Using hardware timestamping to reduce jitter in latency measurement

Christopher Walker, Raul Lopes, Duncan Rand, Tim Chown

Christopher.Walker@jisc.ac.uk

perfSONAR network monitoring

Matrix of tests

- Bandwidth measurements
- One way latency
- Packet loss
- GridPP mesh
  - https://psmad.opensciencegrid.org/maddash-webui/index.cgi?dashboard=UK%20Mesh%20Config

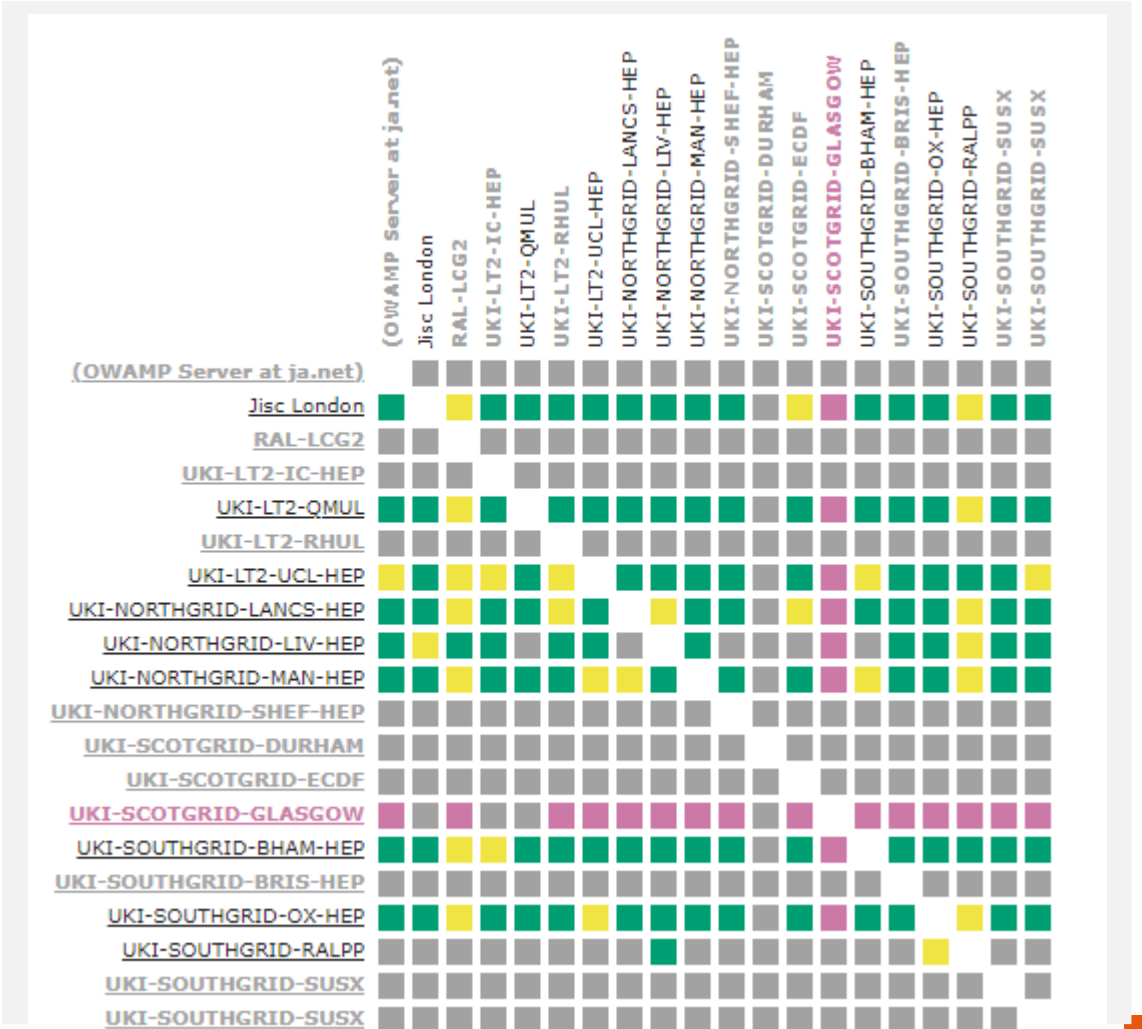

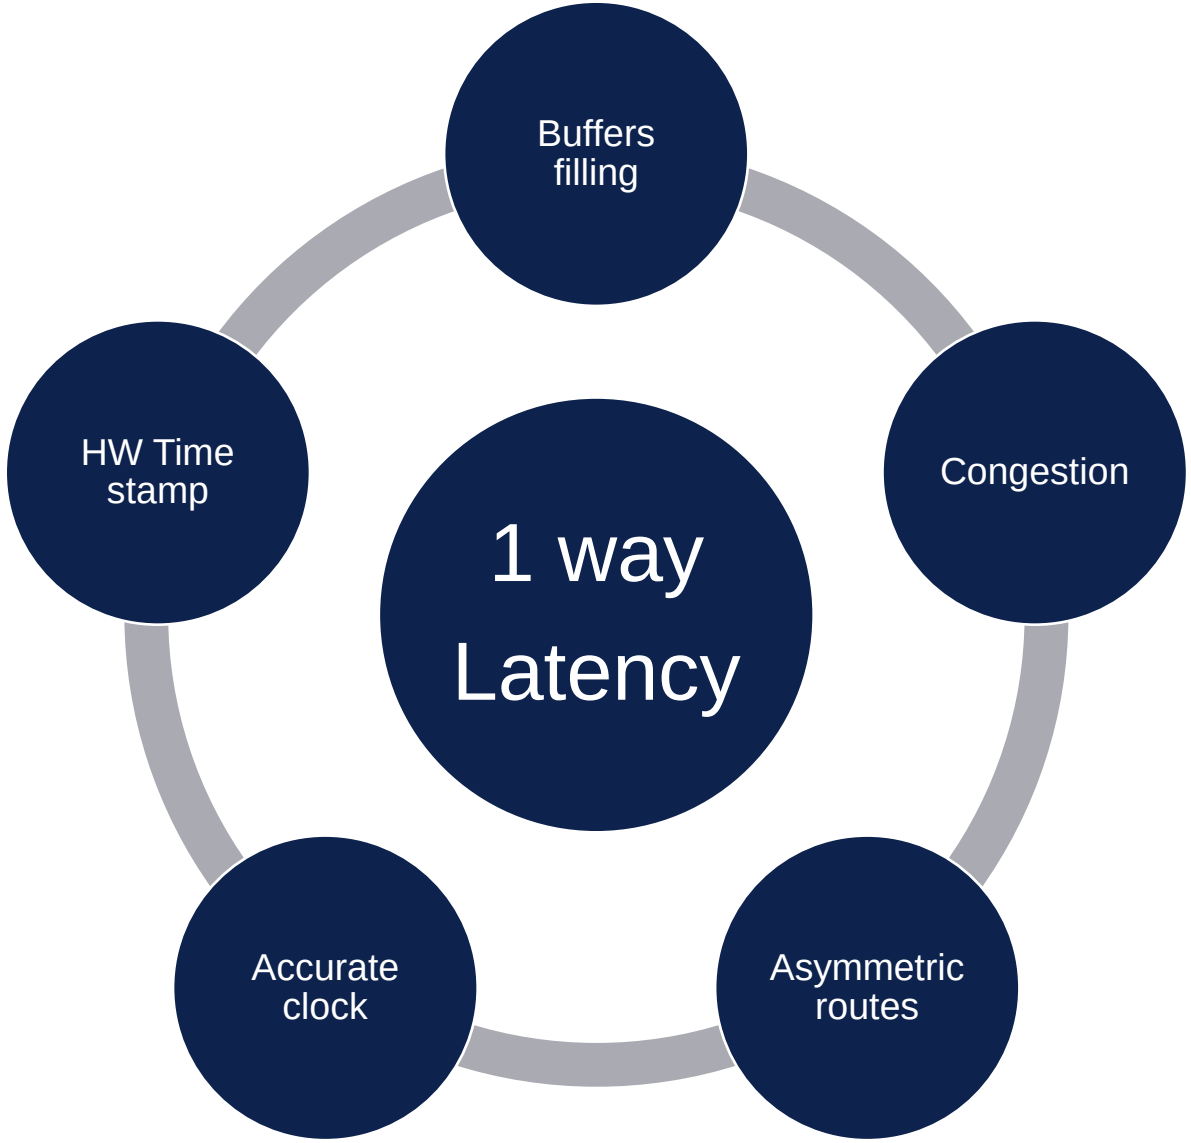
Buffers filling
Congestion
Asymmetric routes
Accurate clock
HW Time stamp
1 way Latency

Perfsonar one way latency measurements

Jisc London to/from Slough (1 month)

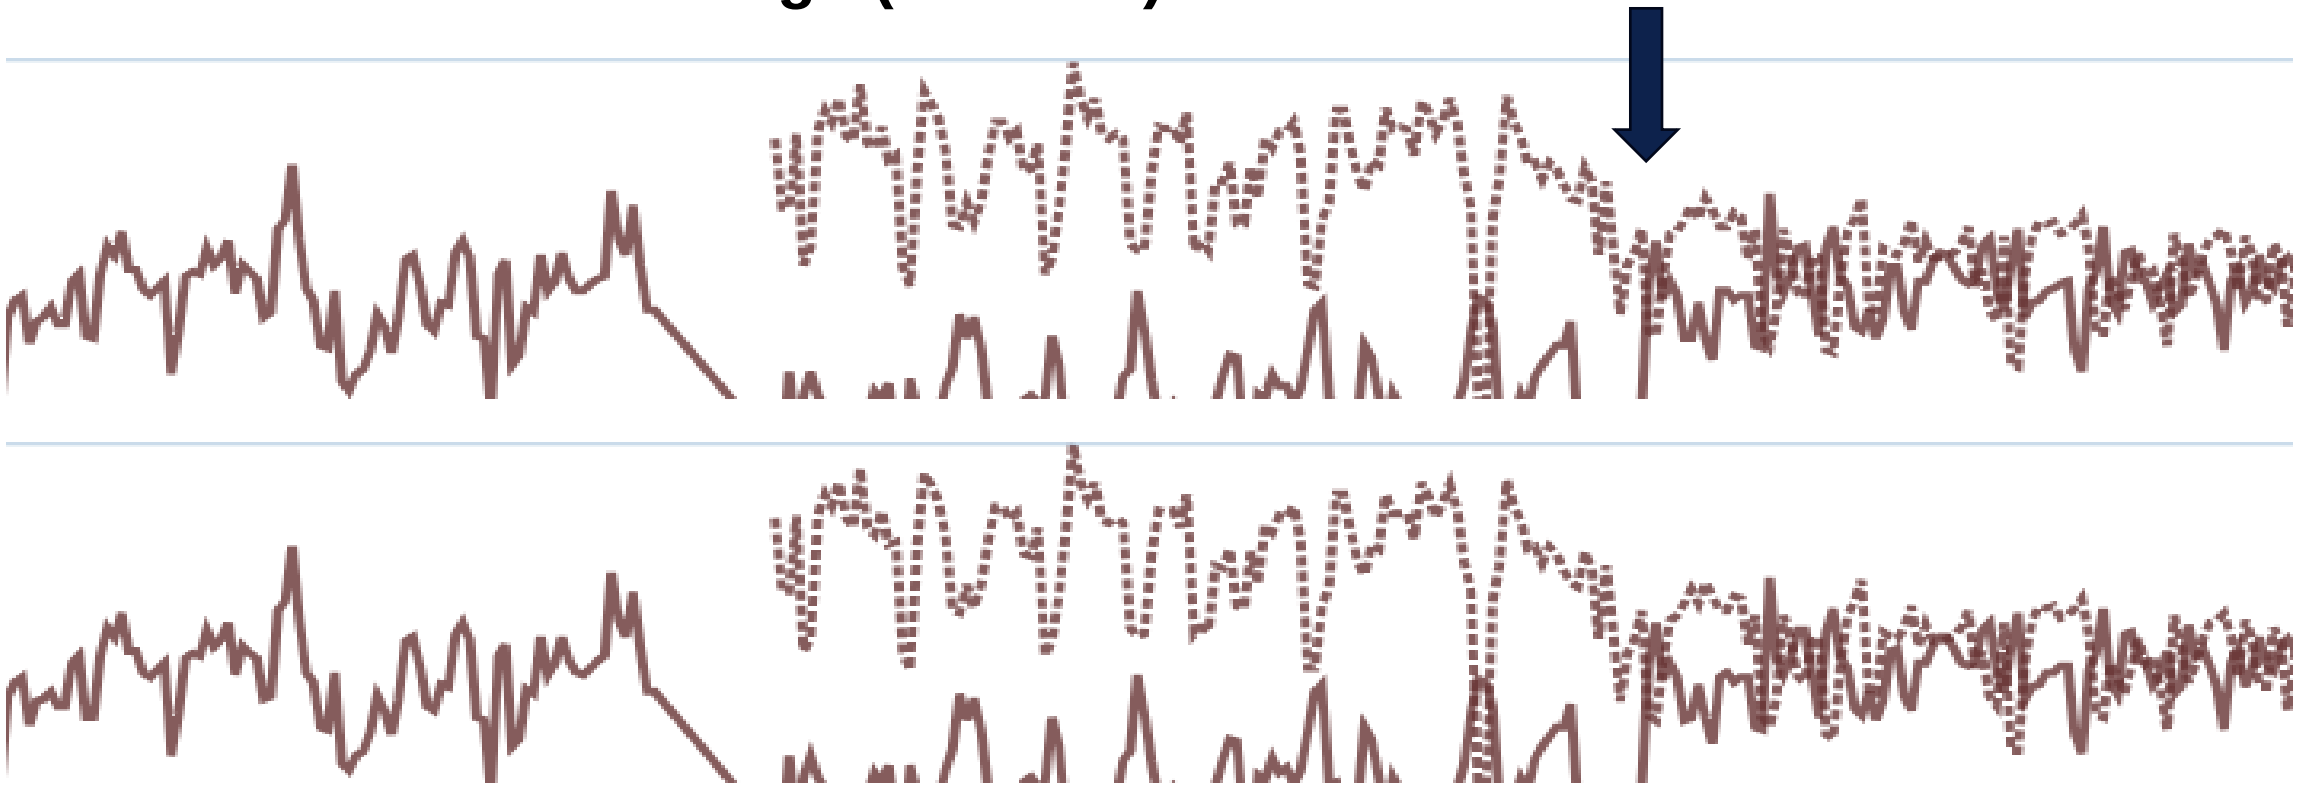

London -> Slough ~0.6 ms

Perfsonar one way latency measurements
Jisc (London) to QMUL

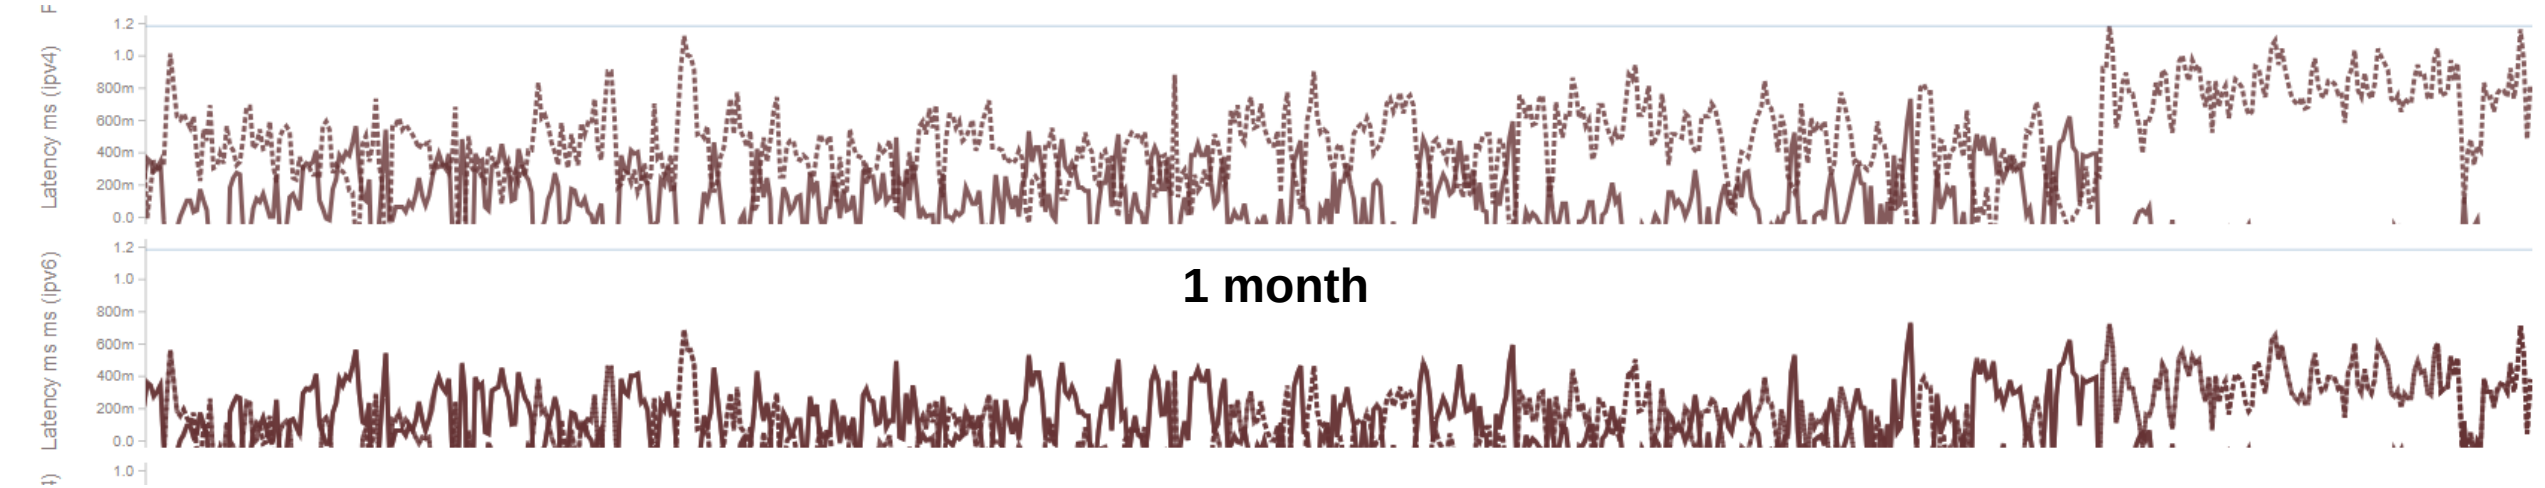

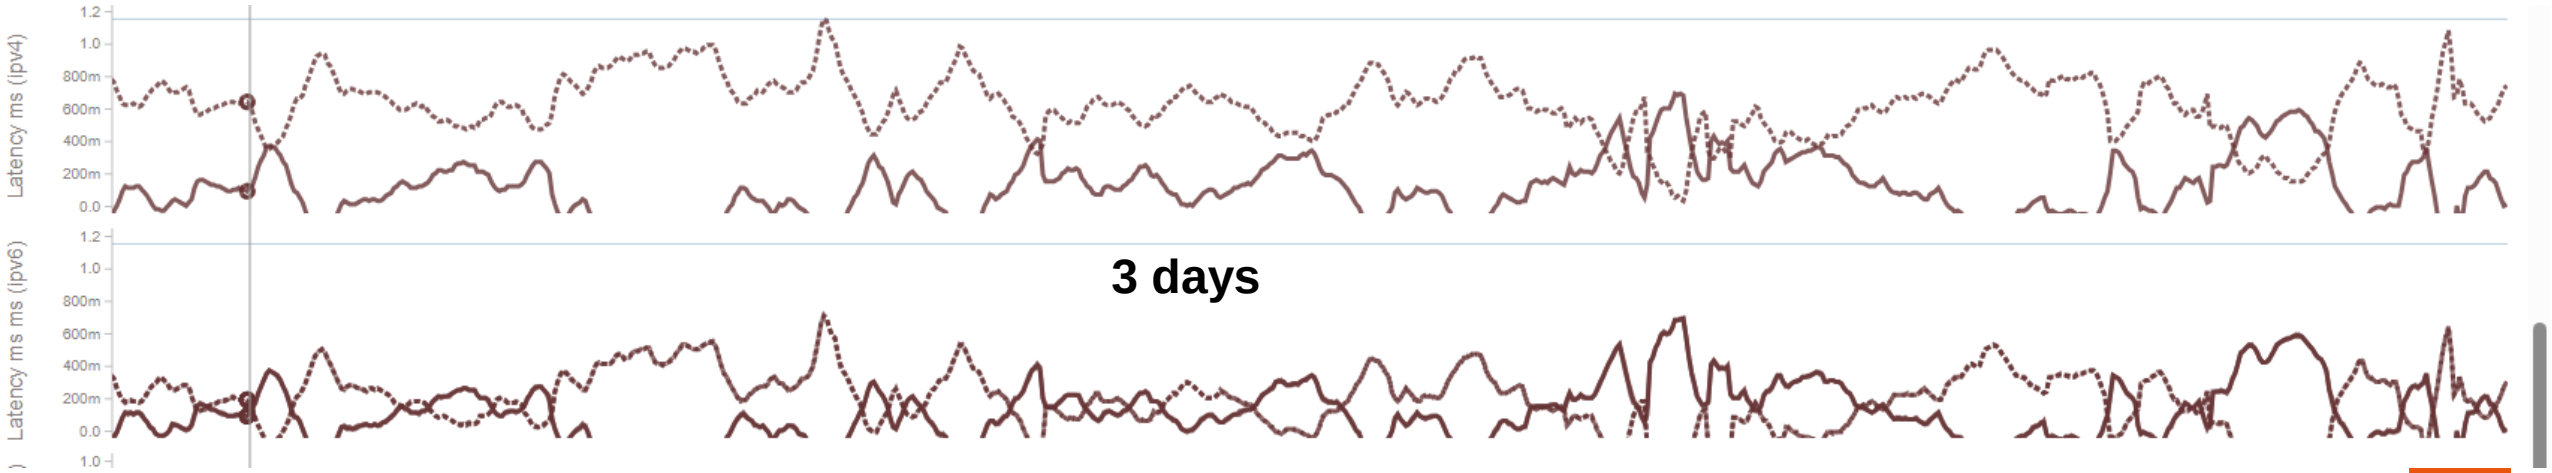

Effect of Clock offset

Latency

= T(Host B) – T (Host A)

Measured latency

= T(Host B) – T(Host A) + Clock offset

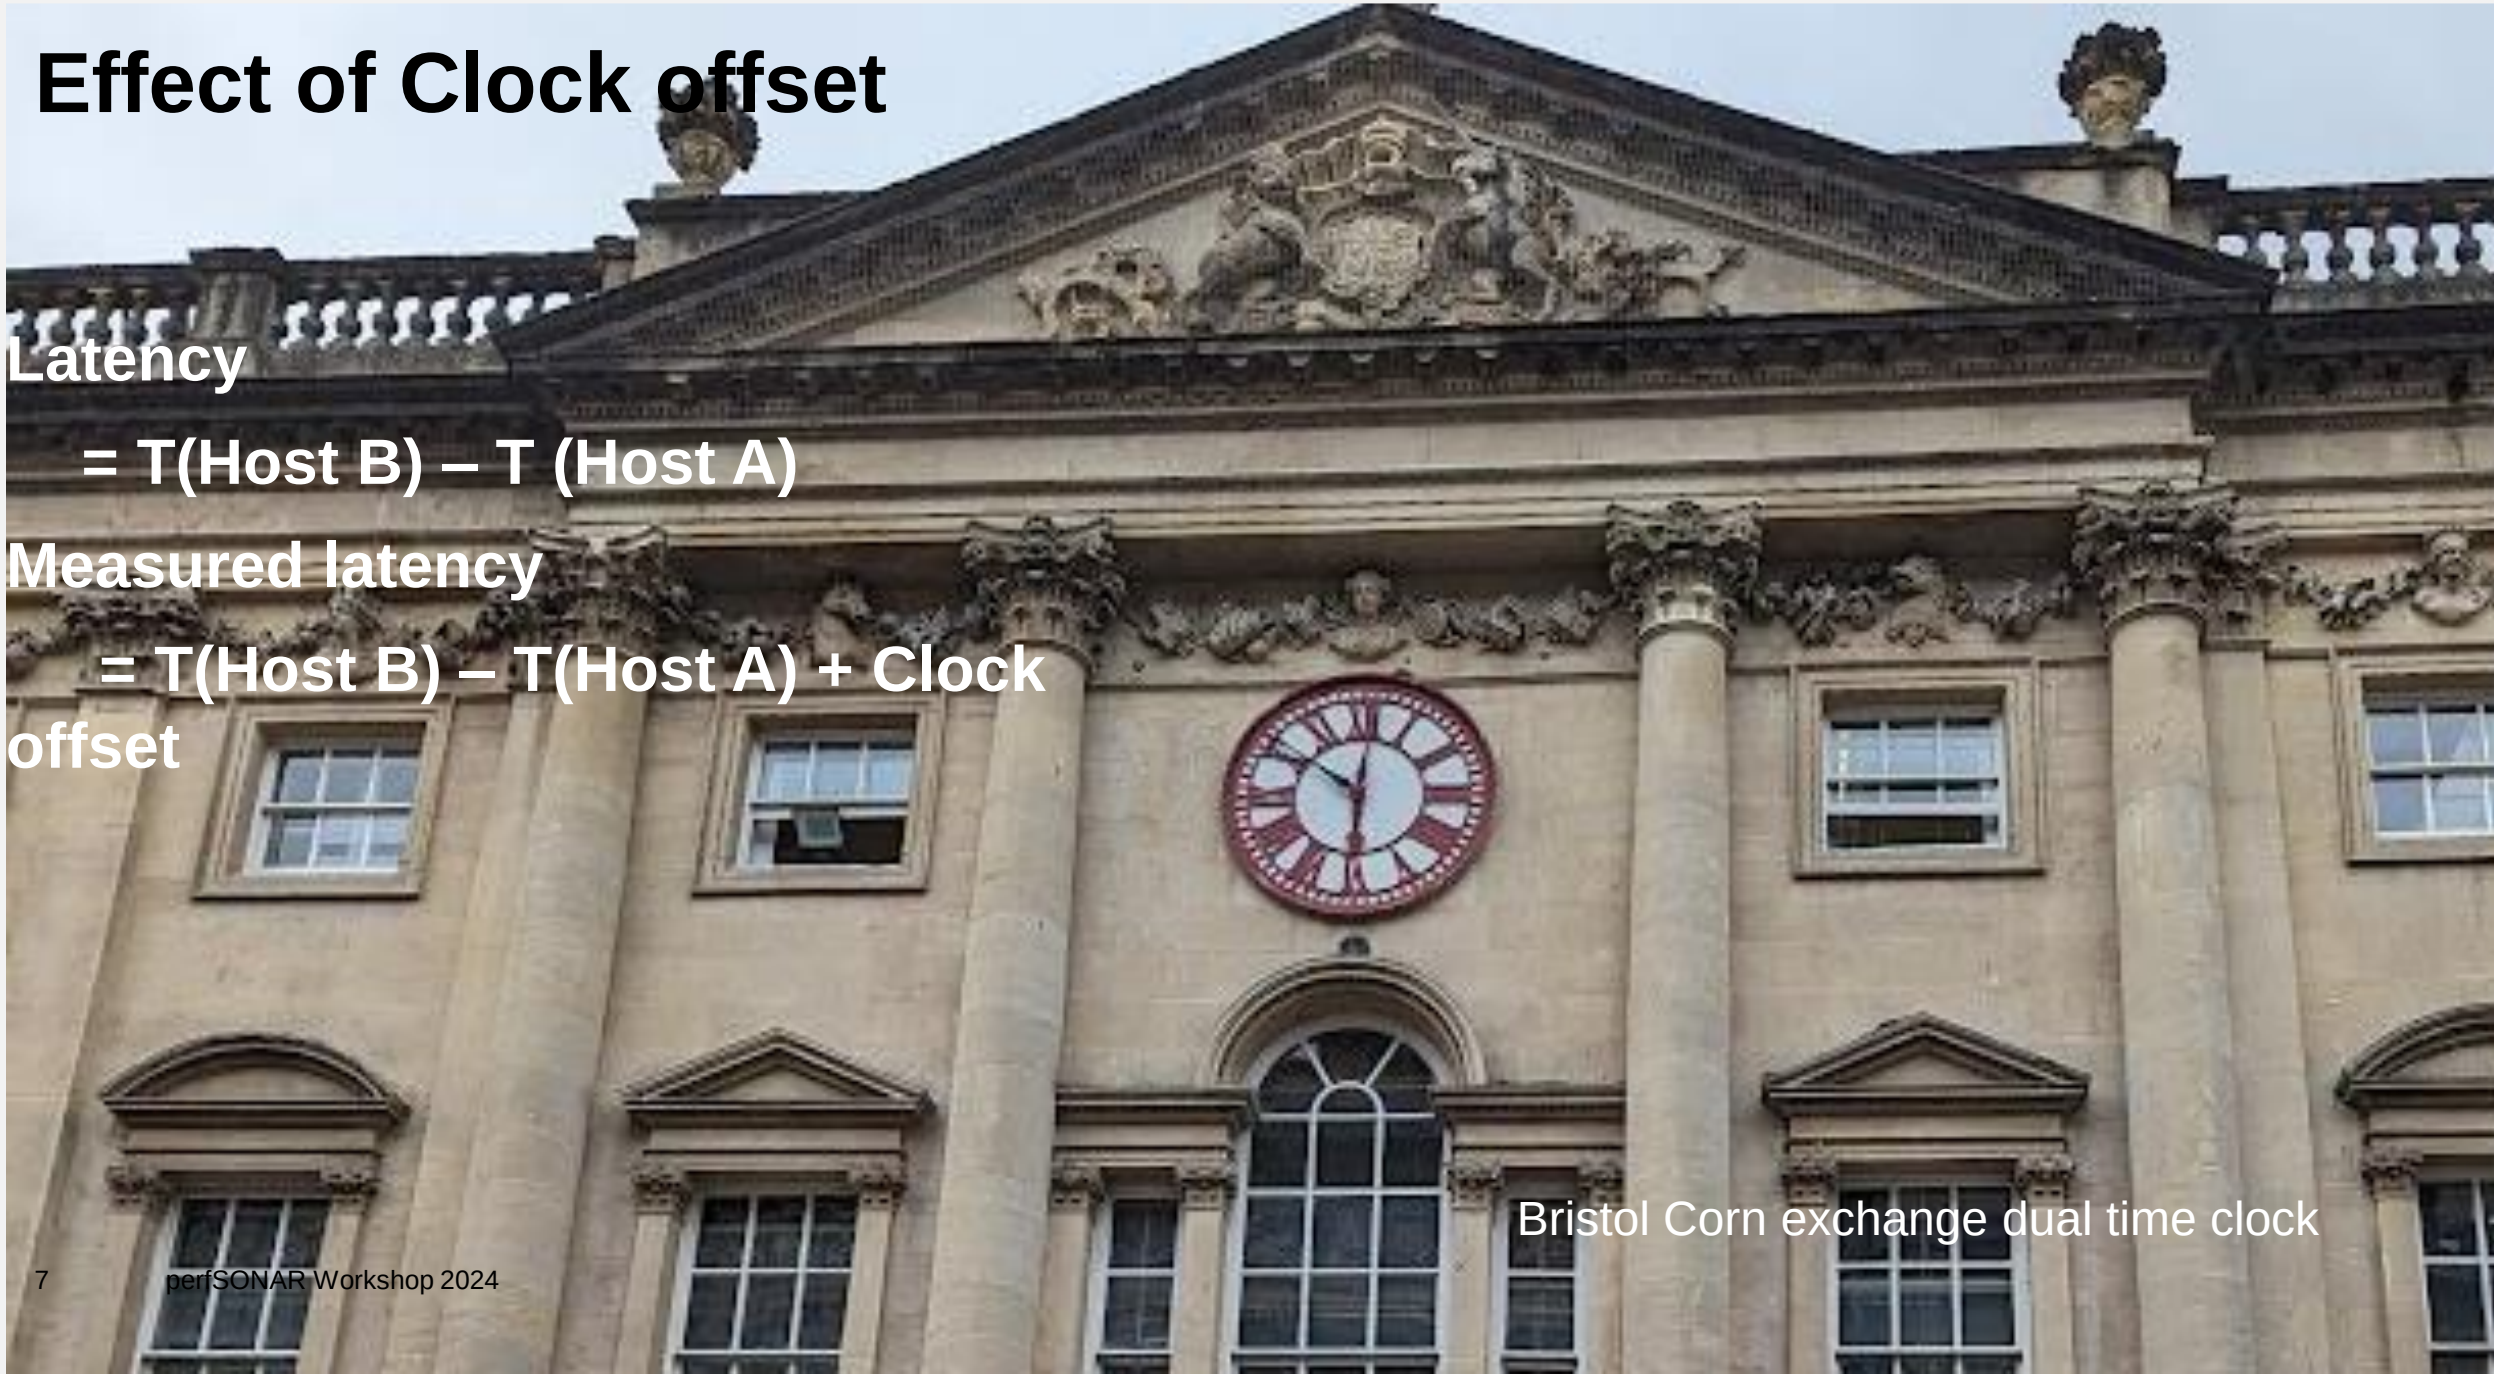

Bristol Corn exchange dual time clock

Network Card timestamping

RHEL 9 “Basic System Settings - Chrony”

- Network card support?

```
# ethtool -T enp175s0f0np0
```

- Output

```
Time stamping parameters for enp175s0f0np0:
Capabilities:
        hardware-transmit
        hardware-receive
        hardware-raw-clock
PTP Hardware Clock: 0
Hardware Transmit Timestamp Modes:
        off
        on
```

Network Card timestamping (2)

- Enabling in **/etc/chrony.conf**

```
# Enable hardware timestamping on all interfaces that support it.

hwtimestamp *
```

- Checking

```
# chronyc ntpdata
…
Offset        : -0.000132777 seconds
Peer delay    : 0.005870986 seconds
Peer dispersion : 0.000003835 seconds
Response time   : 0.000374075 seconds
…
TX timestamping : Hardware
RX timestamping : Hardware
```

Improving clock accuracy/stability

- Accurate time servers
  - How accurate/stable are yours?
    - RIPE ATLAS Measurements?
    - Jisc plans time from NPL
- More accurate time transfer
  - NTP (draft-ietf-ntp-over-ptp-02)
  - PTP
    - Simple PTP at Meta
  - White Rabbit

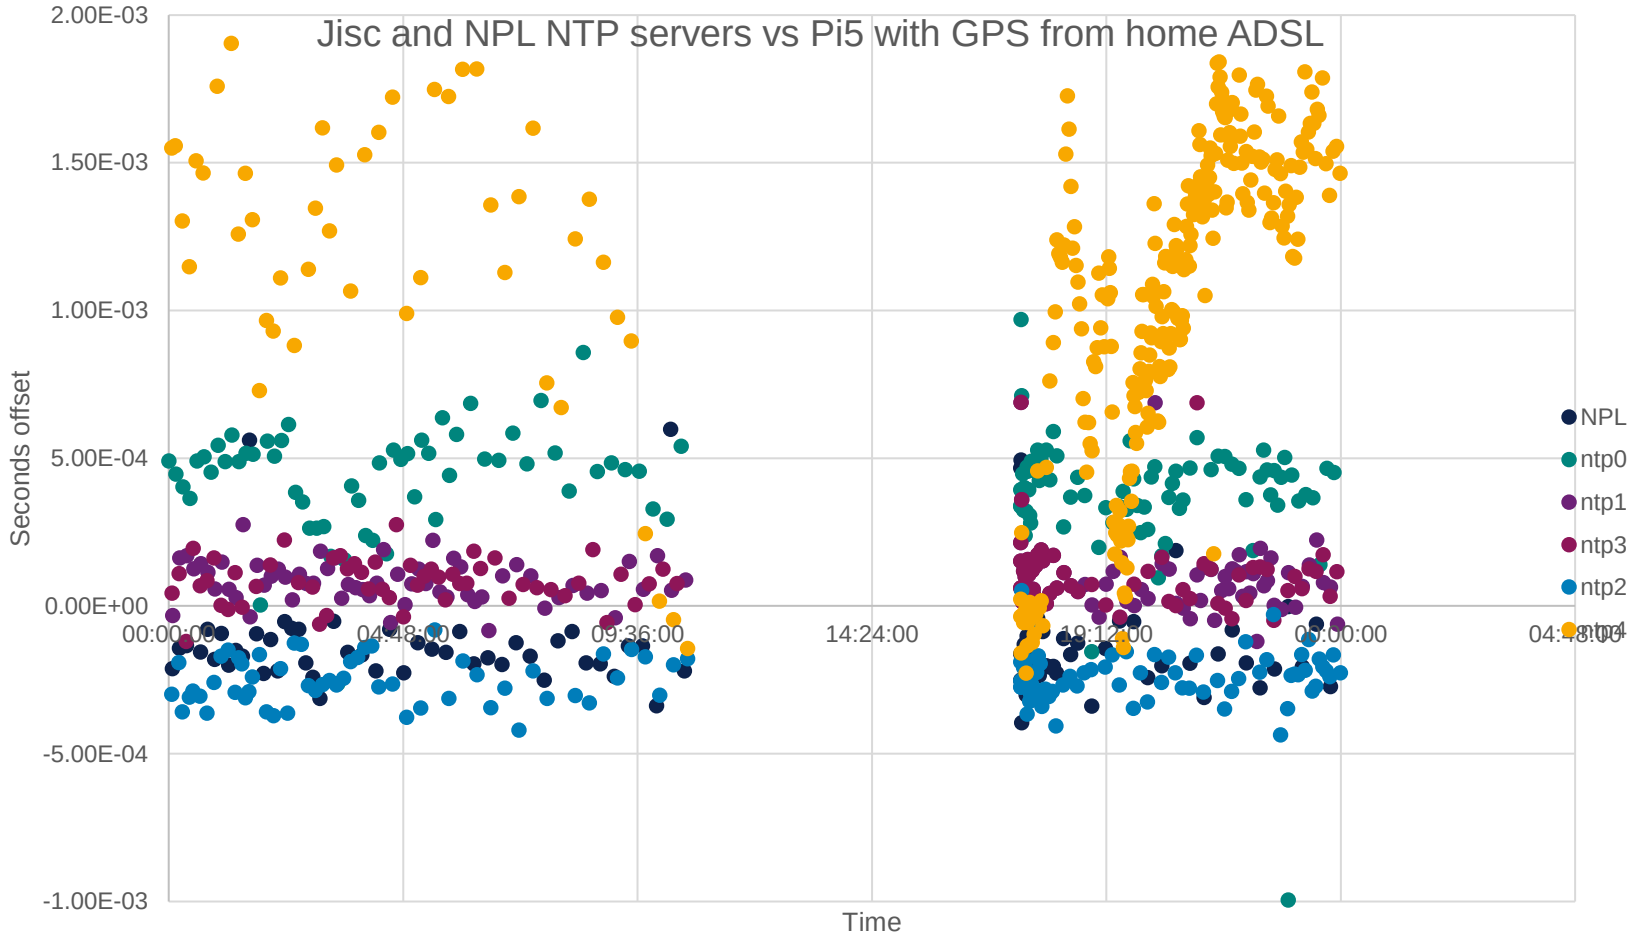

What accuracy is required

Consider 1250 Byte=10,000bit packet

- Packet quantisation time
  - 10Gbit : 1µs
  - 100Gbit : 0.1µs

- Clock stability:
  - ~ 500µs (0.5ms)
    - 500 packets at 10Gig
    - 5000 packets at 100Gig

- What accuracy do we need?

Conclusions

More Accurate measurements

- HW timestamping
  - Quick win and easy with chrony
- Time source accuracy
  - Currently limiting factor for one way latency measurement
- Future work
  - Could/should hardware timestamping be added to OWAMP?
  - Other things to improve latency measurements