Time-Granularity in Measurements

IST-SCAMPI workshop

Remco van de Meent

University of Twente

27-Jan-03

University of Twente
The Netherlands

Rationale

Overprovisioning:

long-term average load + 20% ?

Overview

Research Objectives

Approach

Measurements, case study

Results

Research Objectives

max. time-granularity

short-term “peak” load vs.
long-term average load

Time-Granularity

Common “MRTG measurements” use 5 minute throughput averages; no clue about what happens within this interval

Increase accuracy by decreasing Δt

Limited by timestamp precision, and physics

Approach

We only look at passive measurements, not disturbing normal traffic

Use off-the-shelf hardware & commodity software, for easy reproduction

Case study: measurements on *xxx* Mbit/s networks

Related Work

- DAG cards (special hardware)
- TICKET (special kernel)
- IPmon (special hardware)
- NeTraMeT, NetFlow (flow oriented, granularity)
- …

Case Study

- UT residential network
  2000 x 100 Mbit/s
  300 Mbit/s uplink
  long-term average load ~50%

- Hosting provider network
  50 x 100 Mbit/s
  100 Mbit/s uplink
  long-term average load <10%

UT resnet

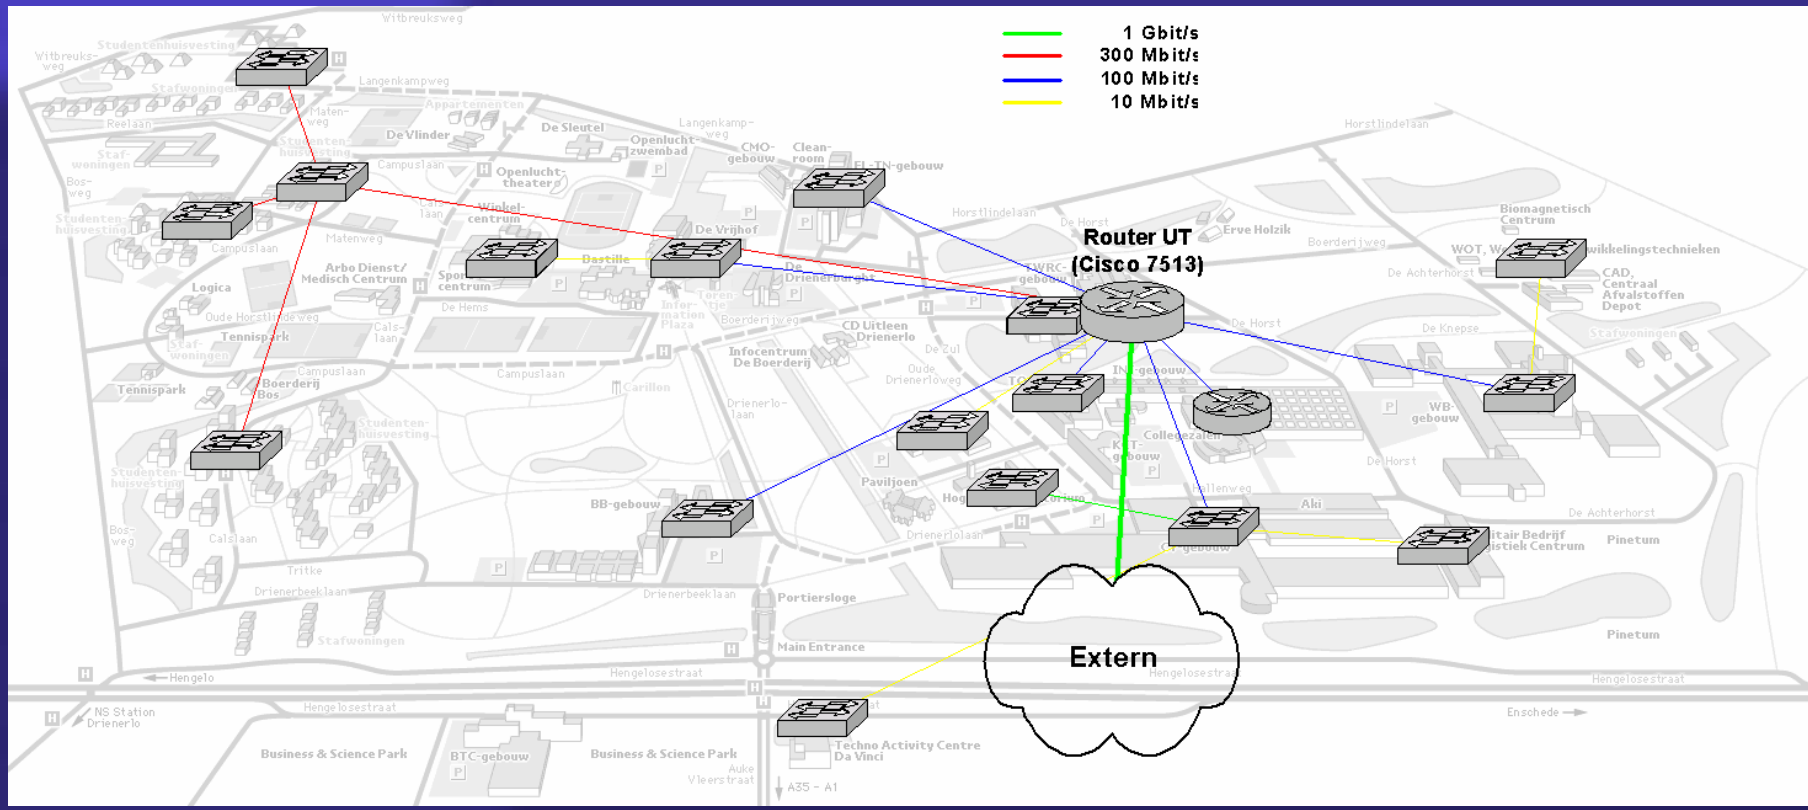
1 Gbit/s
300 Mbit/s
100 Mbit/s
10 Mbit/s
Router UT
(Cisco 7513)
Extern

Measurement Setup

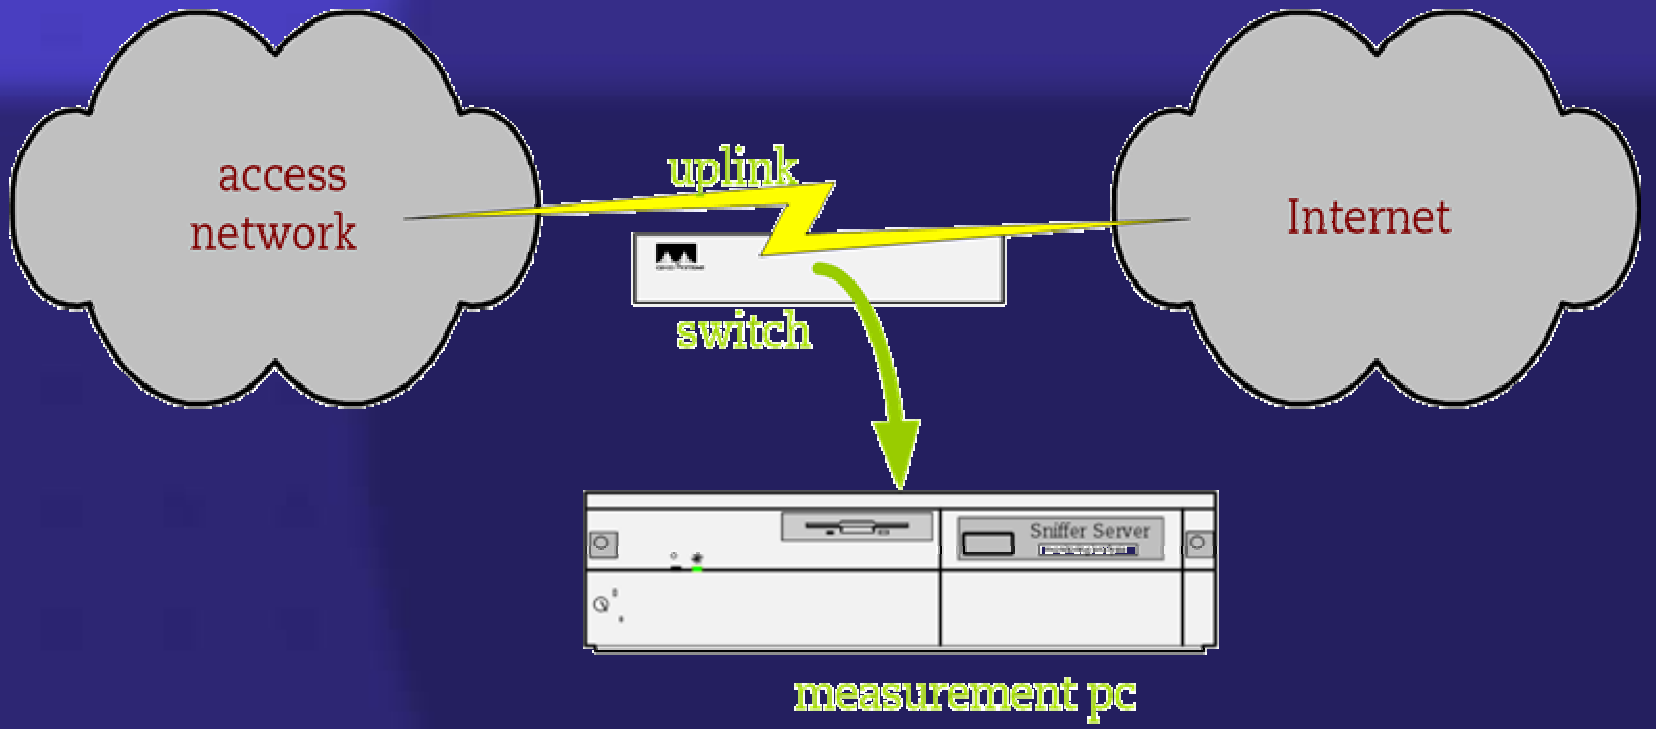

Measurement PC

Pentium III 1GHz
Linux 2.4 kernel
Gbit ethernet
lots of disk

libpcap, tcpdump
timestamping
store headers

Max. Time-Granularity (1)

- Reading IF-MIB counters (with SNMP) takes too much time
- For throughput-measurements with high time-granularity, NeTraMeT / NetFlow is less suited
- Capture and store traffic, process offline

Max. Time-Granularity (2)

- Buffering in switch while copying data to monitor port
- Full-duplex v. half-duplex
- Time difference between forwarding and timestamping
- Resolution of the timestamp

Max. Time-Granularity (3)

- Linux/x86: resolution 100Hz (10ms)
- Timestamped in netif_rx()
- TSC/mmtimer improve this to nanosec resolution, not fully used though
- → time-granularity: 10 ms

UT resnet: throughput (1)

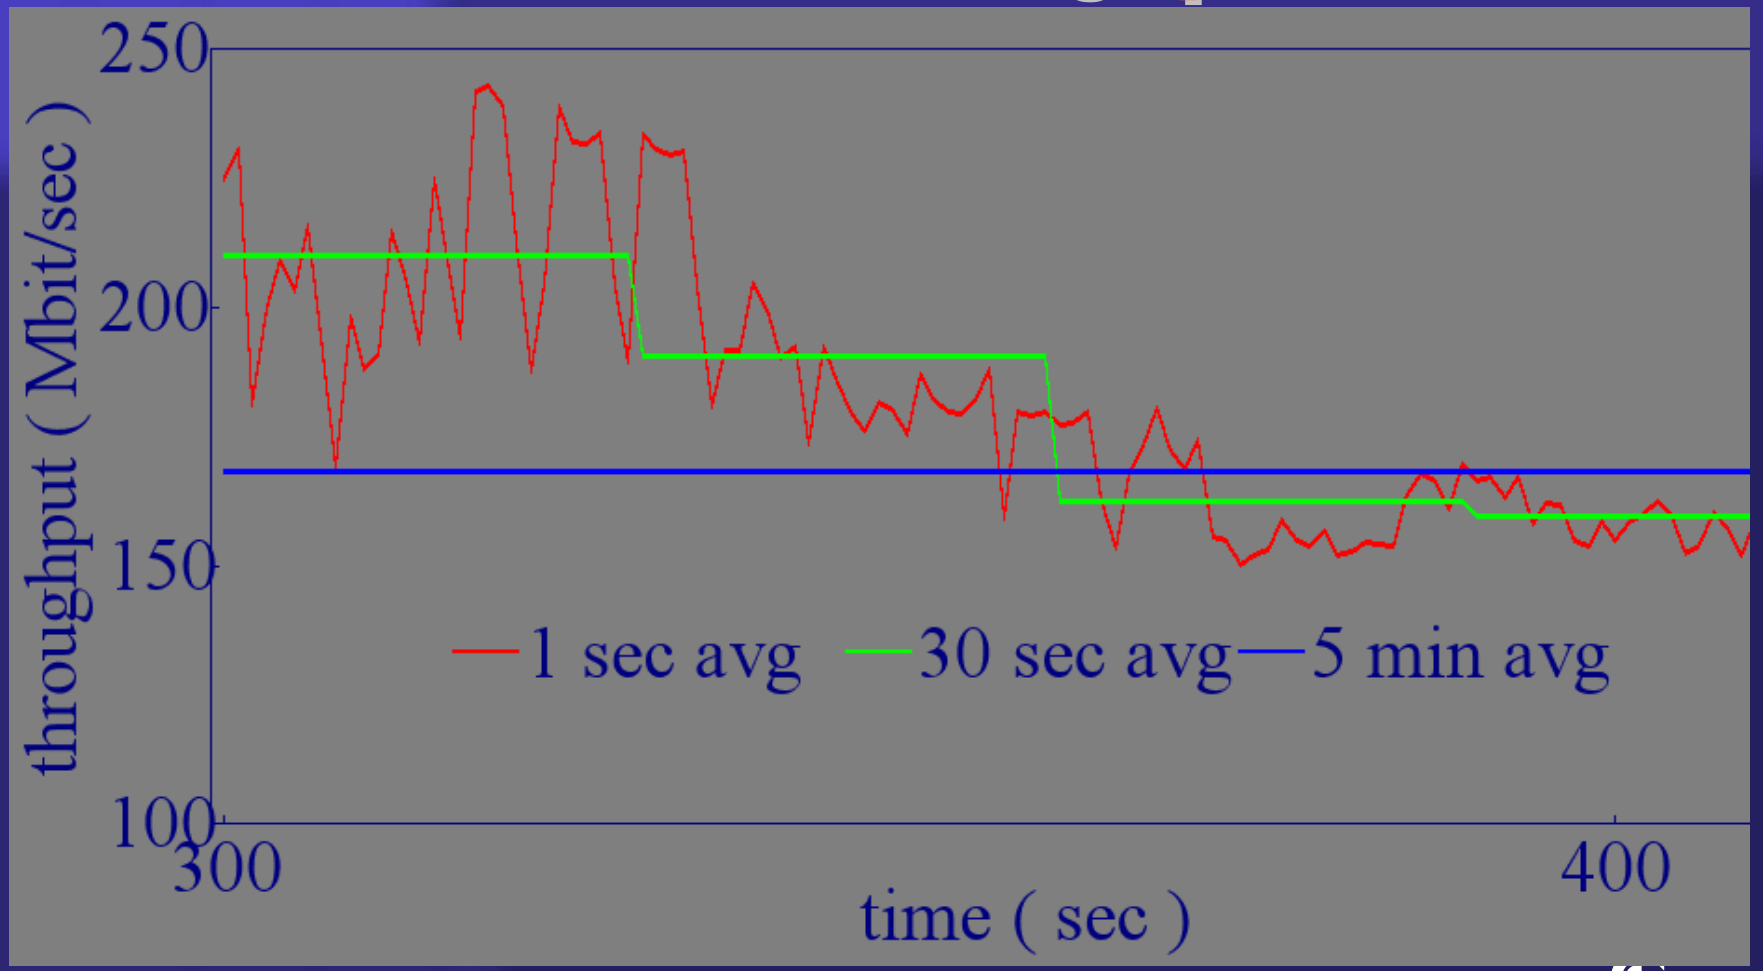

UT resnet: throughput (2)

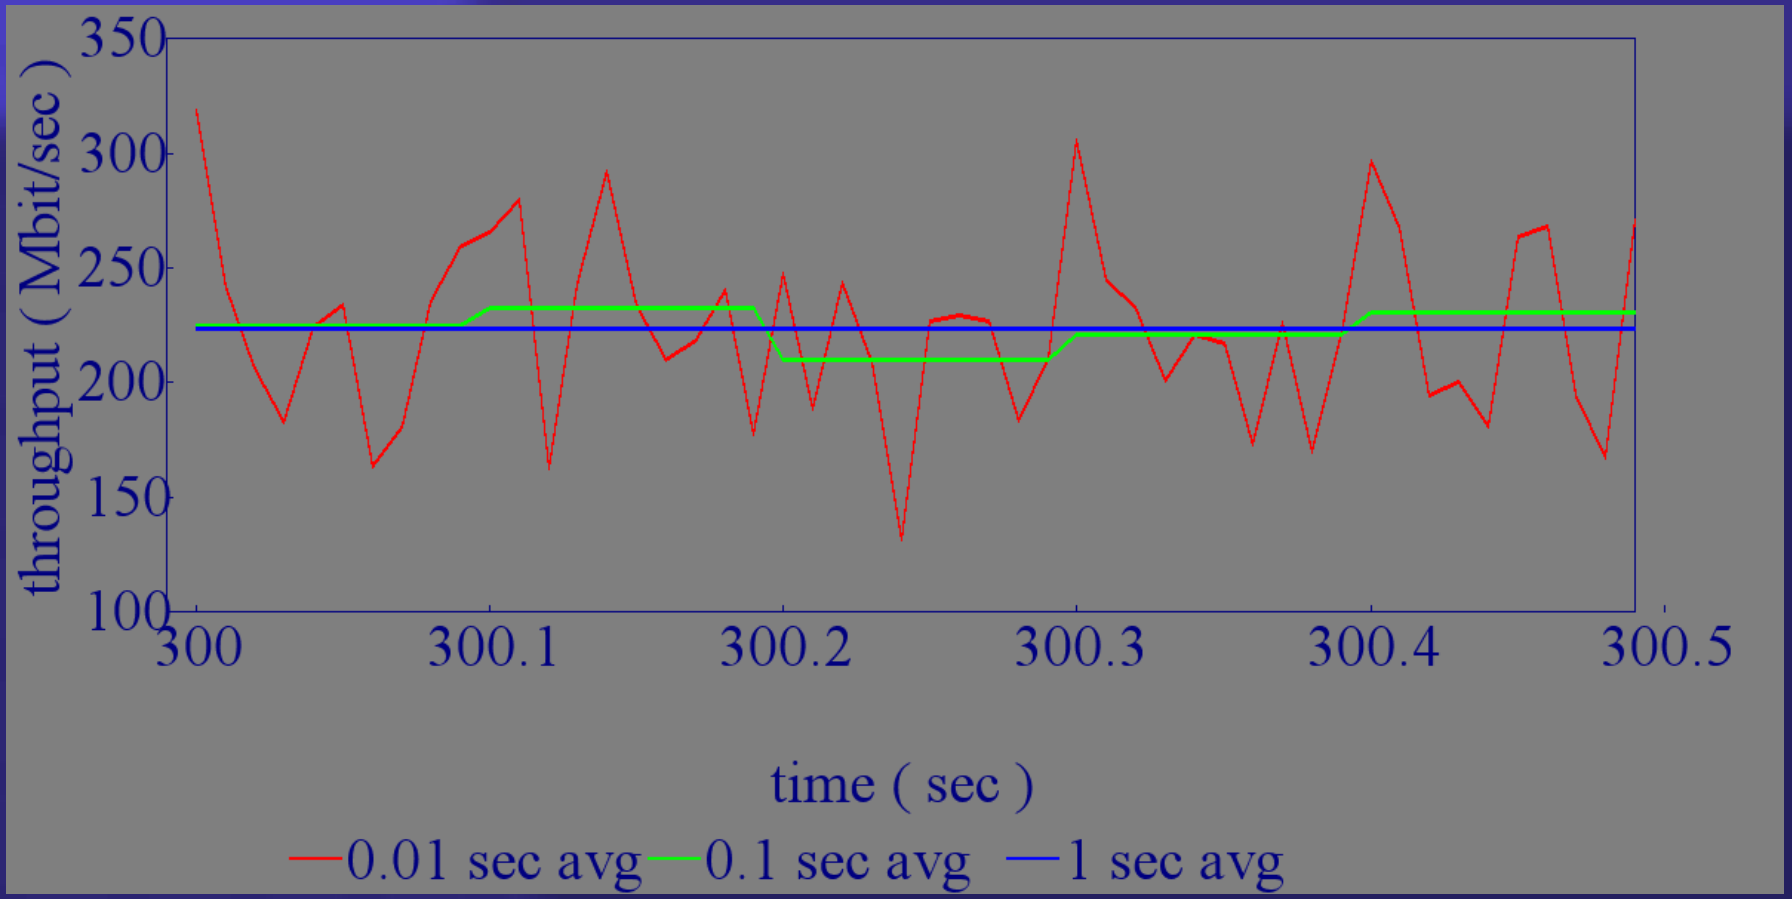

UT resnet: throughput (3)

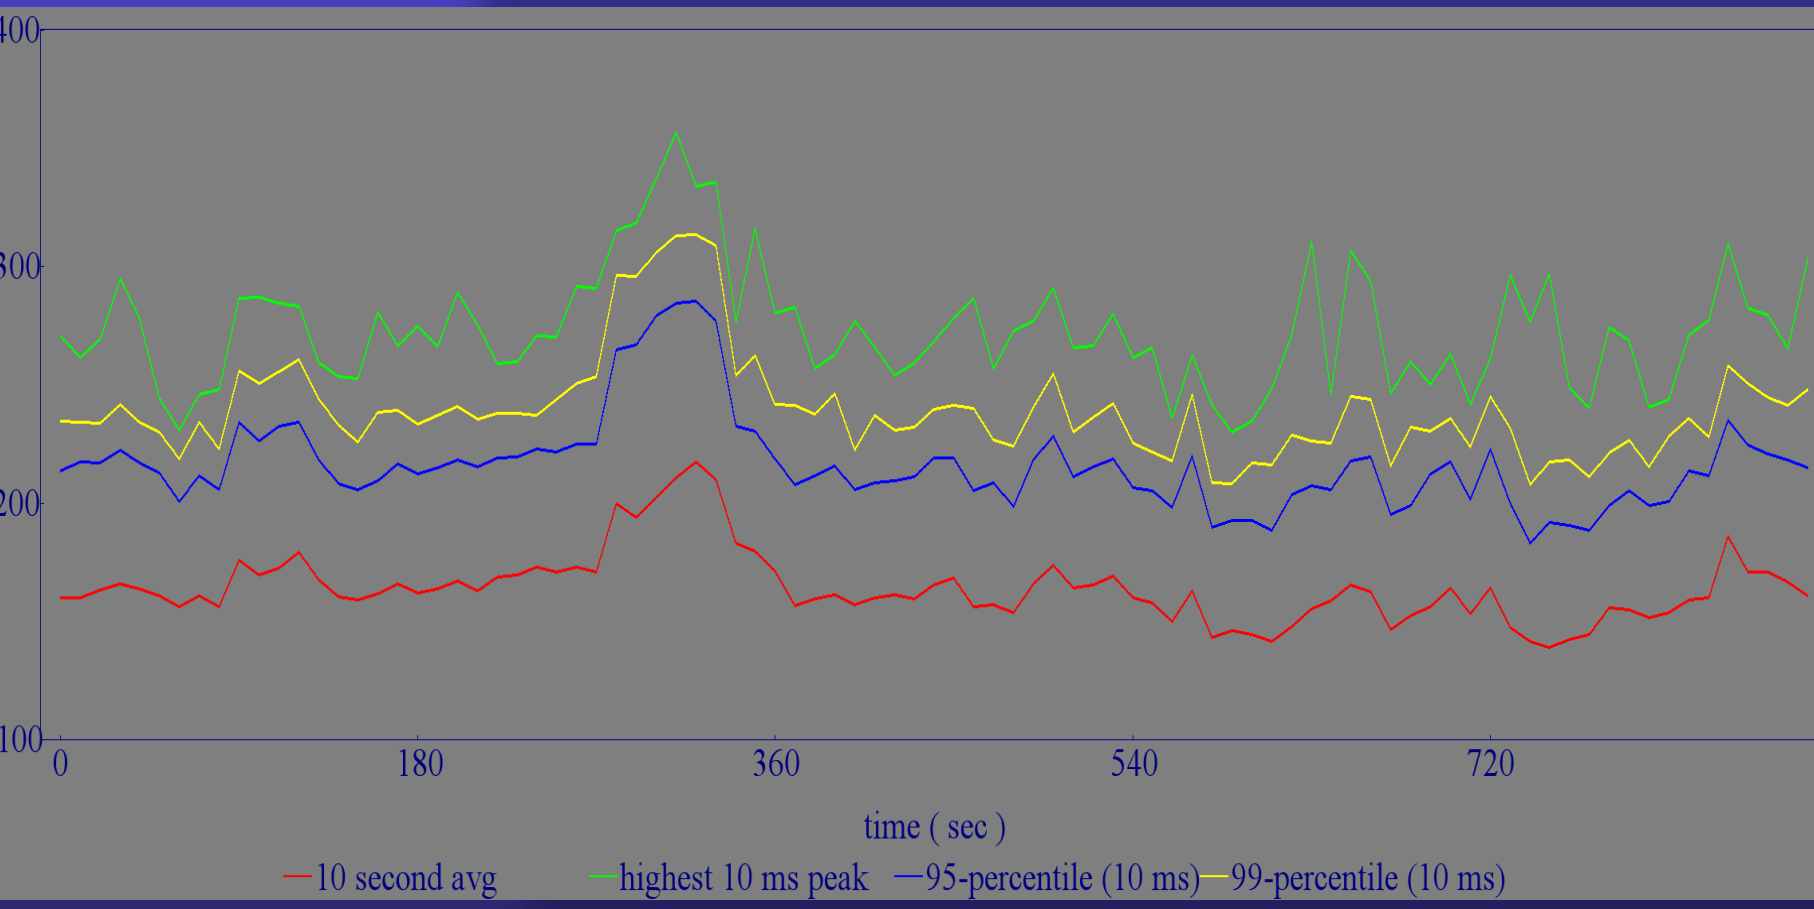

Hosting Provider network: throughput

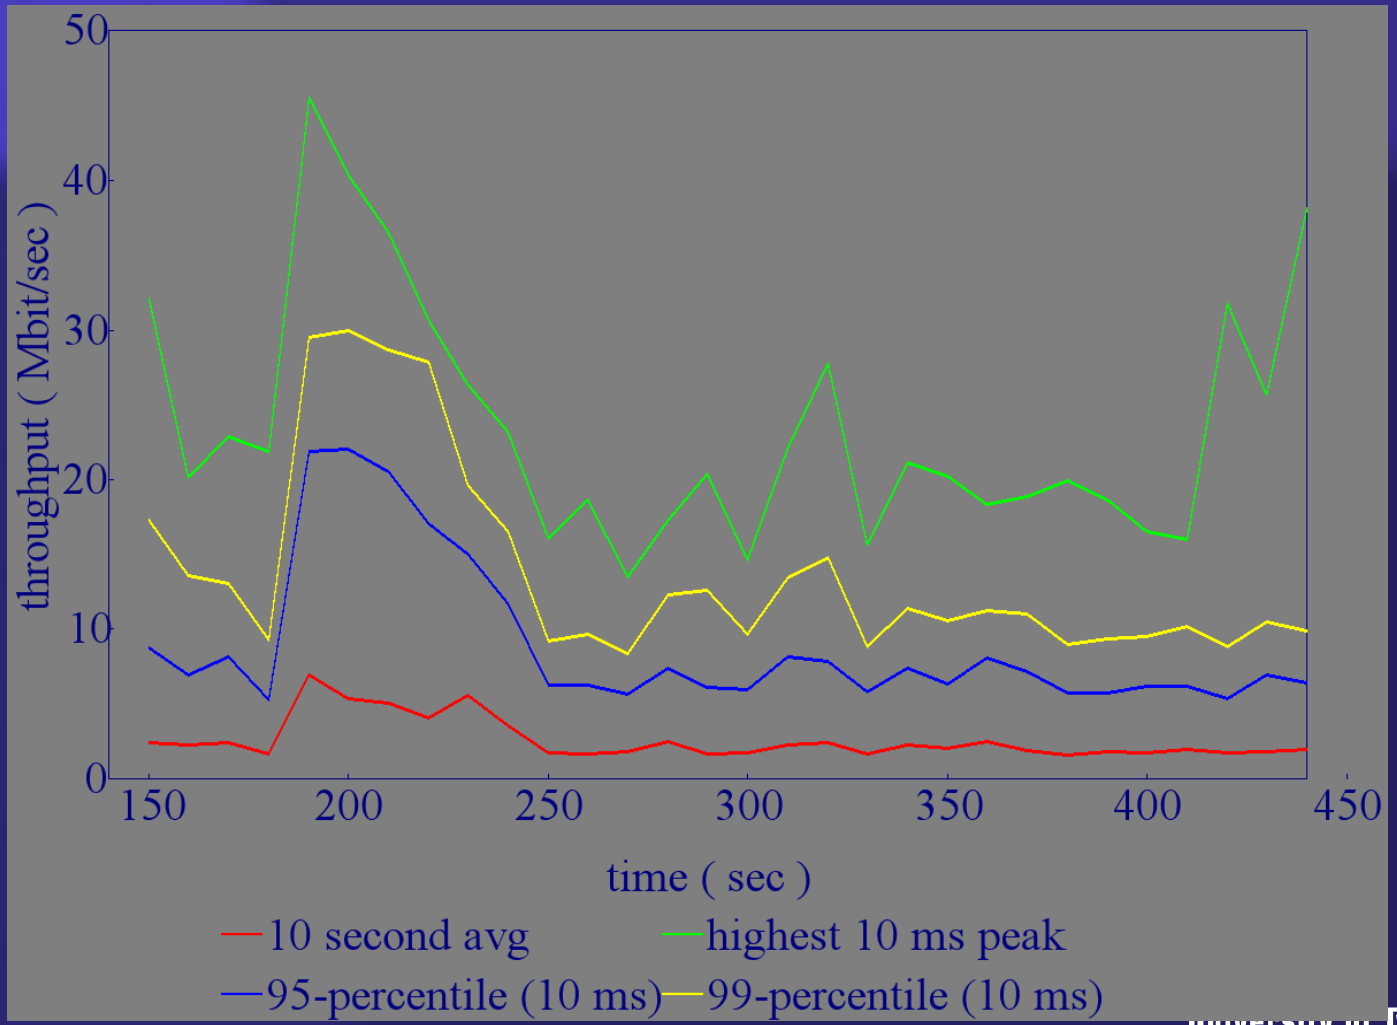

Conclusions

- 10ms time-granularity is achievable
- Considerable difference between ``peak'' and ``average''
- *Intelligent overprovisioning*: we don't know yet