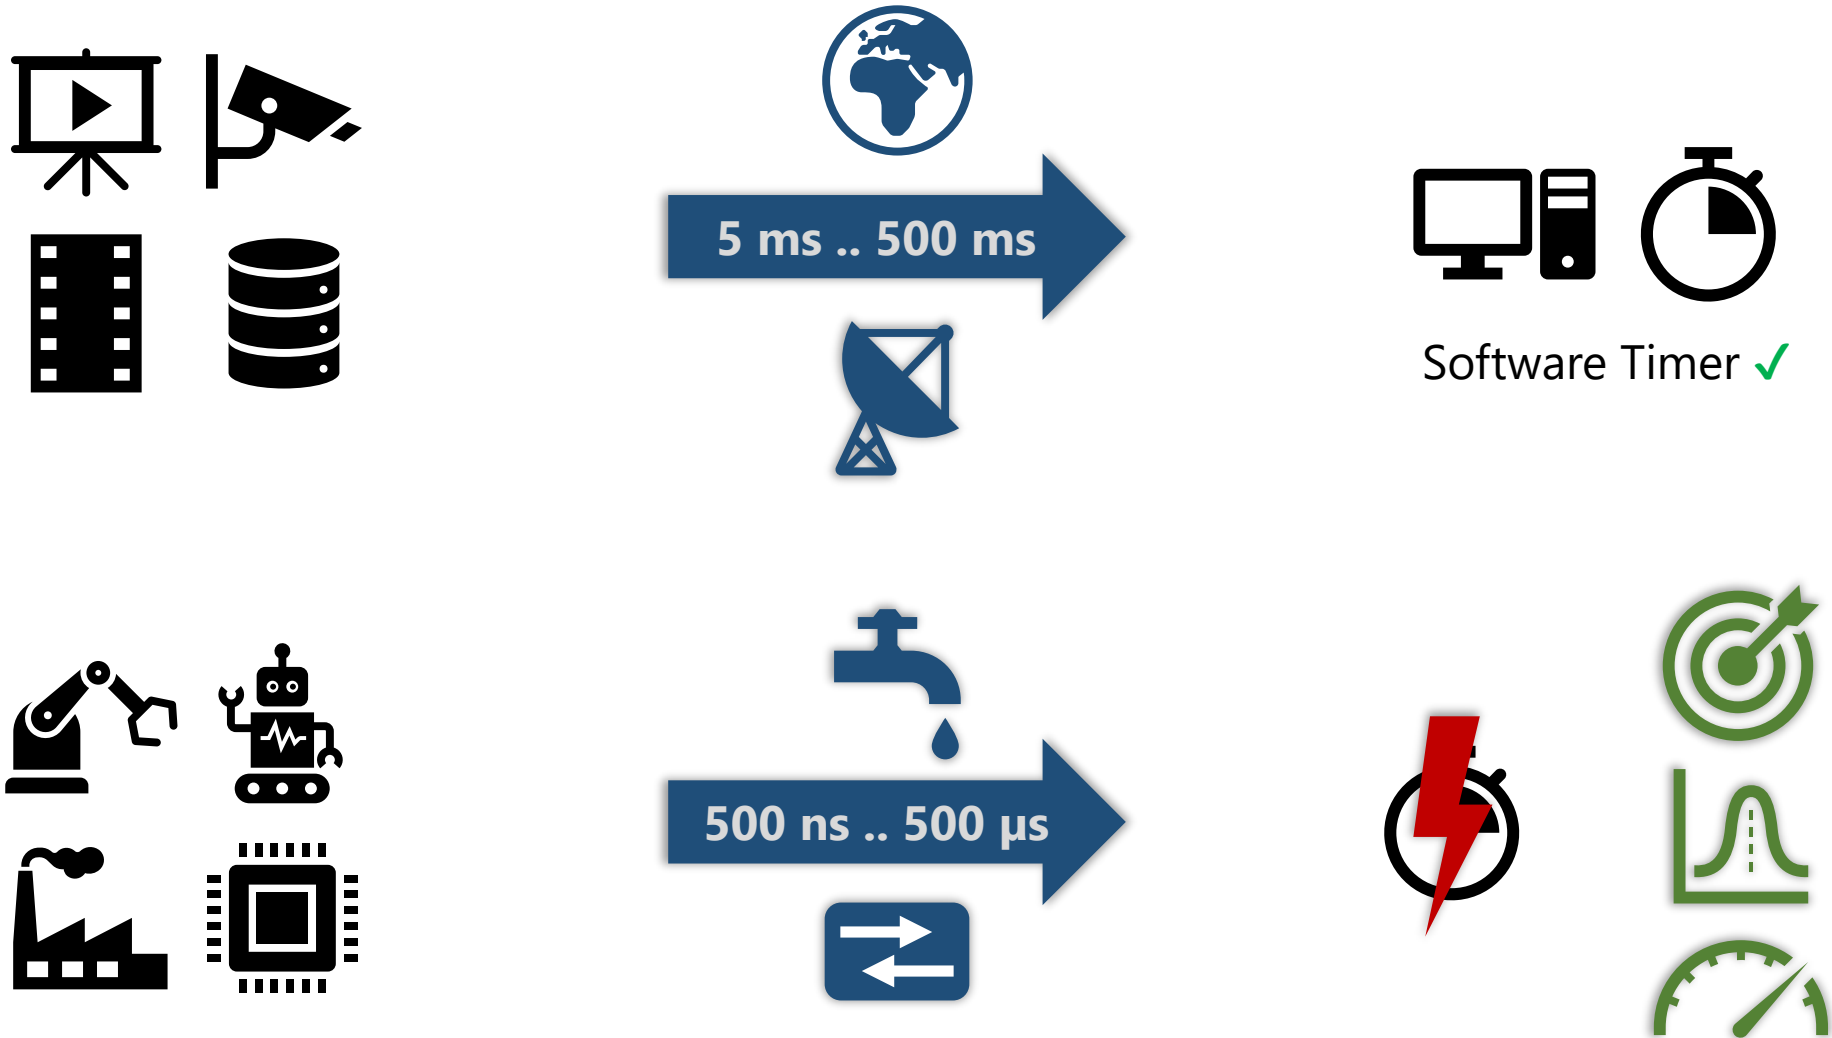
Alexej Grigorjew, Lukas Kilian Schumann, Philip Diederich,
Tobias Hoßfeld, Wolfgang Kellerer

alexej.grigorjew@uni-wuerzburg.de

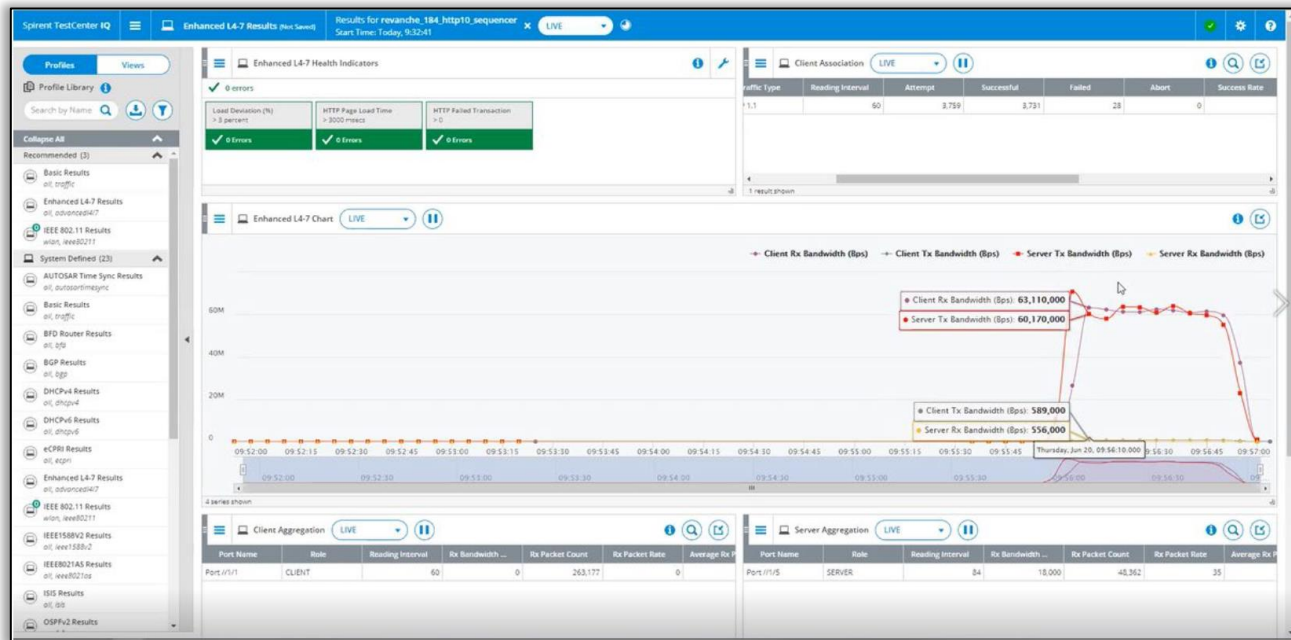
We All Do Measurements




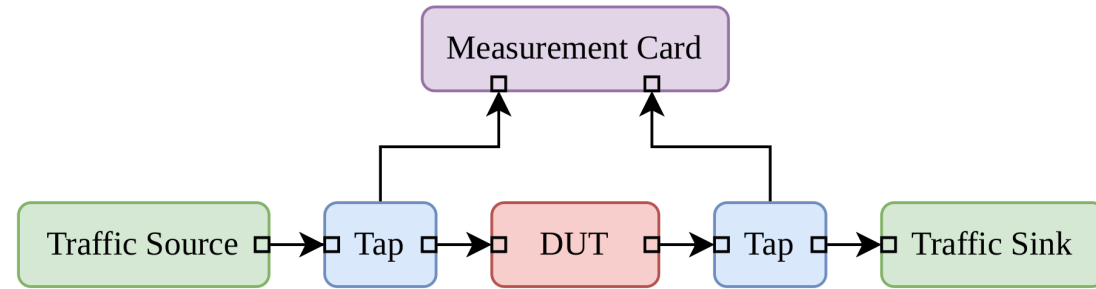
What Do We Need From Our Measurements?



Measurement Equipment & Software can be Expensive




Affordable Alternative: Hardware Timestamping with Commodity NICs



```
demo@host ~ % lspci | grep Ethernet
00:1f.6 Ethernet controller: Intel Corporation ... I219-V
03:00.0 Ethernet controller: Intel Corporation I350 ...
03:00.1 Ethernet controller: Intel Corporation I350 ...
05:00.0 Ethernet controller: Intel Corporation I350 ...
05:00.1 Ethernet controller: Intel Corporation I350 ...
demo@host ~ %
```

How Good Is It? Measurement Setup & Parameters



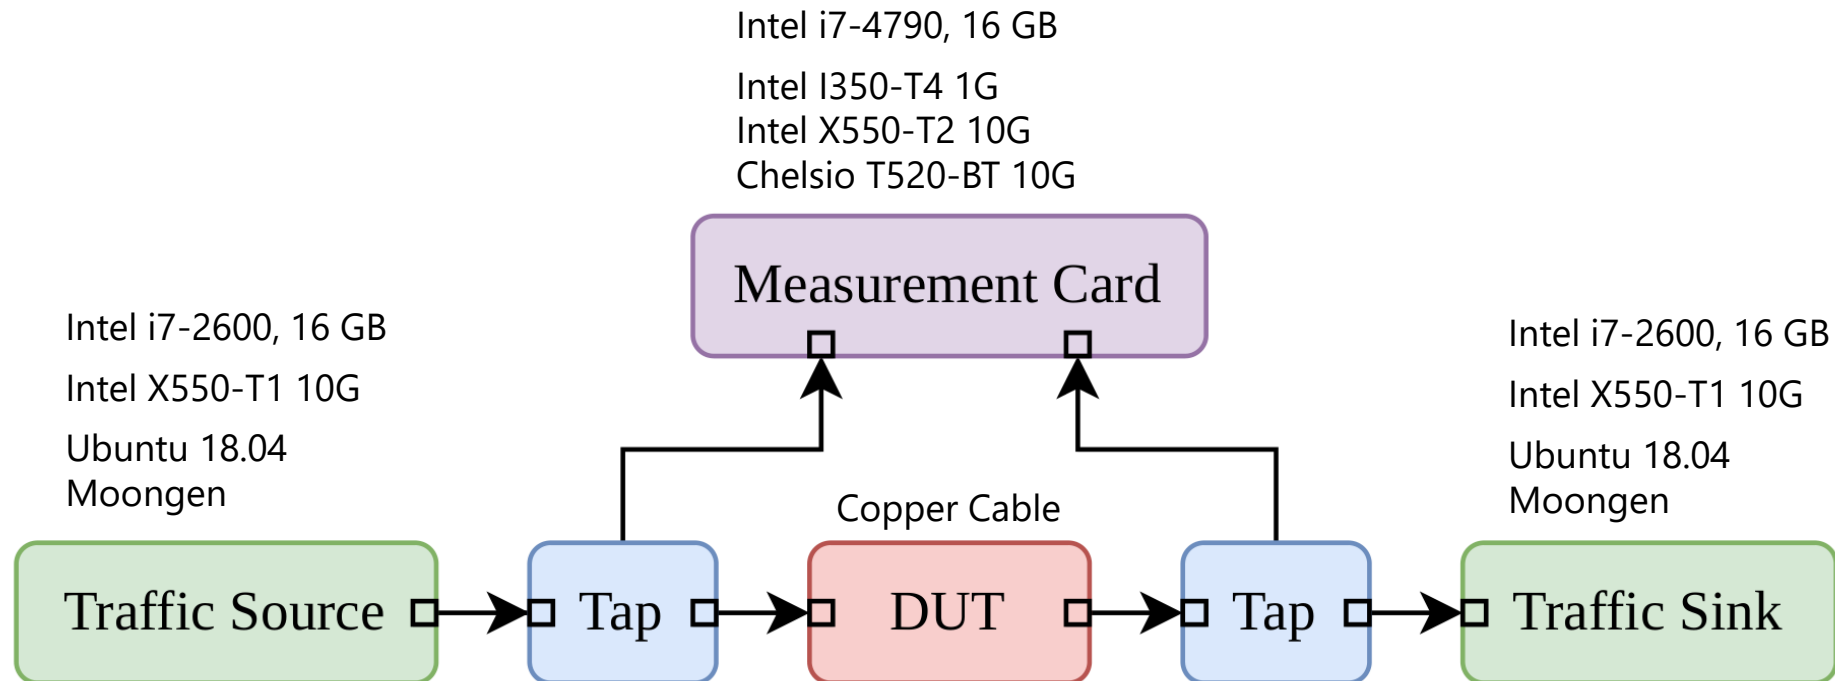
Technique

rx_ring, hw/sw

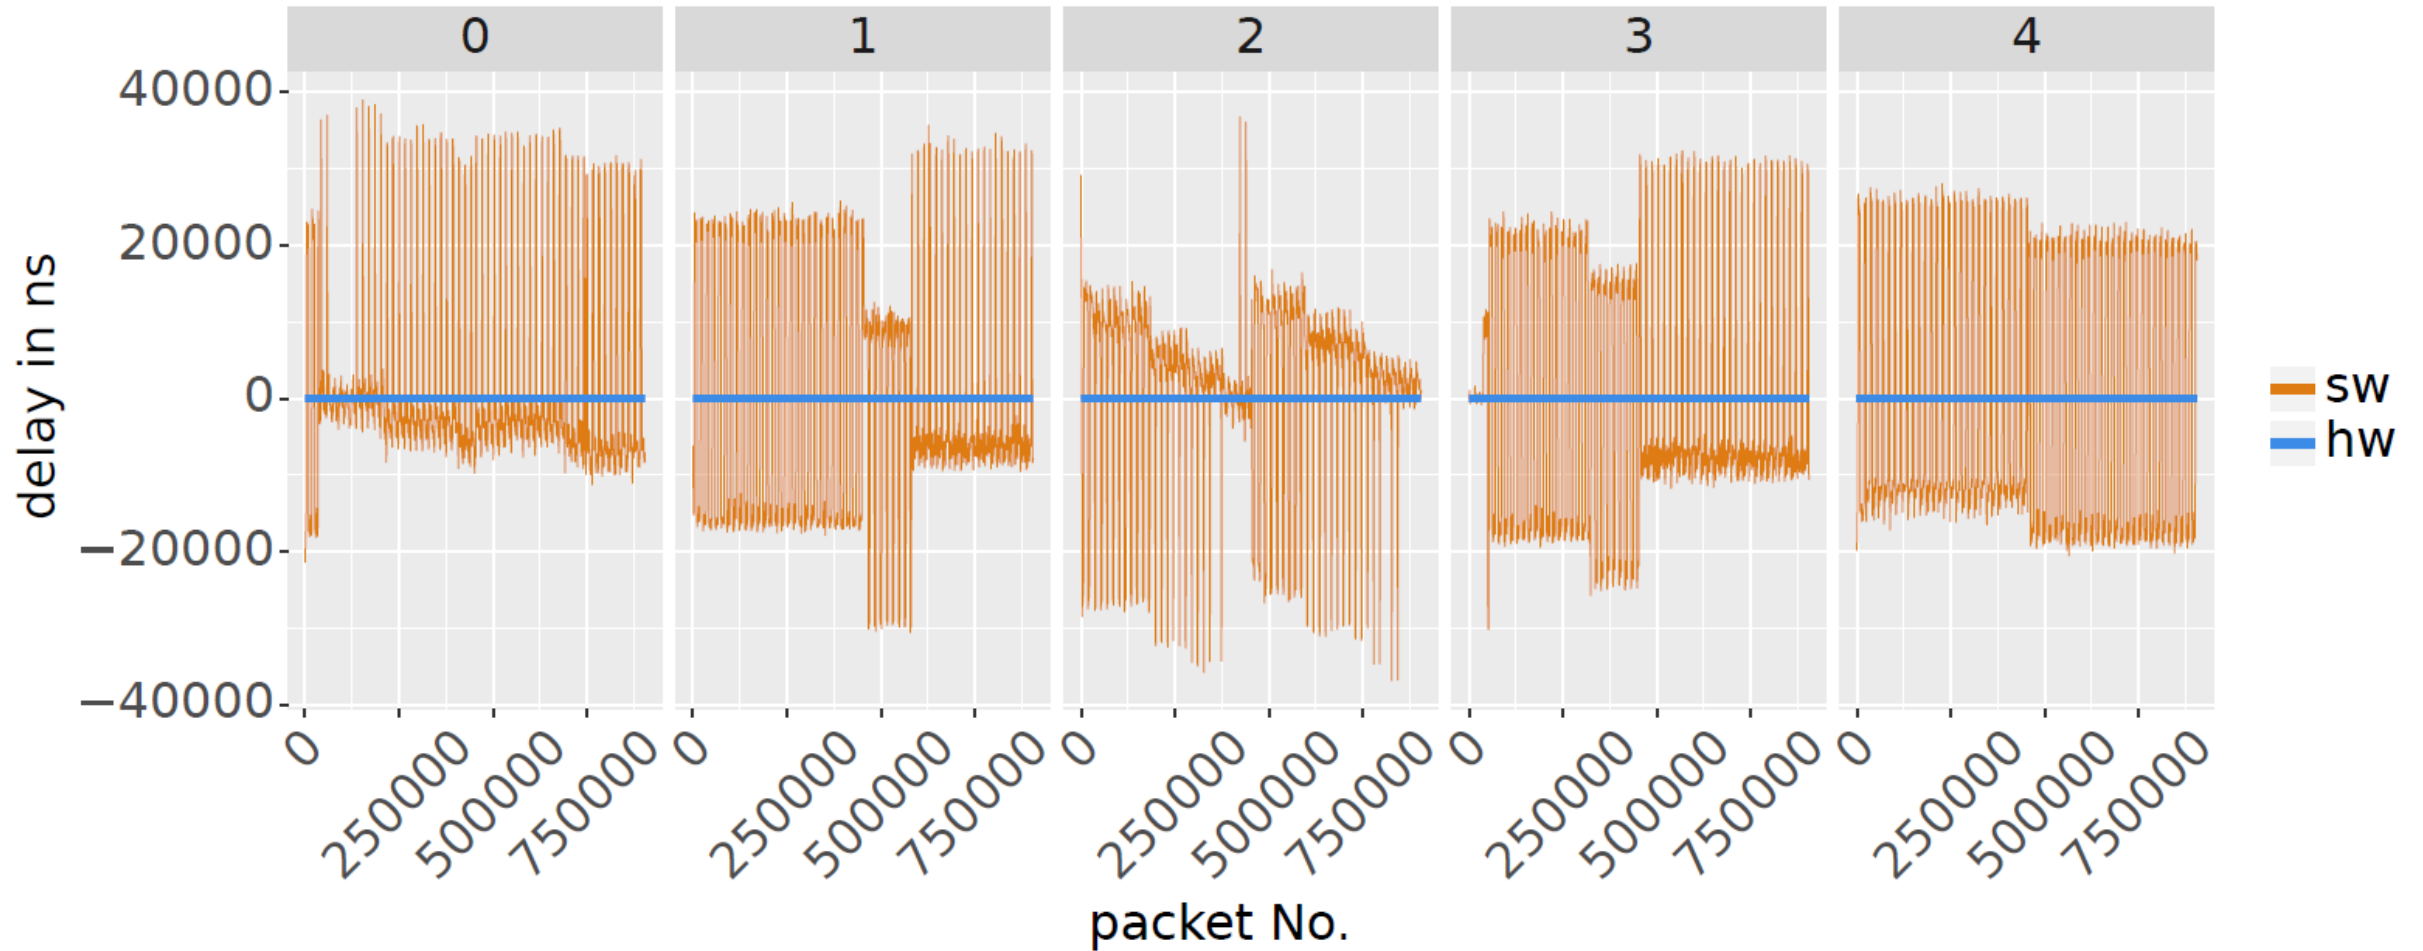
recvmmsg, hw/sw

libpcap, hw/sw


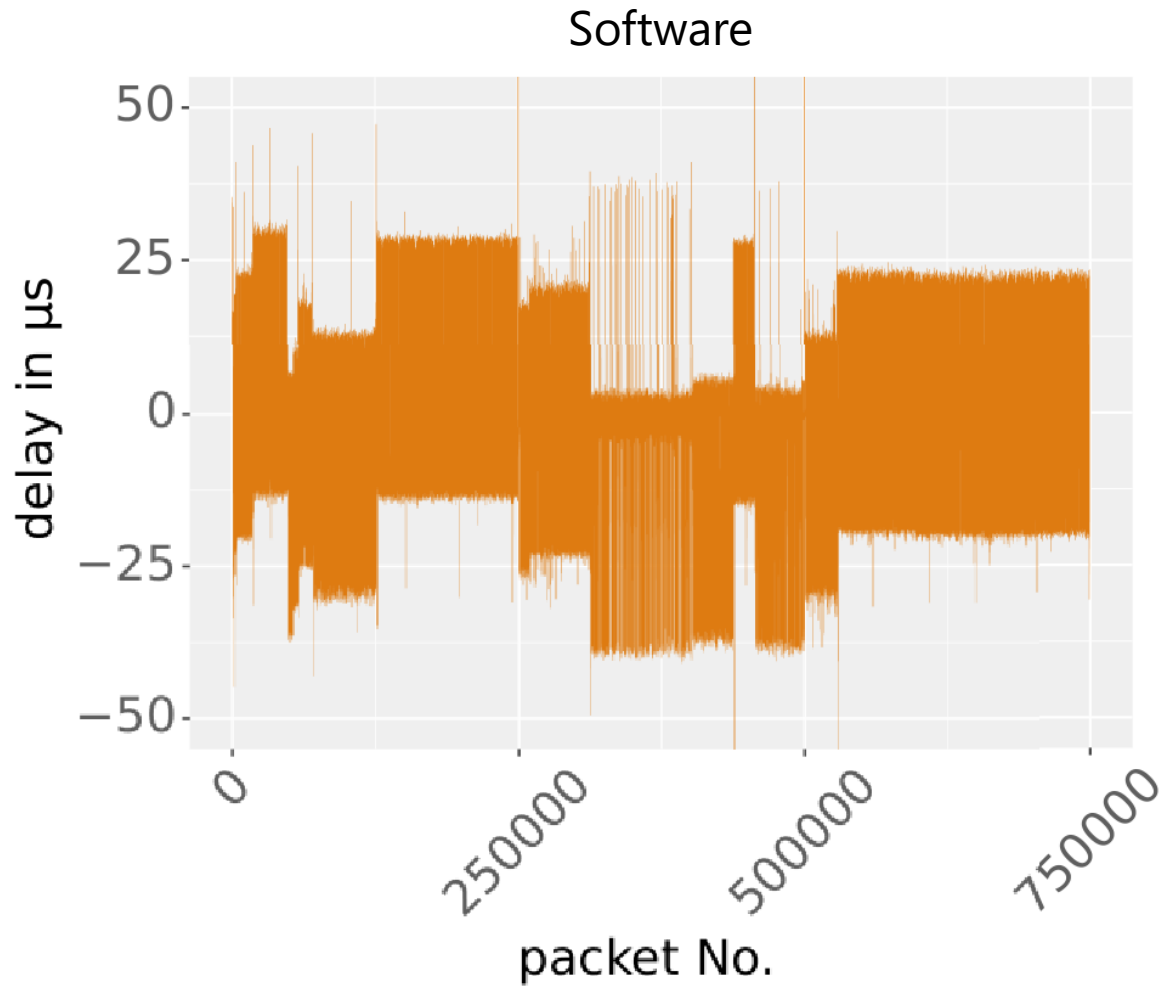
xdp



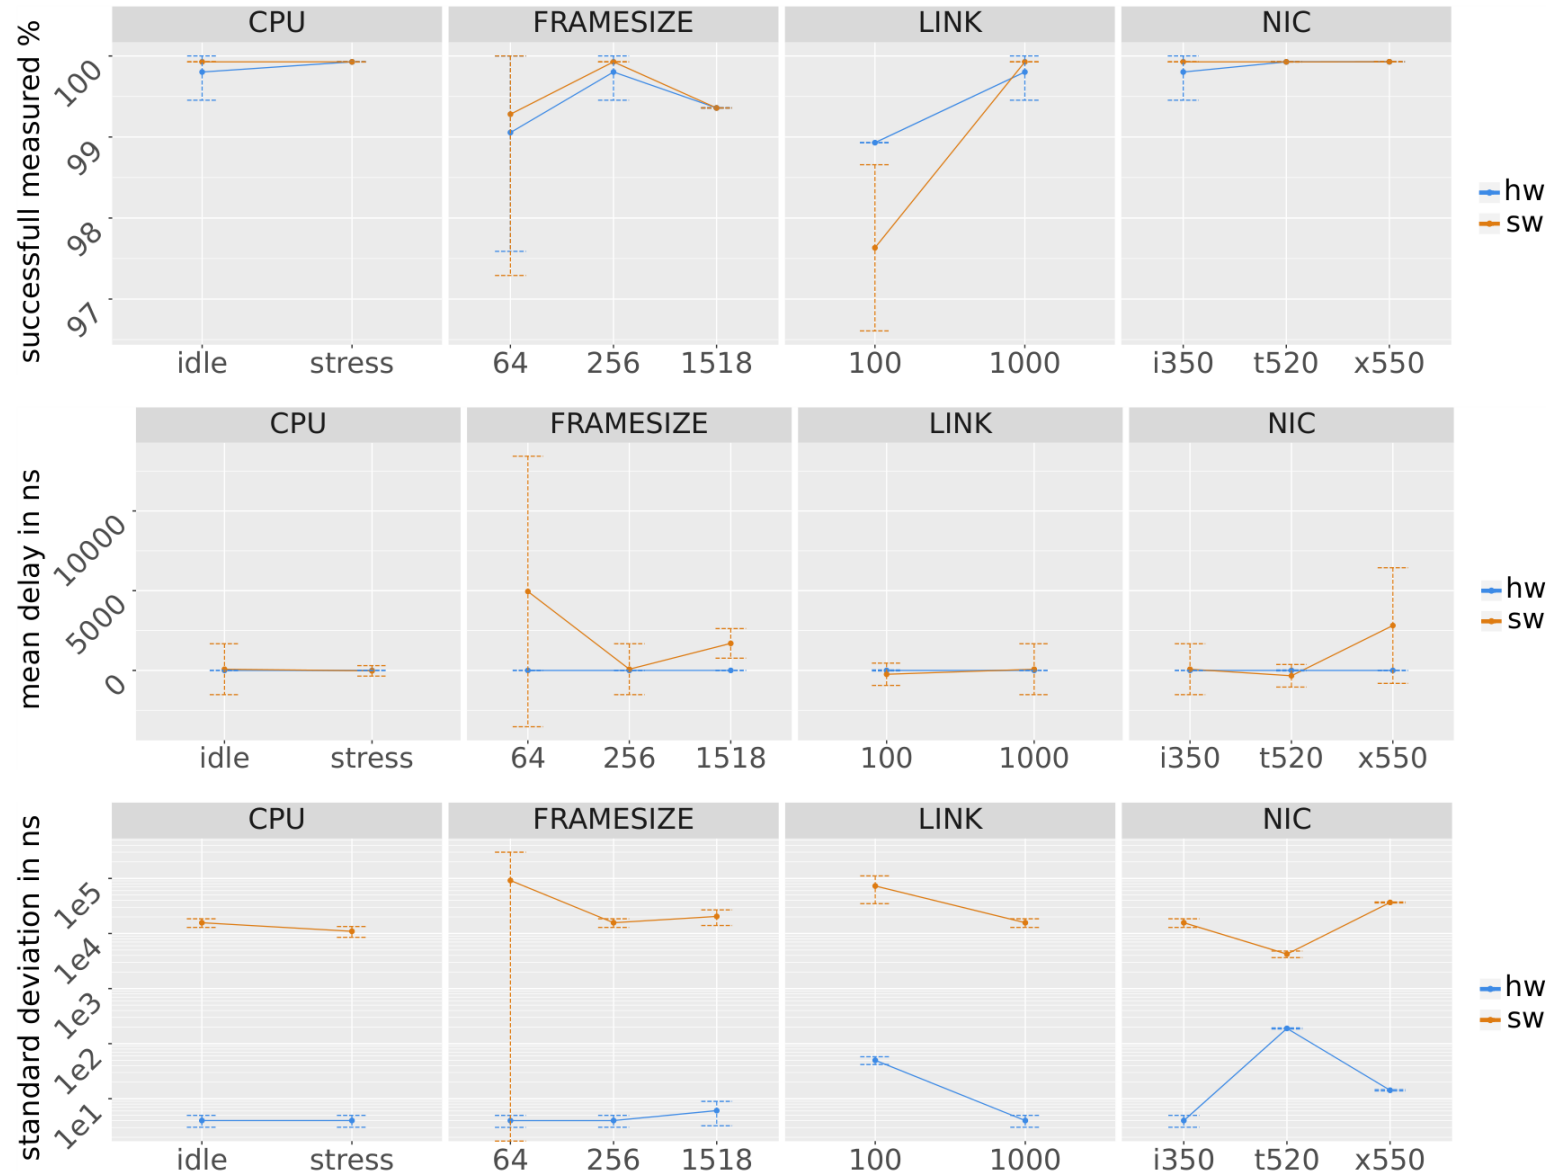
Time Series Plots – Software vs. Hardware Timestamps



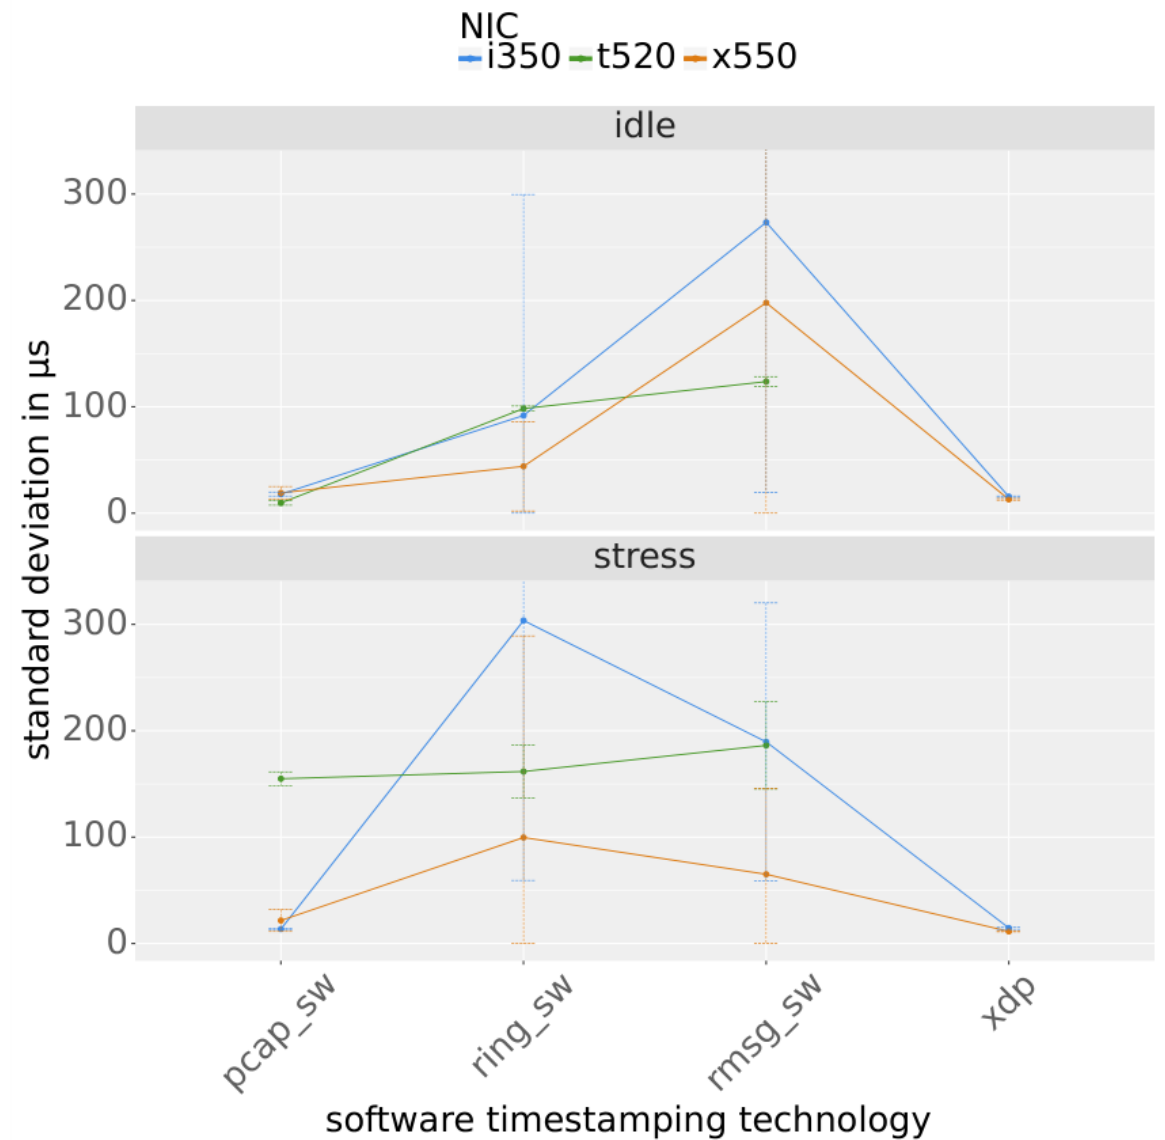
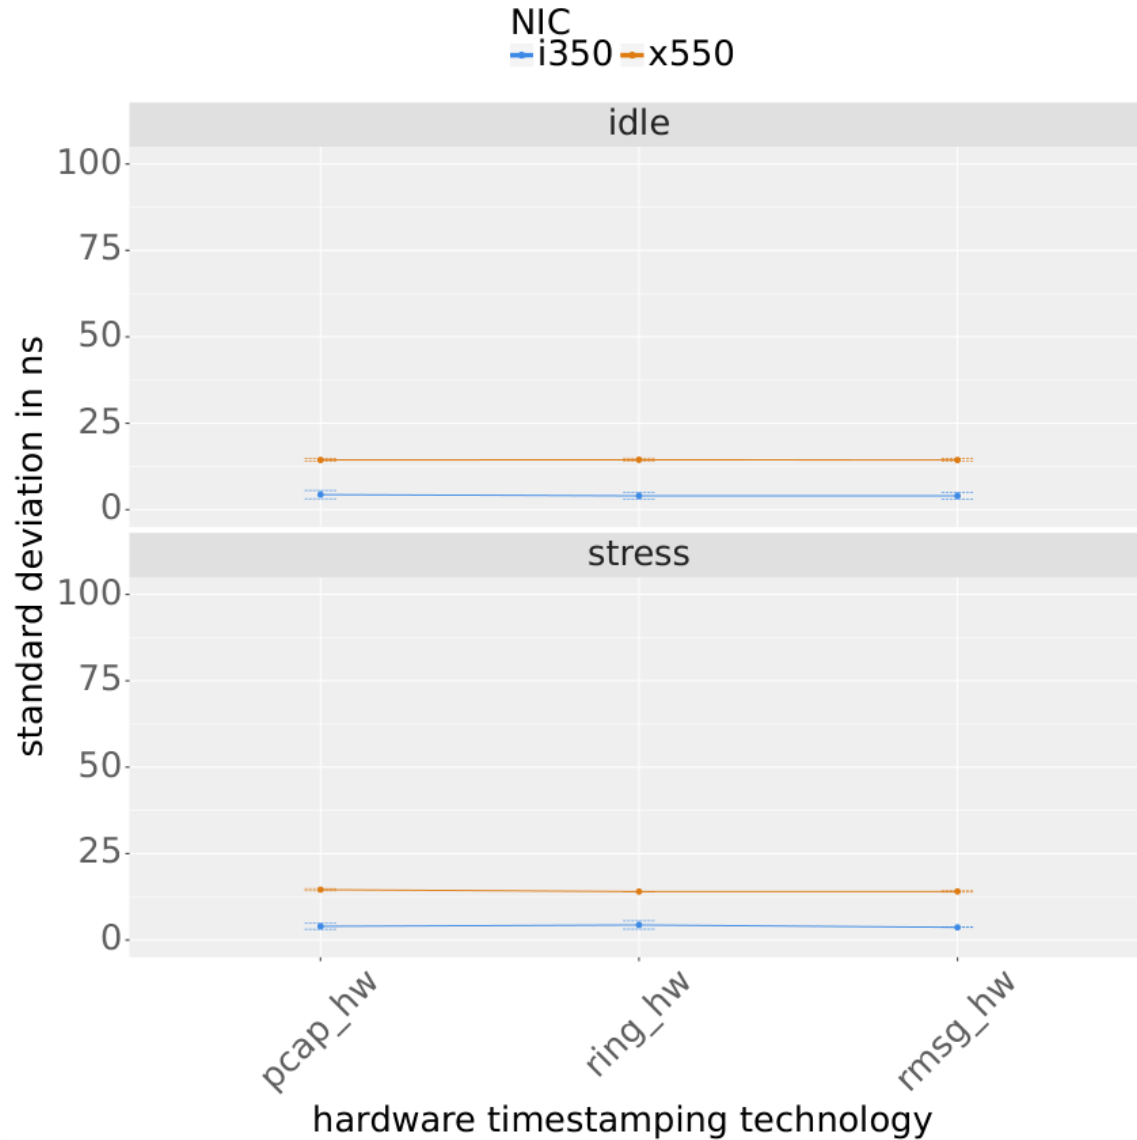
Artificial, Stable Source of Delay: Varying Cable Lengths



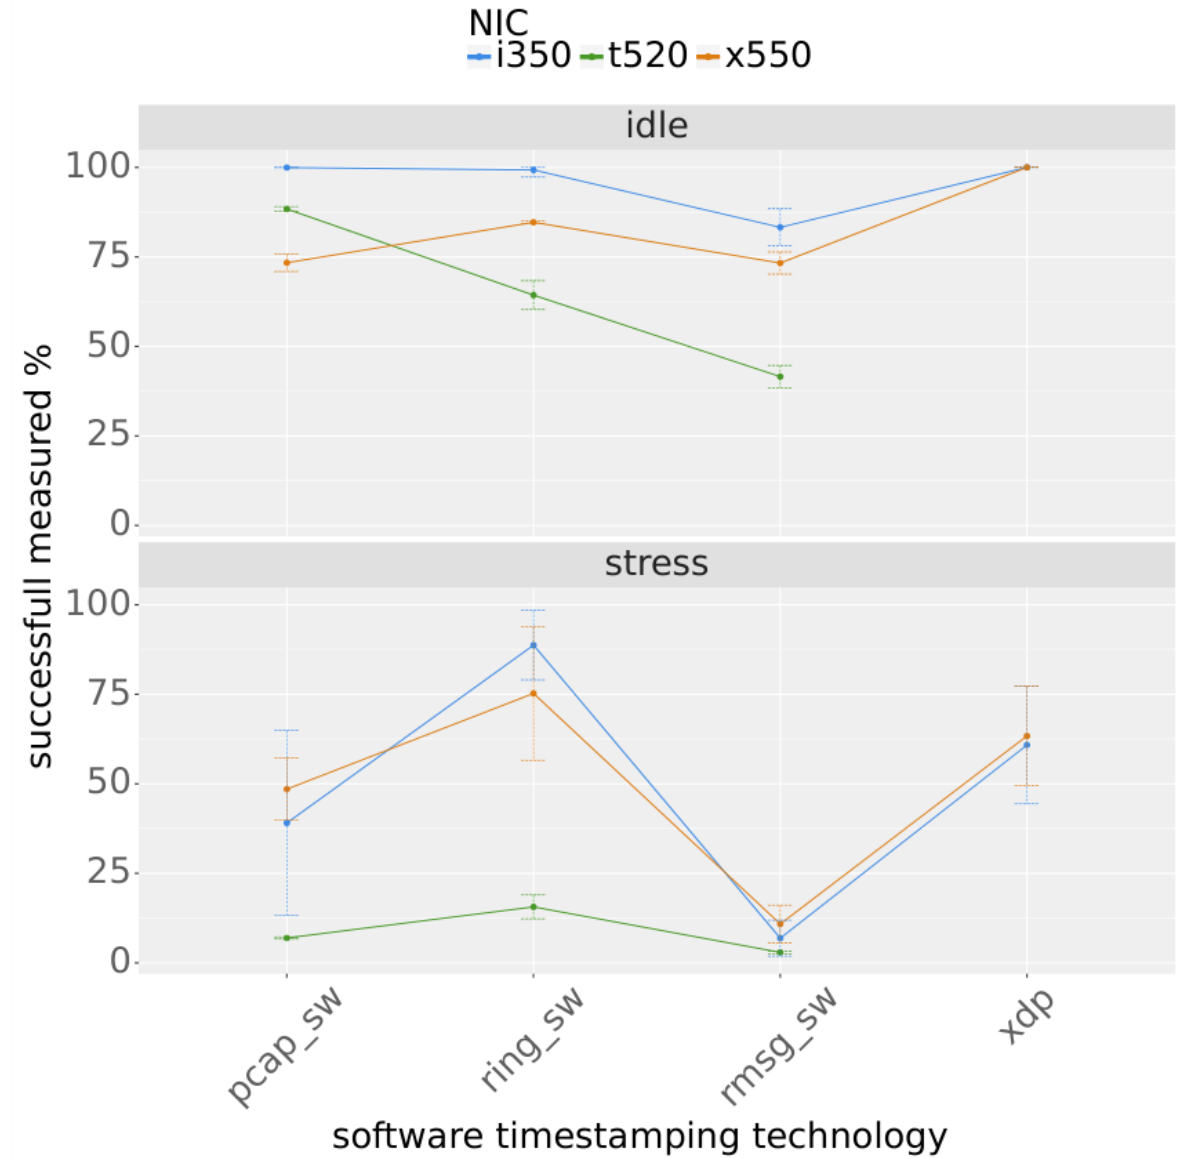
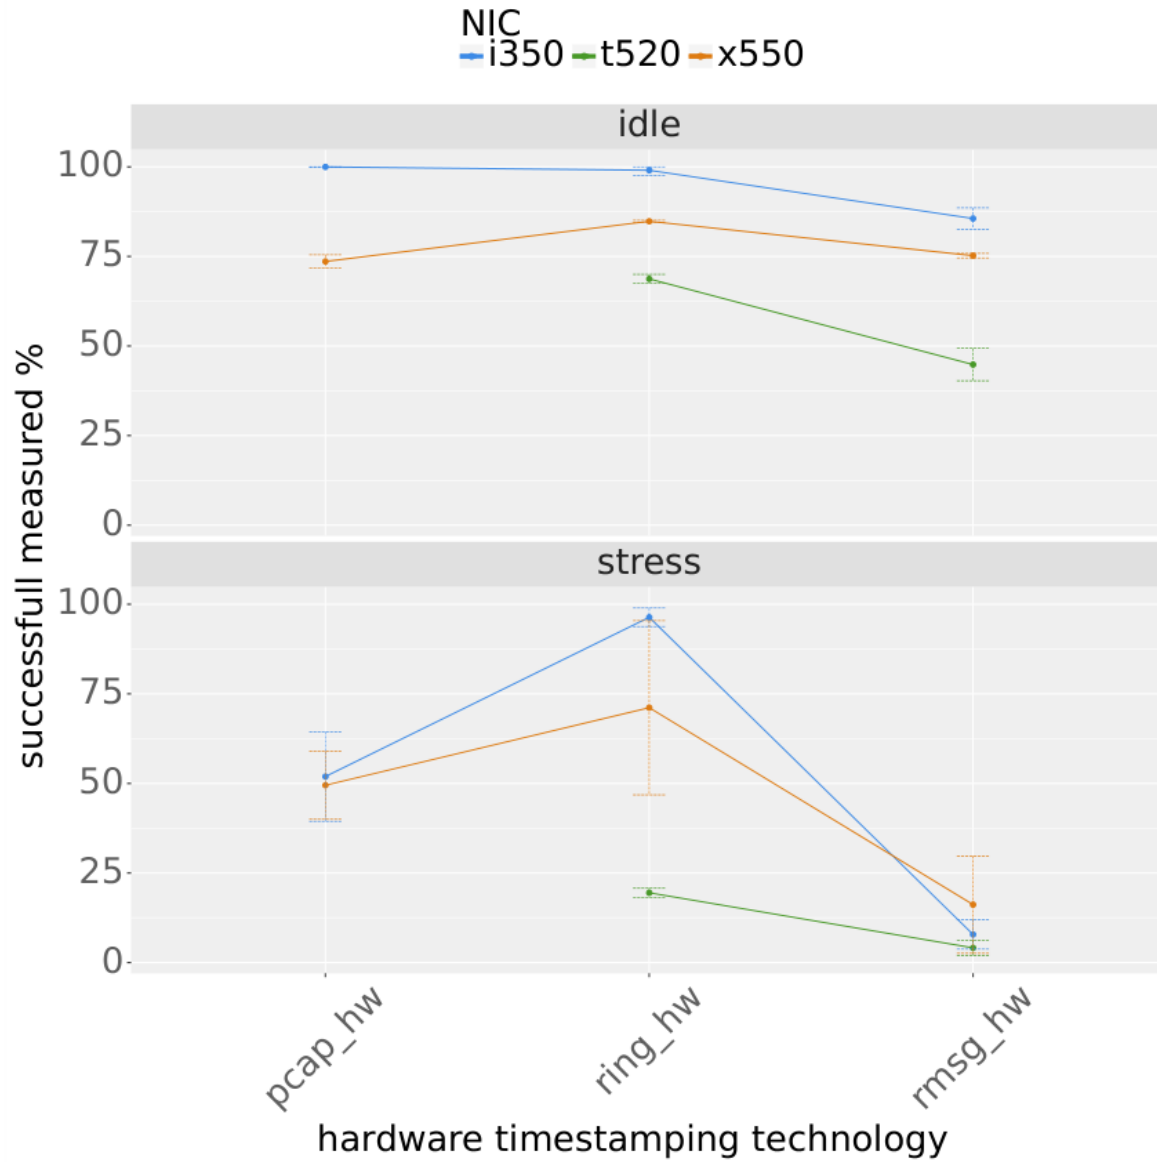
Main Effects



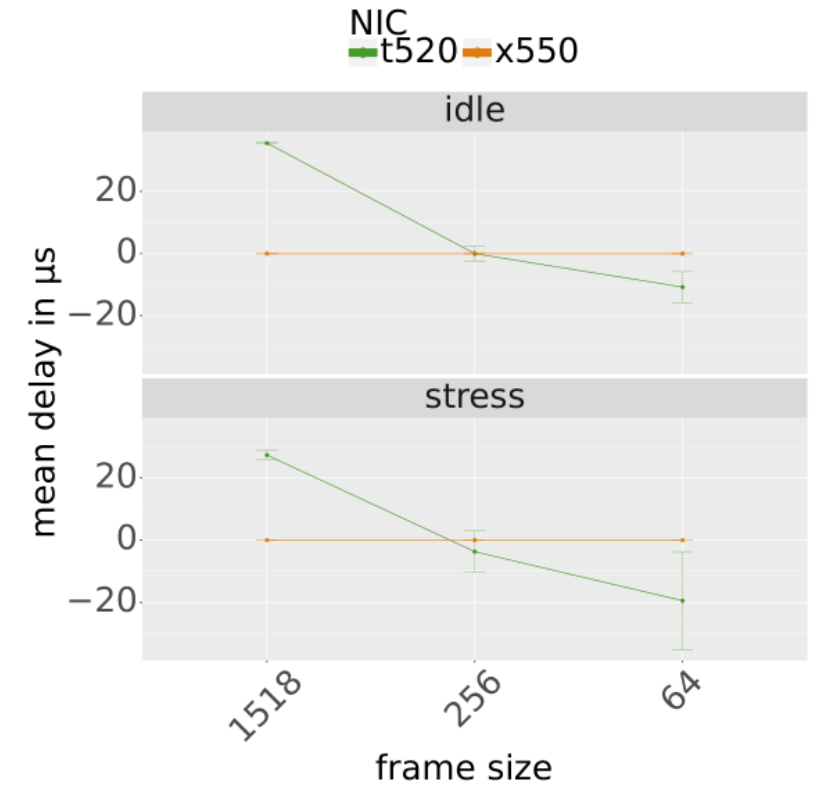

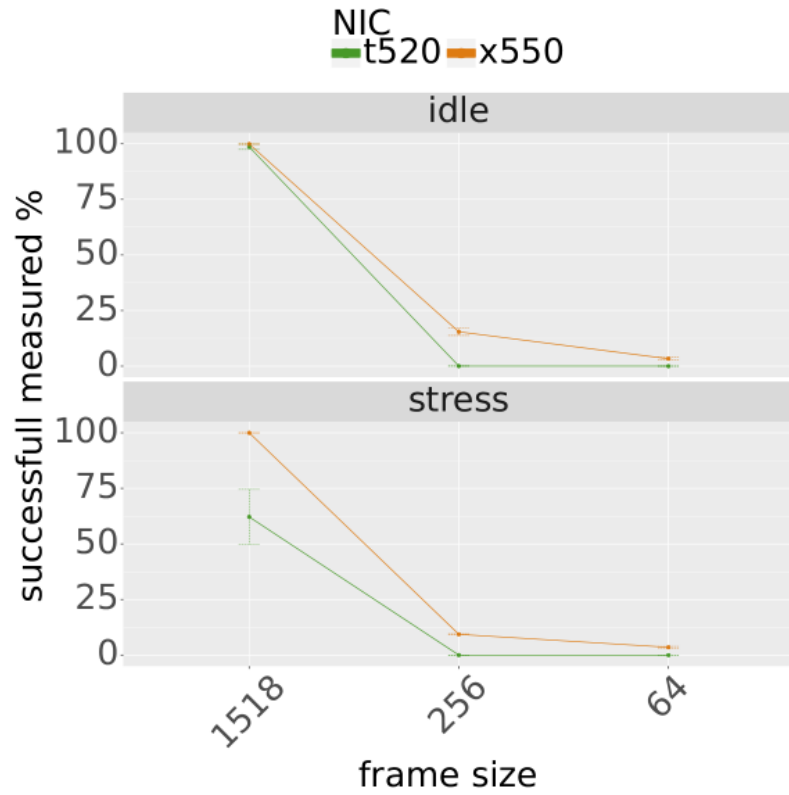
Standard Deviations



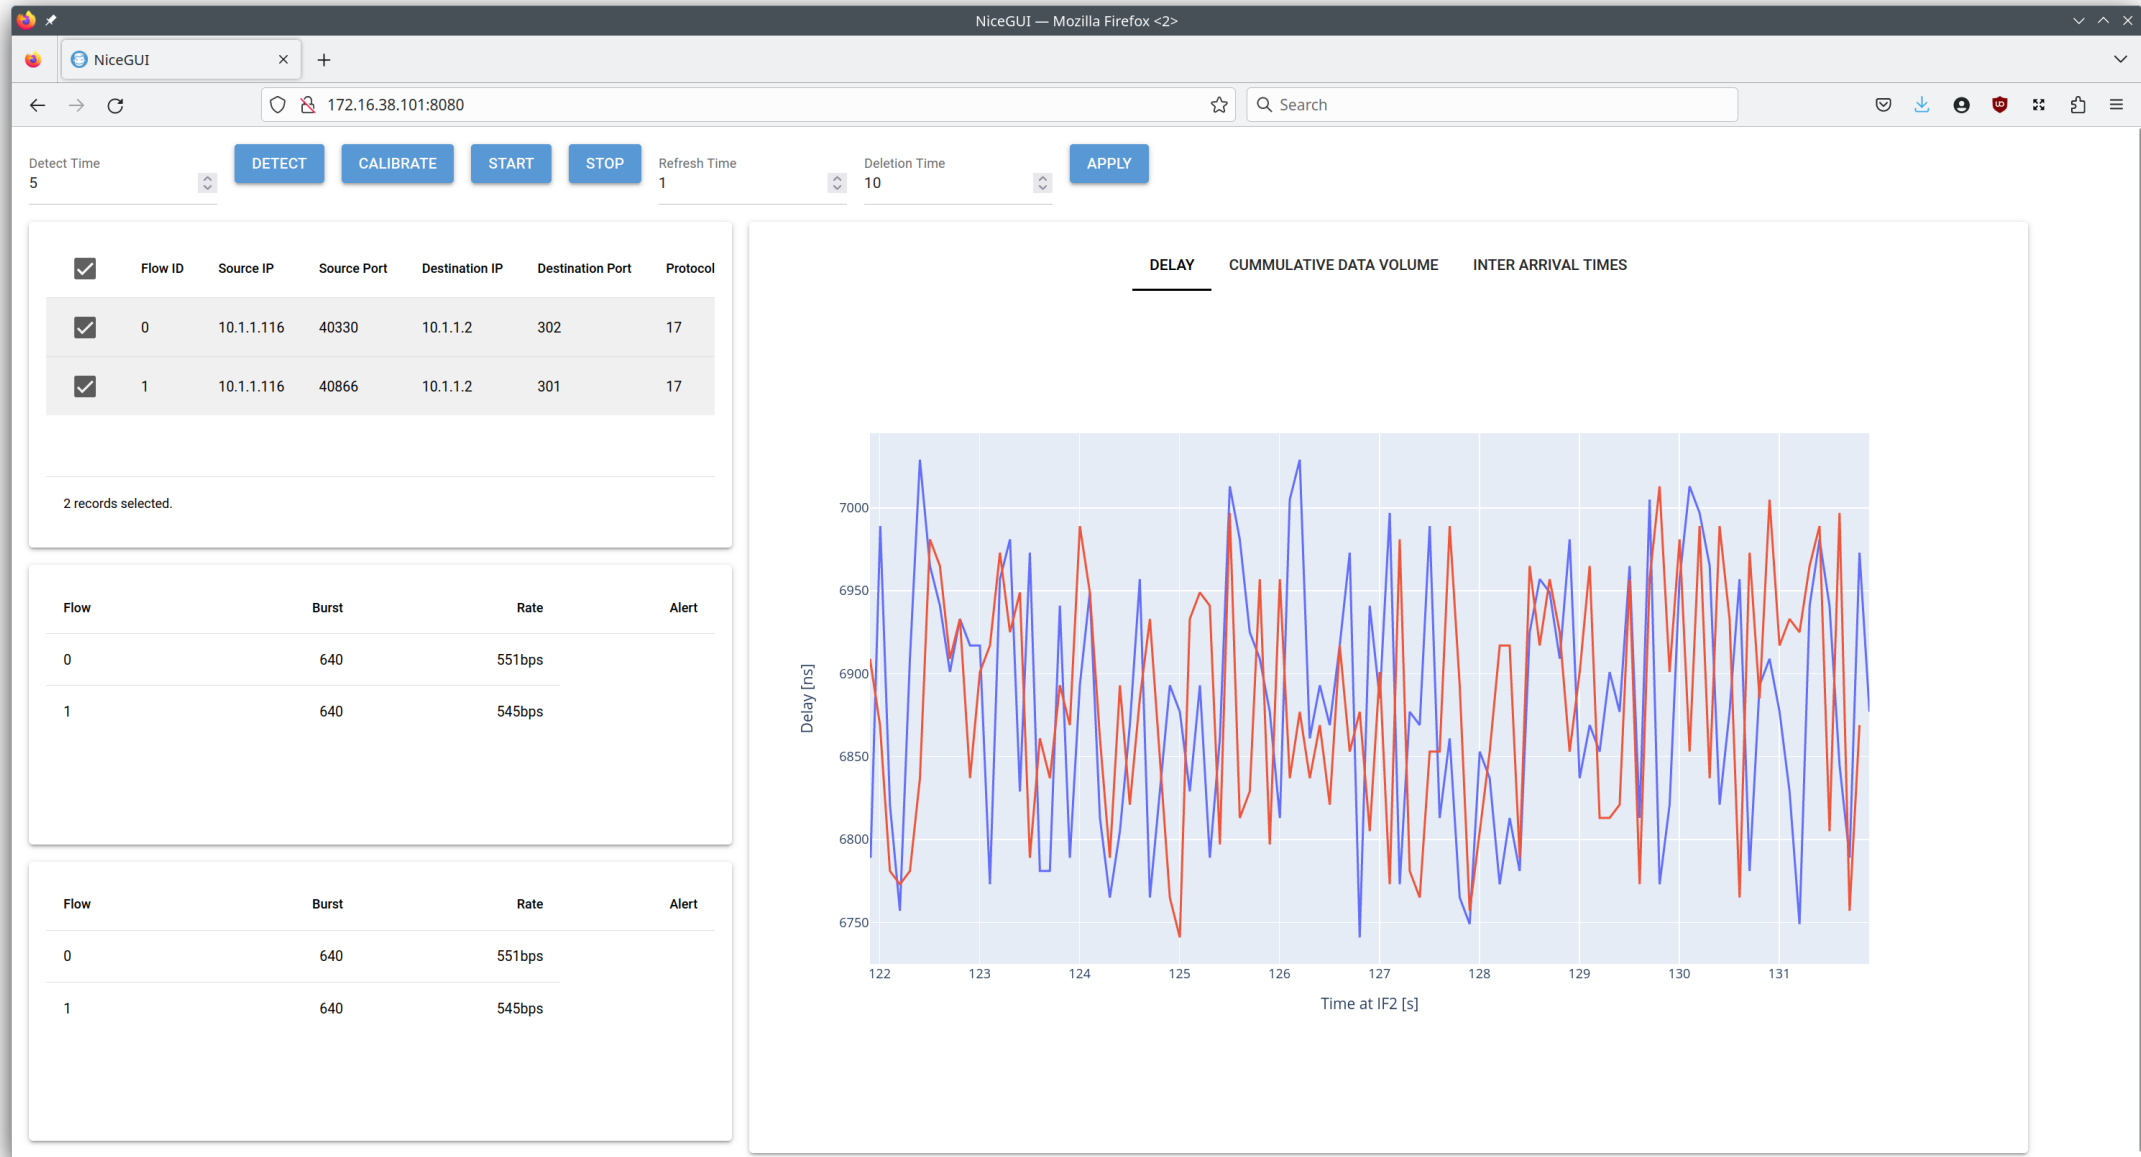
Throughput Comparison



What's Missing? → 10G Performance




What's Coming Next? → Demo Testbed




Conclusion

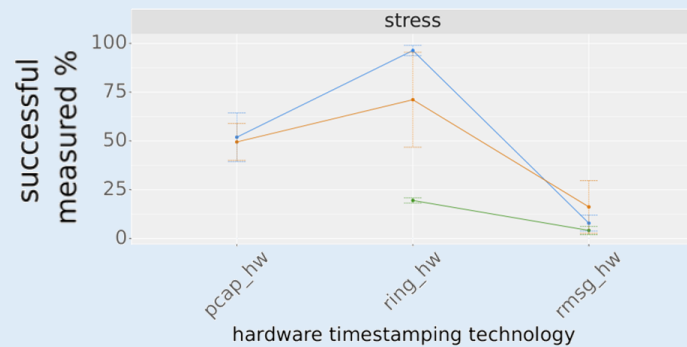
Crazy High Accuracy



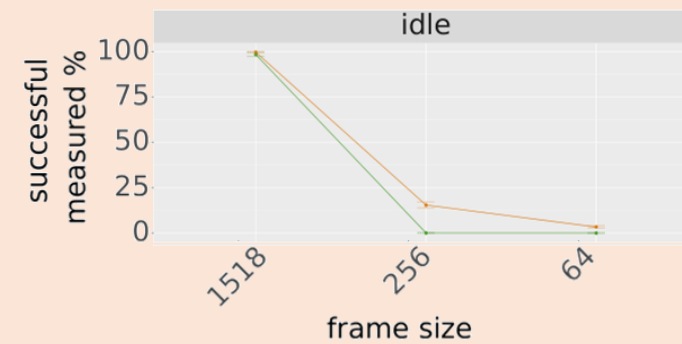
Cheap



Avoid Many SysCalls



10G Needs Work



THANK YOU!

Questions, comments, suggestions?



Alexej Grigorjew

University of Wuerzburg

Chair of Communication Networks

Email: alexej.grigorjew@uni-wuerzburg.de