Network Measurements for Data Centres

Diana Andreea Popescu
Supervisor: Andrew Moore

Why is network latency important?

How to provide network guarantees to applications?

Is bandwidth guarantee enough for an application performance?

What happens with performance if the applications' requests are delayed?

1. Understand impact on application performance
2. Measurement instrumentation to determine network latency dynamically
3. Scheduling mechanism for flows/apps across the data centre

Why is network latency important?

Experiments to show how application performance is influenced when adding delay using NetEm

Apps: Memcached, MySQL database, Apache web server

Why is network latency important?

Apache web server with Apache Benchmark

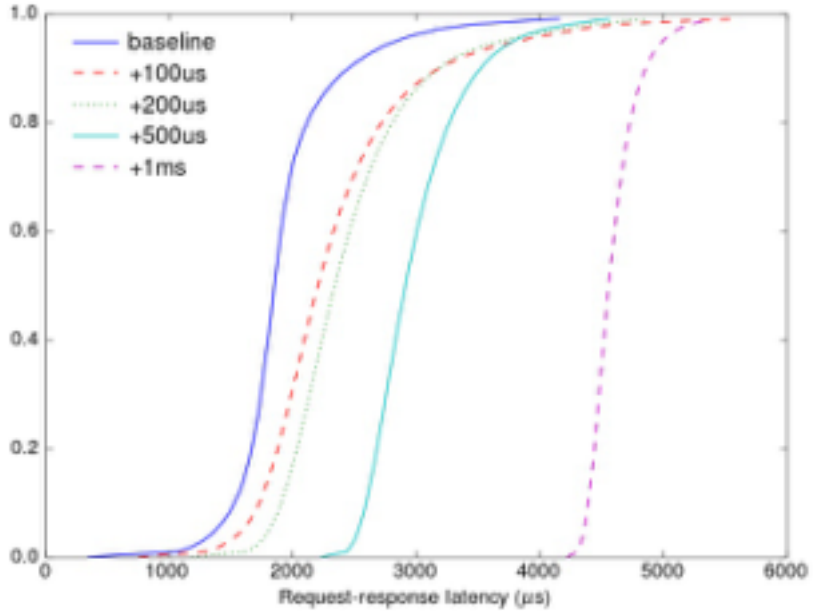

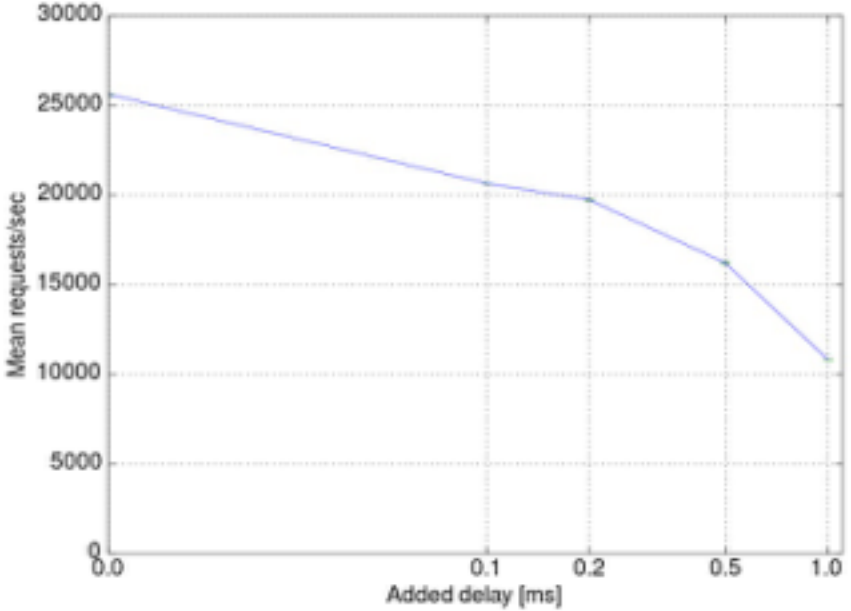

Why is network latency important?

MySQL database server with TPC-C benchmark

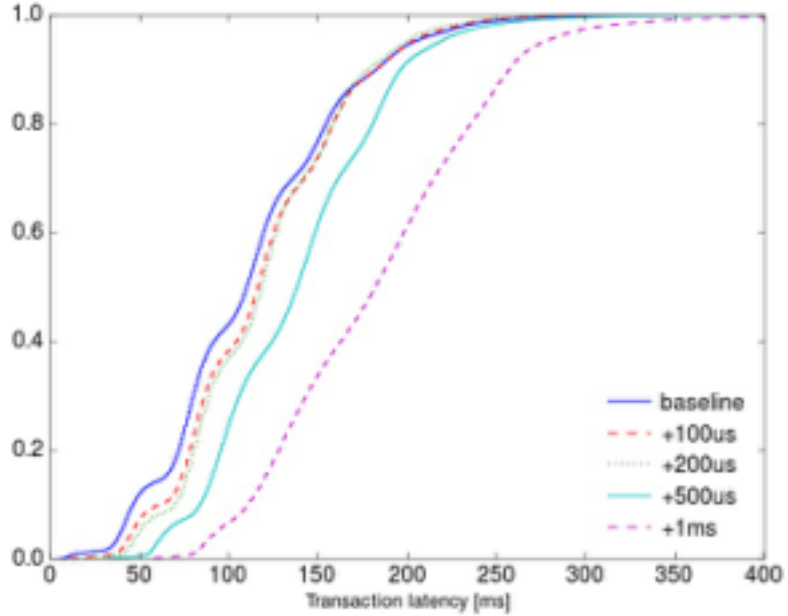

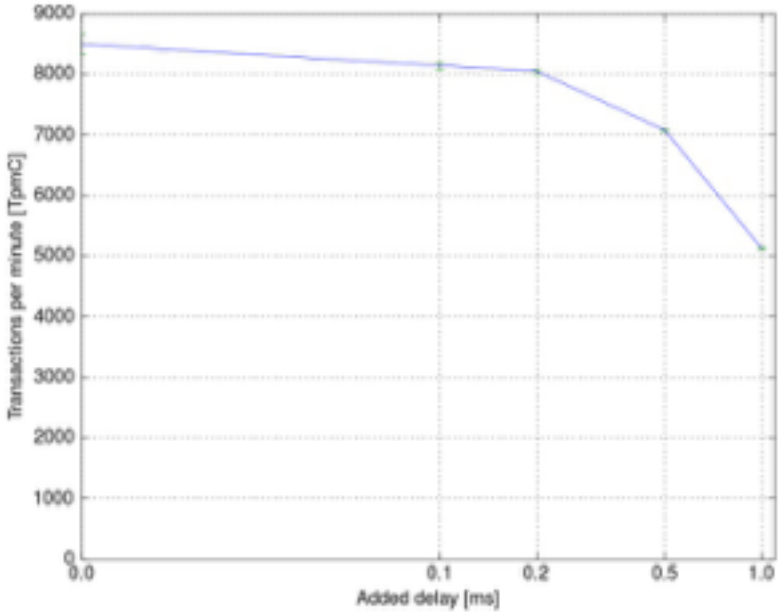

Why is network latency important?

Memcached server with memaslap client

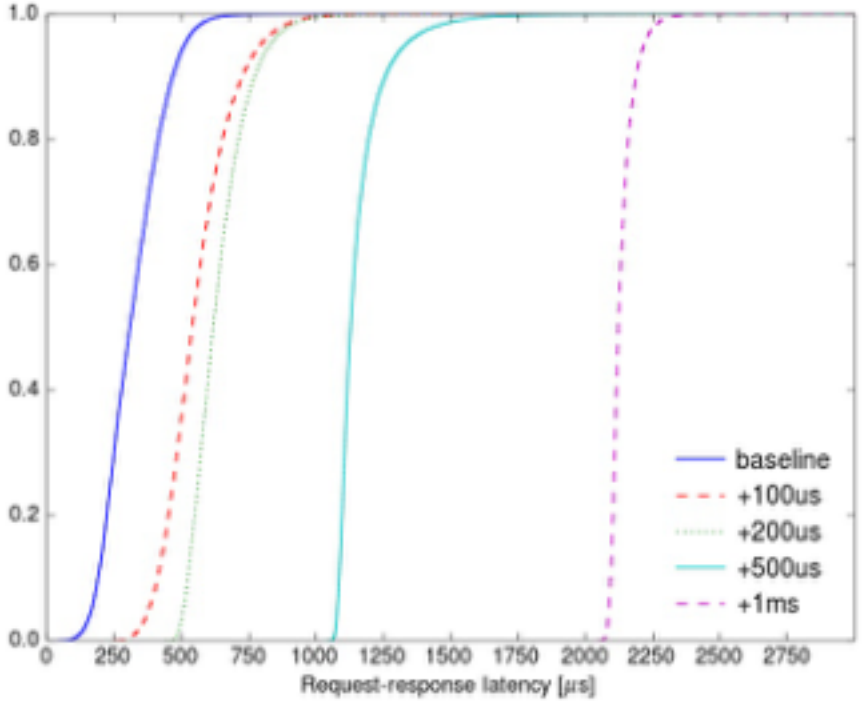

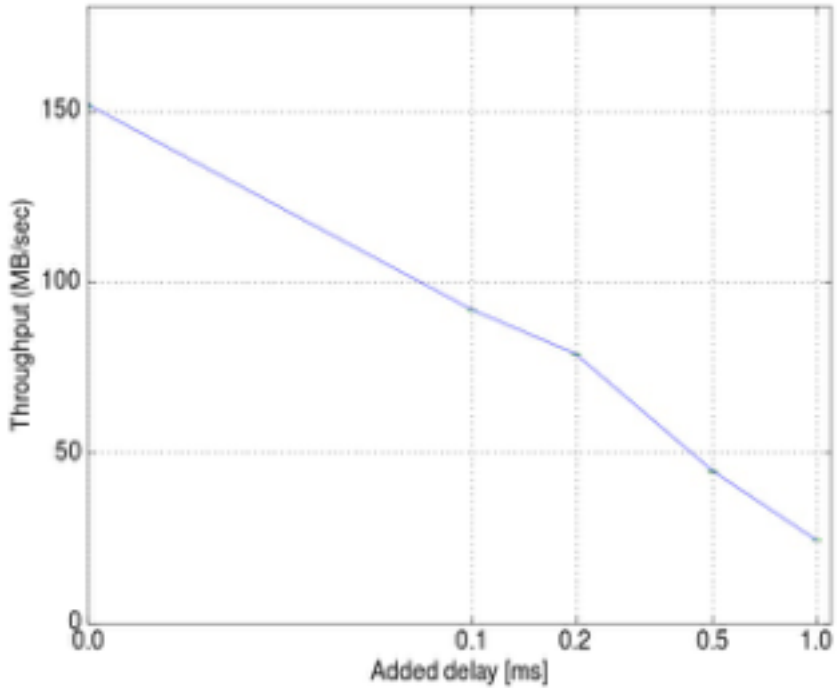

Omniscient

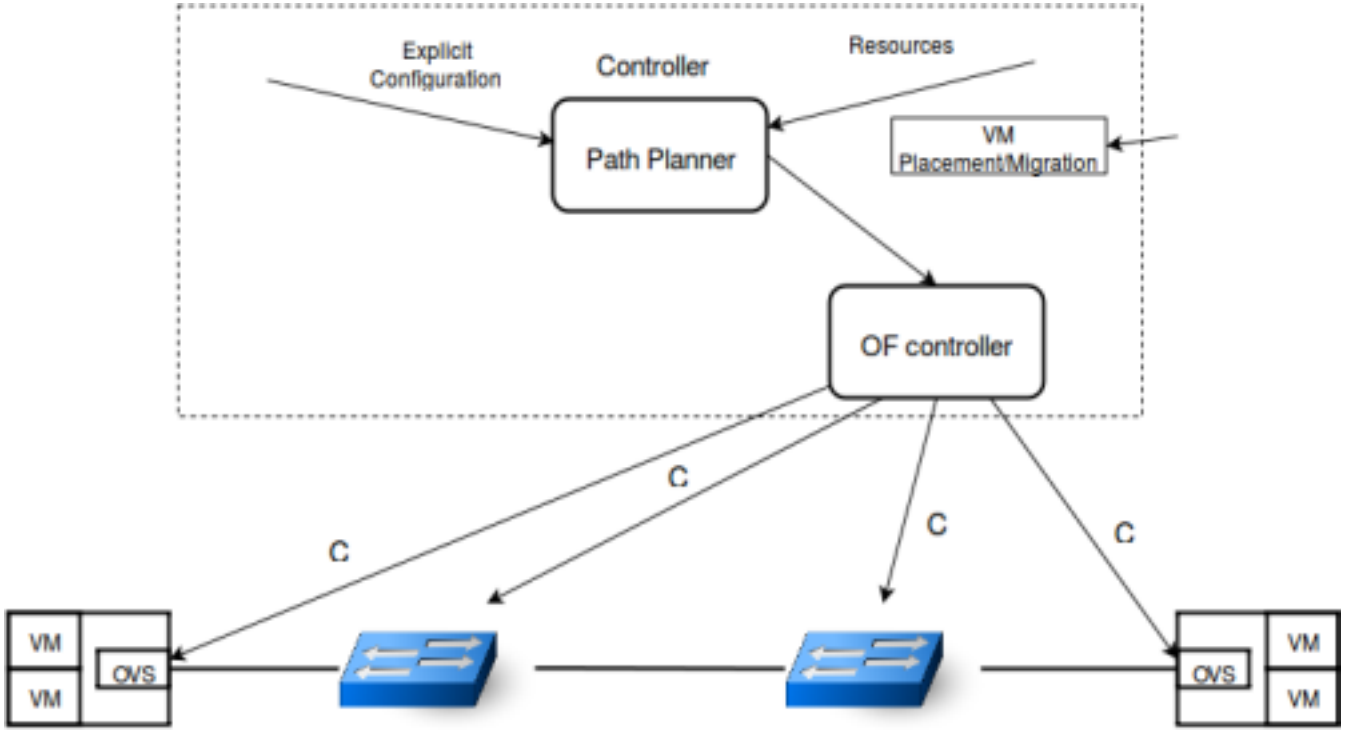
Explicit
Configuration
Controller
Resources
Path Planner
VM
Placement/Migration
OF controller
C
C
C
C
VM
VM
OVS
OVS
VM
VM

Omniscient

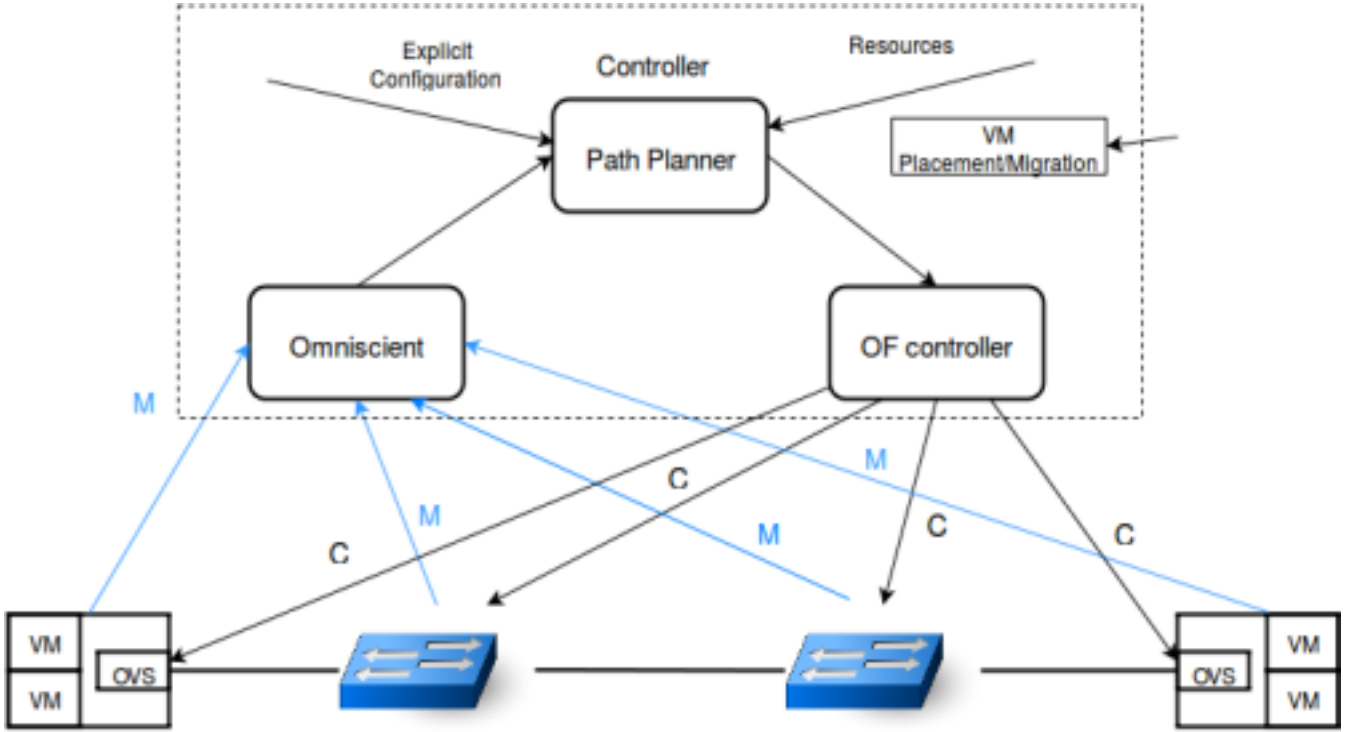
Explicit
Configuration
Controller
Resources
Path Planner
VM
Placement/Migration
Omniscient
OF controller
M
C
M
C
M
C
M
C
VM
VM
OVS
OVS
VM
VM

How can we measure network latency in a data centre?

- Each host has GPS receiver
  - Not scalable
- Ping
  - Overhead is high, accuracy is not good
- DTP (Datacentre Time Protocol) [SIGCOMM'16]
  - Hardware-based solution, not immediately deployable
- PTP (Precision Time Protocol)
  - Moderate overhead, sub-us precision

How can we measure network latency in a data centre?

- Each host has GPS receiver
  - Not scalable
- Ping
  - Overhead is high, accuracy is not good
- DTP (Datacentre Time Protocol) [SIGCOMM'16]
  - Hardware-based solution, not immediately deployable
- PTP (Precision Time Protocol)
  - Moderate overhead, sub-us precision

PTP (Precision Time Protocol)

- Protocol for synchronising clocks over the network
- T1' - T1 = offset + d and T2' - T2 = -offset + d, assumes d symmetric
- offset =(T1

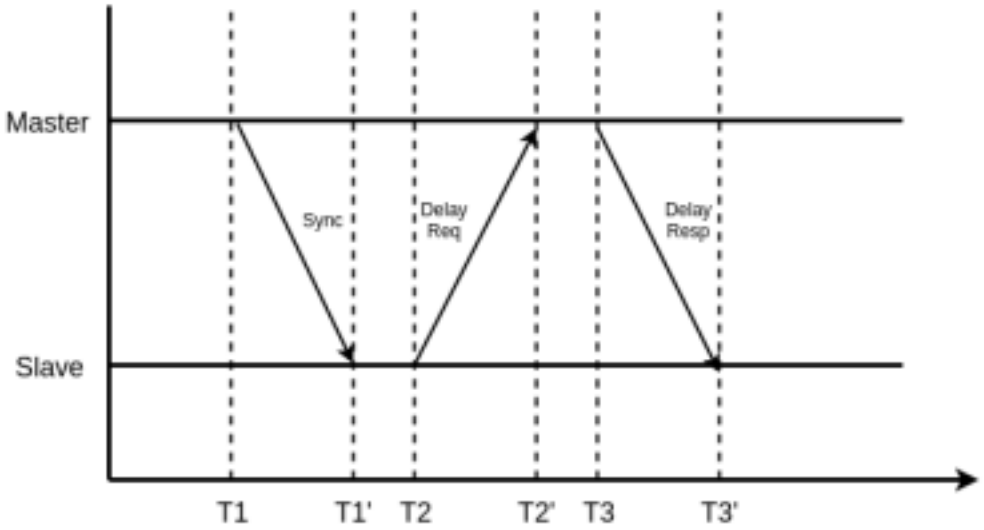

Experiments with iperf, ptp and app

- Run iperf continuously
- Run iperf alternating with periods when its not running - how are the apps affected by jitter

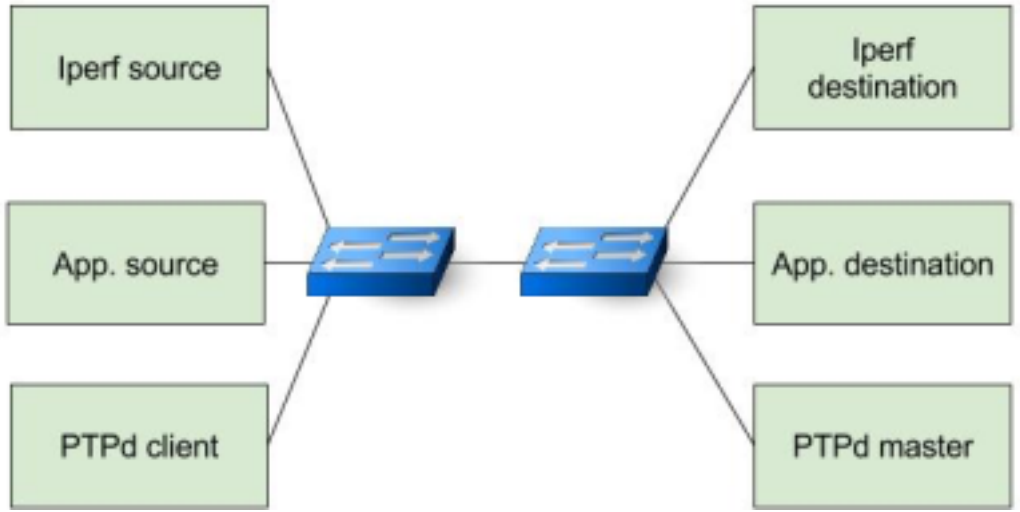

Apache benchmark + Iperf + PTPd

App

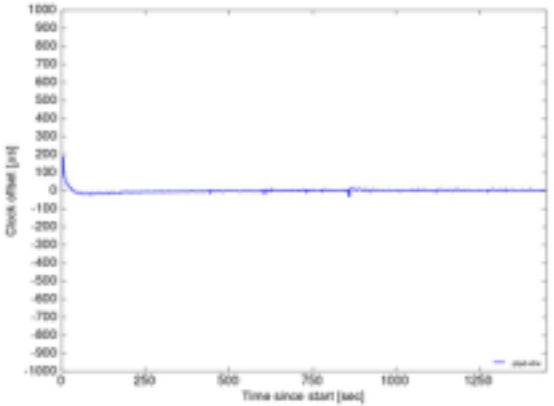

App + Continuous Iperf

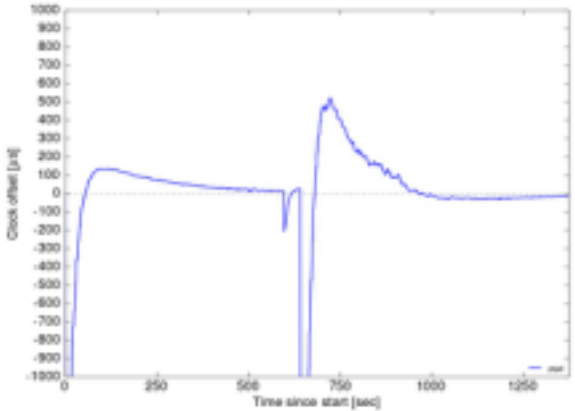

App + Periodic Iperf

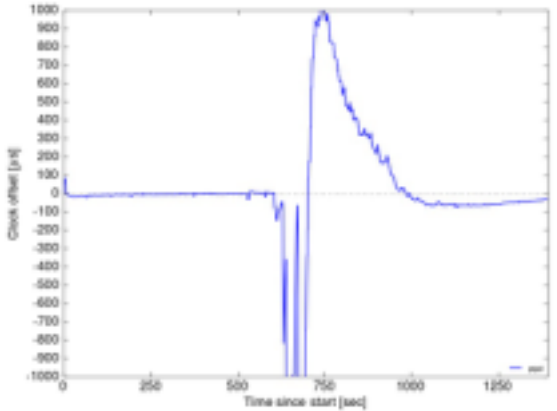

TPC-C benchmark + Iperf + PTPd

App

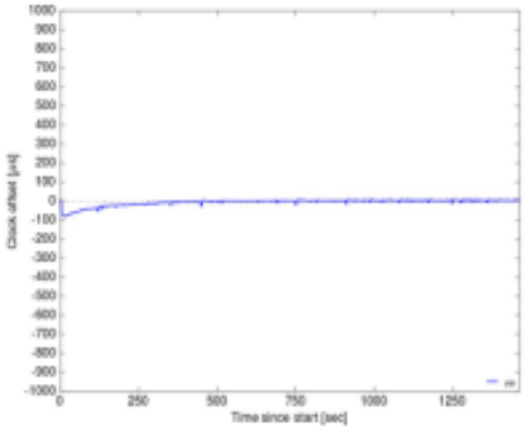

App + Continuous Iperf

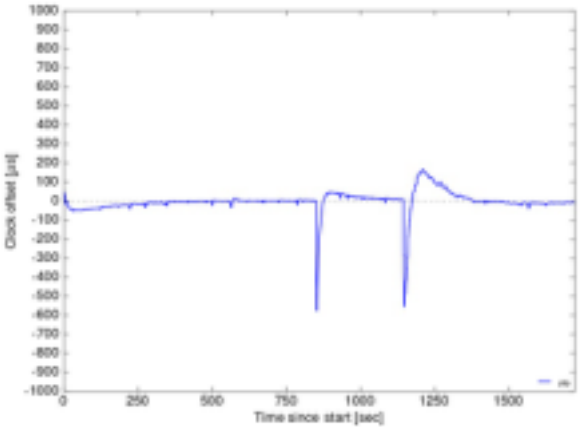

App + Periodic Iperf

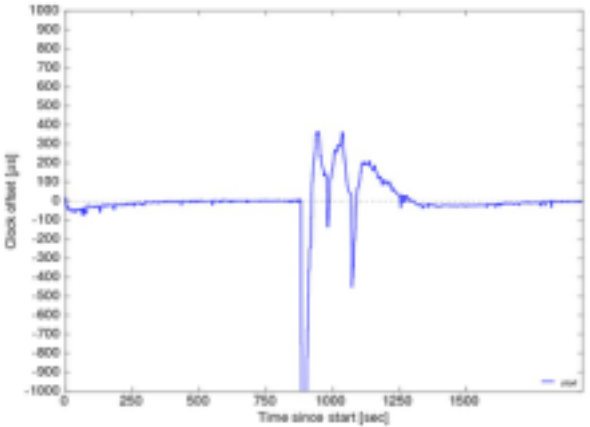

Memcached client + Iperf + PTPd

App

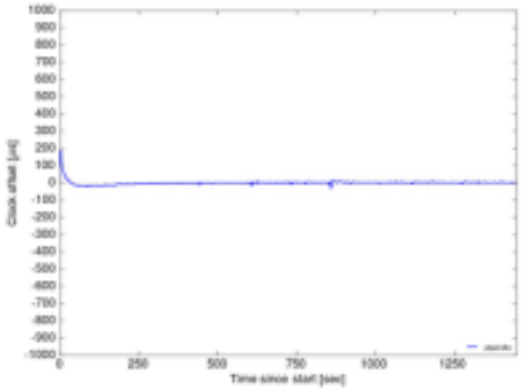

App + Continuous Iperf

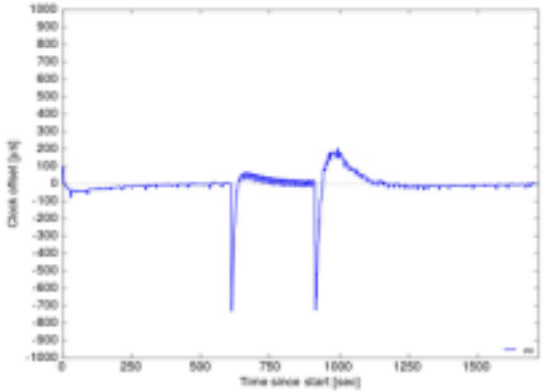

App + Periodic Iperf

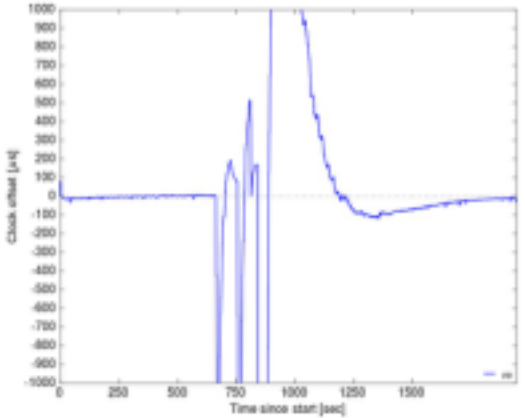

Clock offset vs. Network Latency

- Tap TX and RX to timestamp packets on each direction
- Timestamp packets using ExaNIC with 6.2 ns timestamp resolution
- Correlate timestamp at the entry and exit of switch on each direction
- Correlate delays w mps

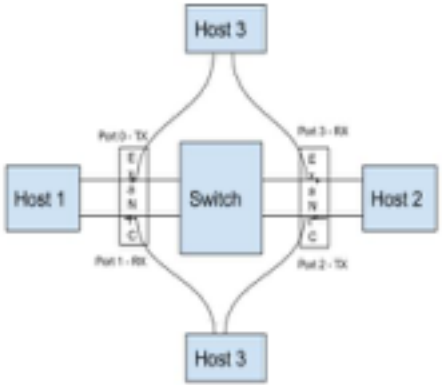

Preliminary results

- PTP is affected by the traffic at the end host

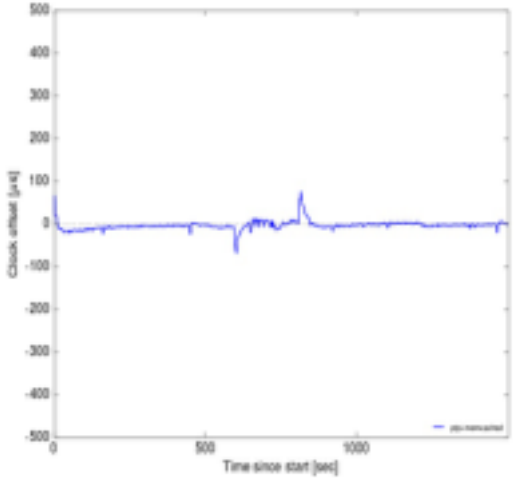

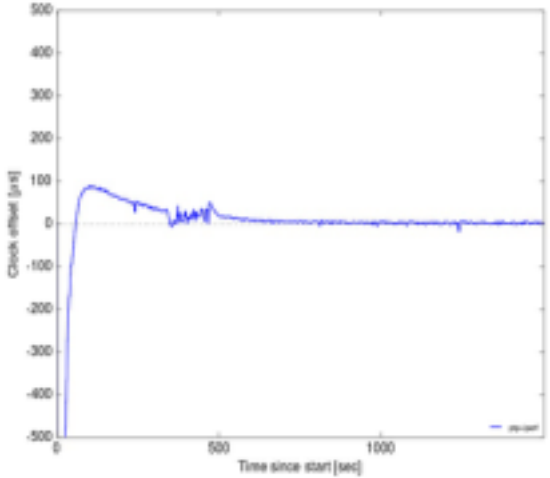

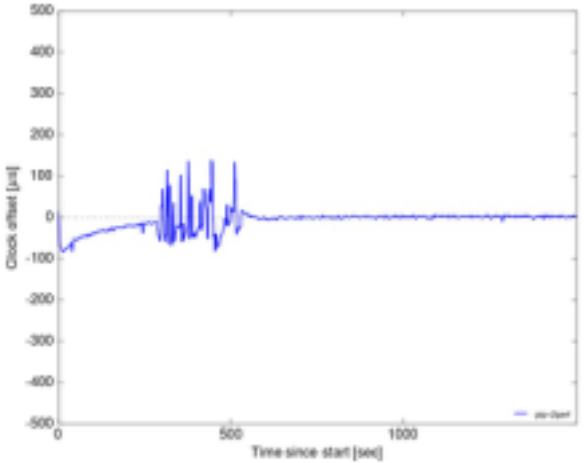

- Should eliminate this interference or account for it in our analysis of the overall path delay

Next steps

- ExaNIC to be used for the setup with two switches in order to avoid end-host interference from results as a first step
- Quantify end-host effect
- Comparison with ping in terms of load and accuracy
- Implement the rest of the controller

Conclusion

- Application performance is affected by increased network latency
- We study PTP as a mechanism to determine dynamically the latency
- We want to provide apps network latency guarantees