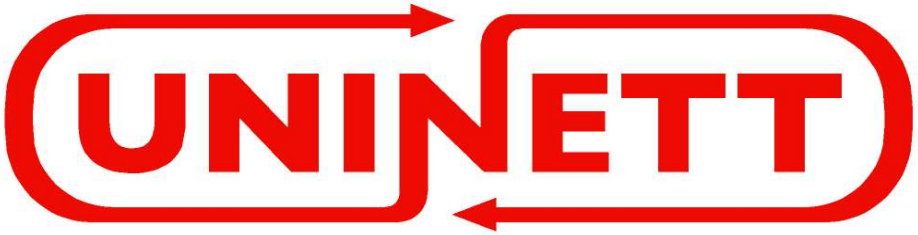

Micro-measurements in high capacity networks

21. NORDUnet Network Conference

August 25th, 2003

Jon Kåre Hellan

What do we know about our traffic?

- We monitor link traffic
- UNINETT statistics:

  Year, month, week, day

  Resolution: week – 2 minutes

http://drift.uninett.no/stats/tbl-report.html

http://drift.uninett.no/kartg/last/uninett/norge/geo/nuh

UNINETT

A day in the life of a gigabit link

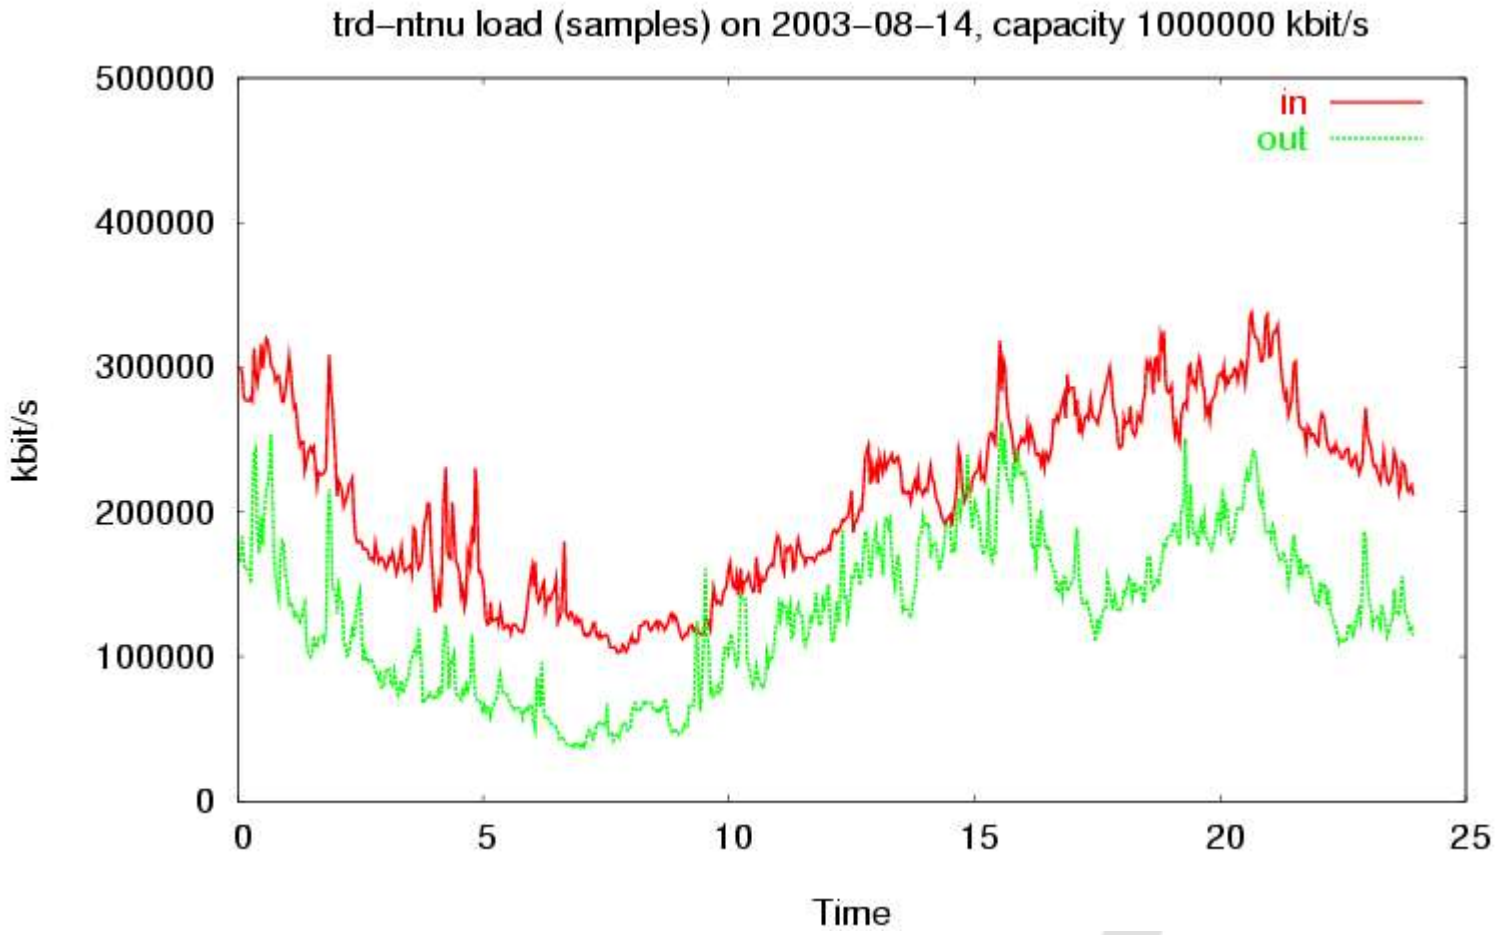

UNINETT

The microscopic view

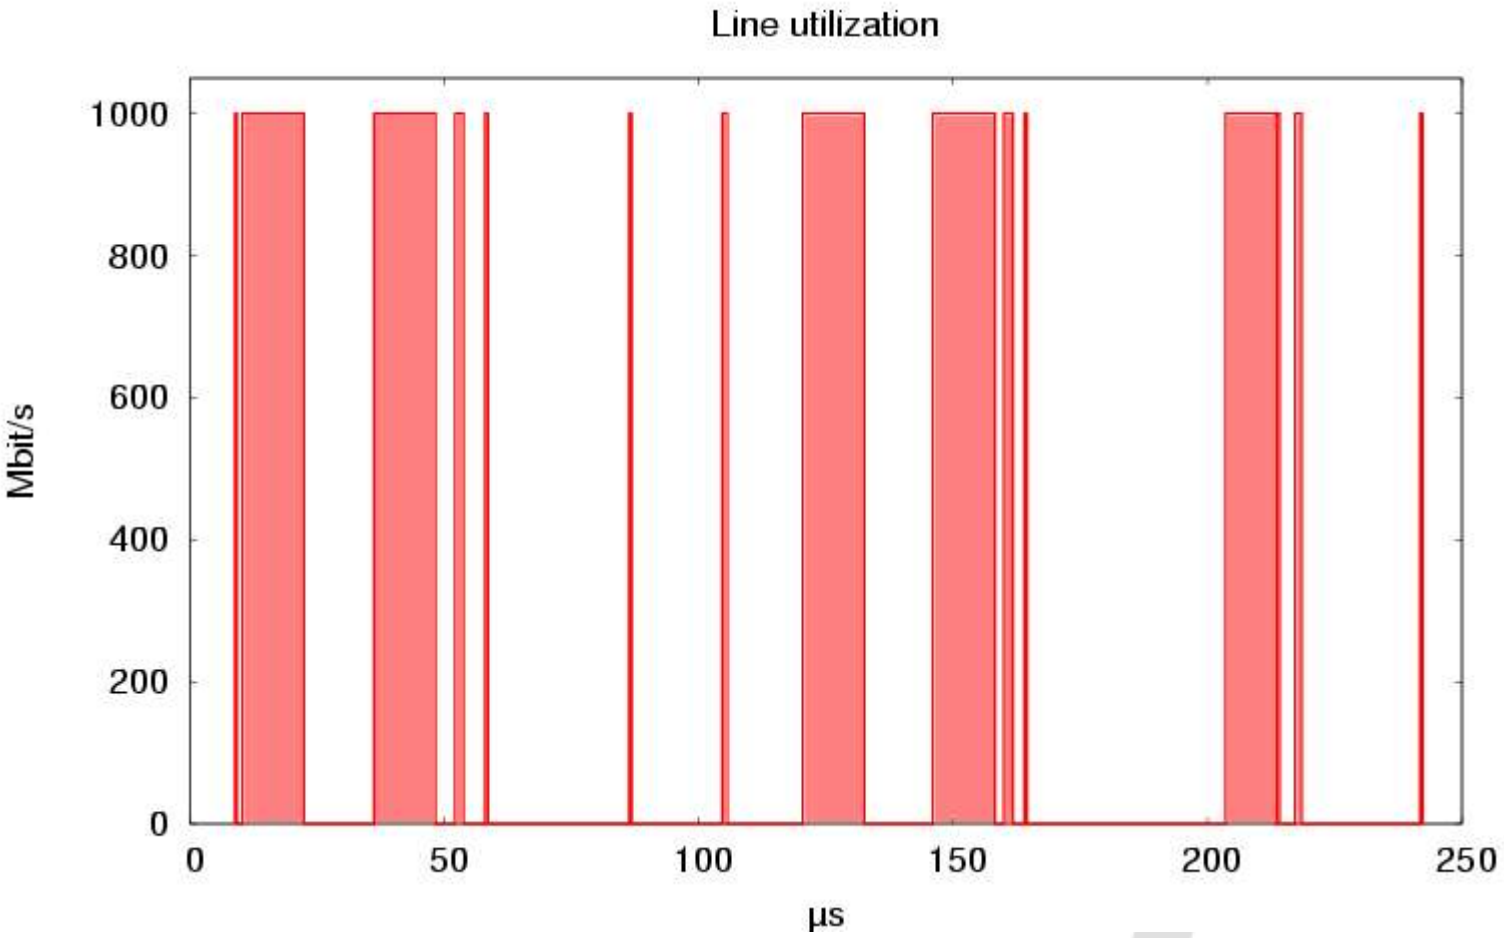

The load is low. Are we safe?

- What happens between the samples?
- Network traffic is bursty
  - Models which assume statistical independence are too optimistic
- Self similarity has been observed
  - Traffic patterns are said to look the same when studied at very different time scales
  - Self similar traffic sources do not aggregate well
  - Studies from the 90's. Still relevant?

Measurement tools

- SNMP counters

  Used to make UNINETT traffic statistics

- Hardware probes
  - DAG cards from Endace

    Up to 2.5 Gbit/s

    UNINETT has 3
  - COMBO-6, Masaryk University

    In development, part of SCAMPI project.

UNINETT

SNMP polling with a twist

- "Di-daaah" polling
  - For each regular sample, also take a shorter sample
- But
  - SNMP counters are updated infrequently
  - Updating them has low priority
- So on Cisco core routers, measuring < 5s yields unreliable results

The cloud should be around the lines

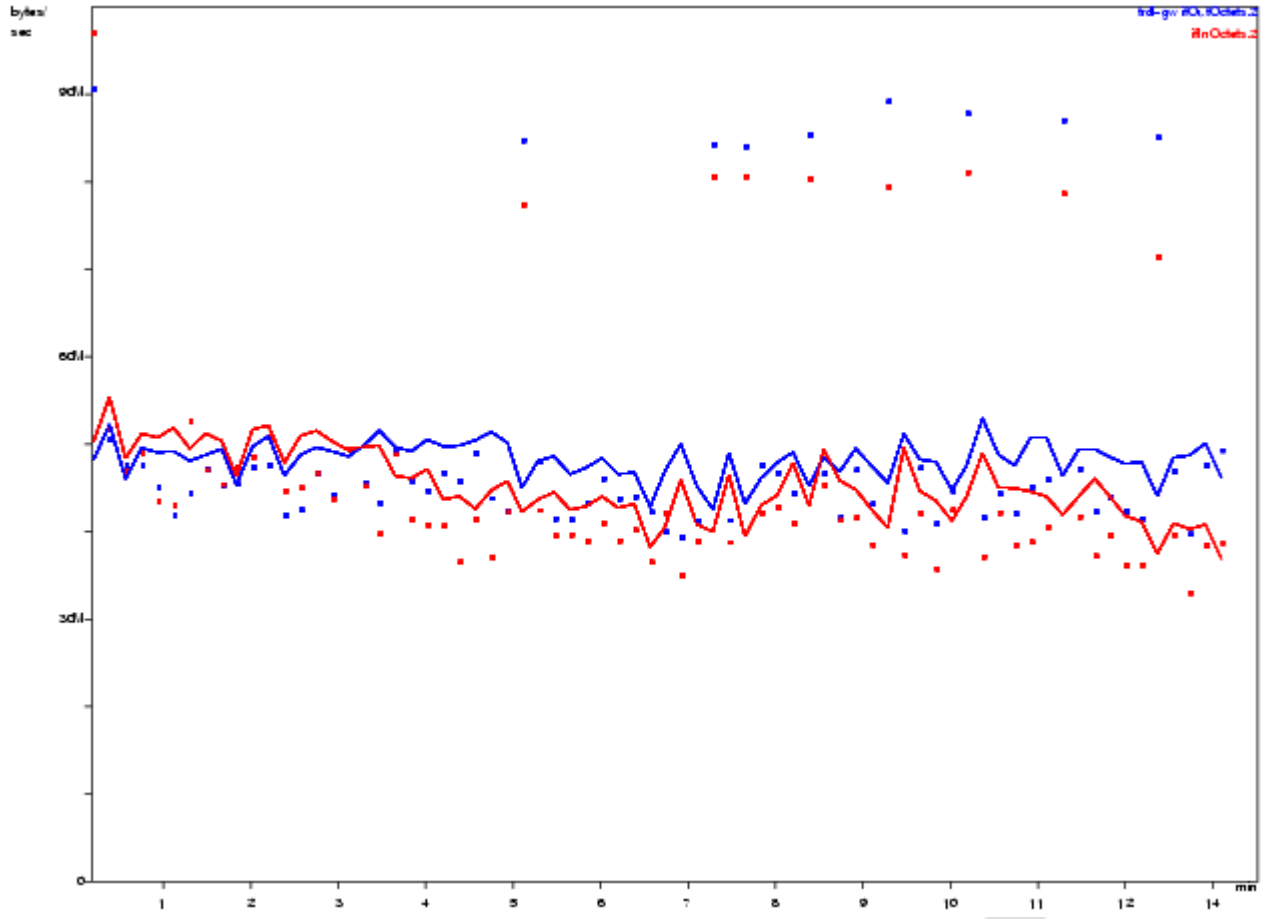

What went wrong?

- Imagine that router updates counters once every long sampling period.
- Two cases
  - Update outside short sampling period
    - looks like zero traffic
  - Update inside short sampling period
    - looks like all traffic happens here
- Another idea: measure queues, but same problem

Is there a better way?

- Hardware probes
- PC card
  - Time source
    - Crystal
    - GPS
  - Optical splitter
- UNINETT uses DAG from Endace
- Masaryk University is developing Combo 6

UNINETT

Zooming in, shortening periods

10 s

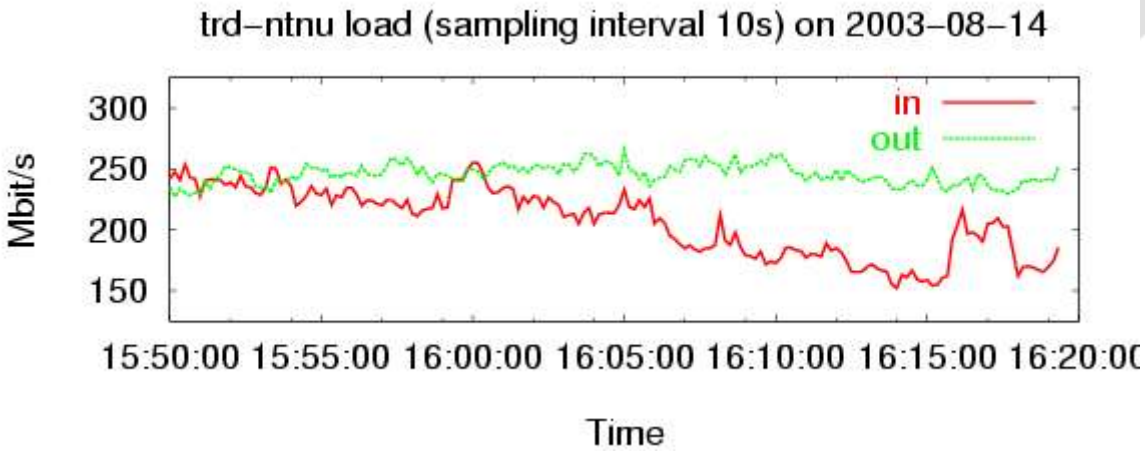

1 s

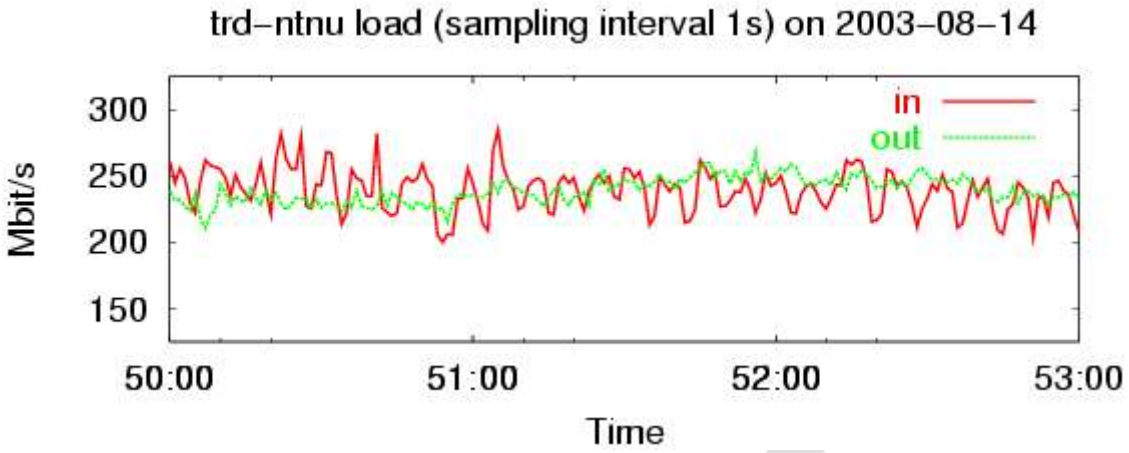

Zooming in (different trace)

1 s

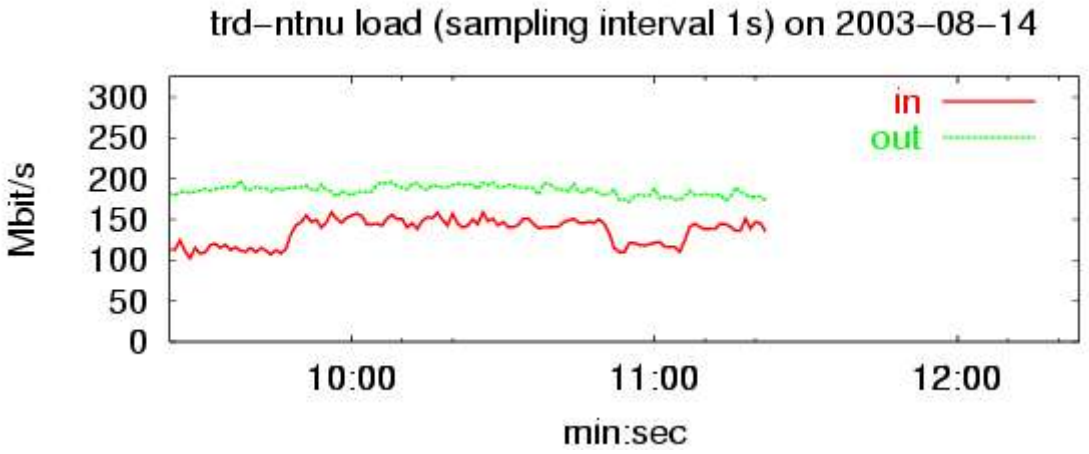

100 ms

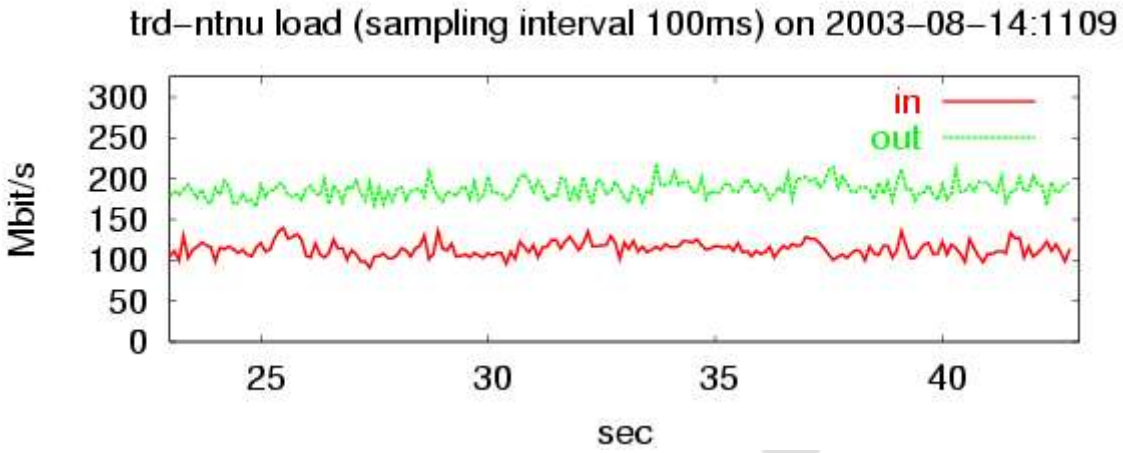

UNINETT

Still zooming

100 ms

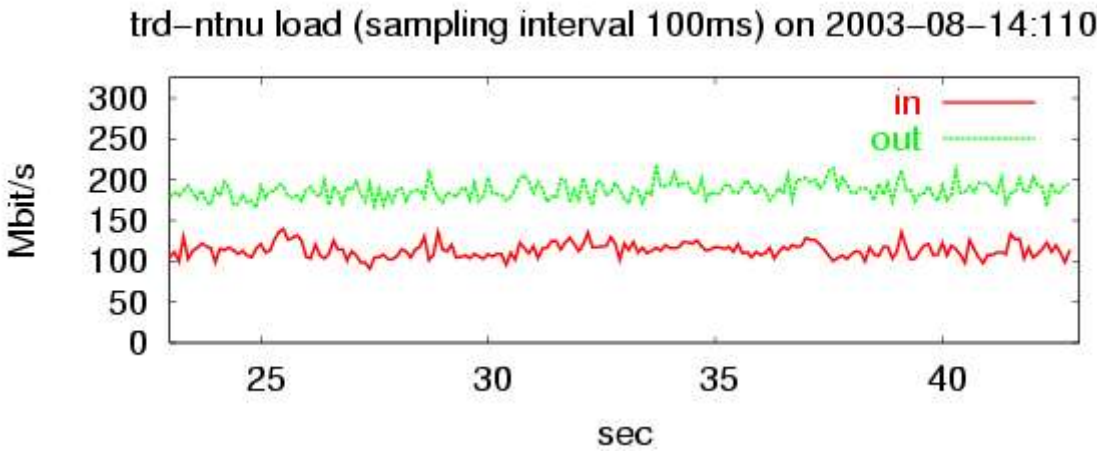

10 ms

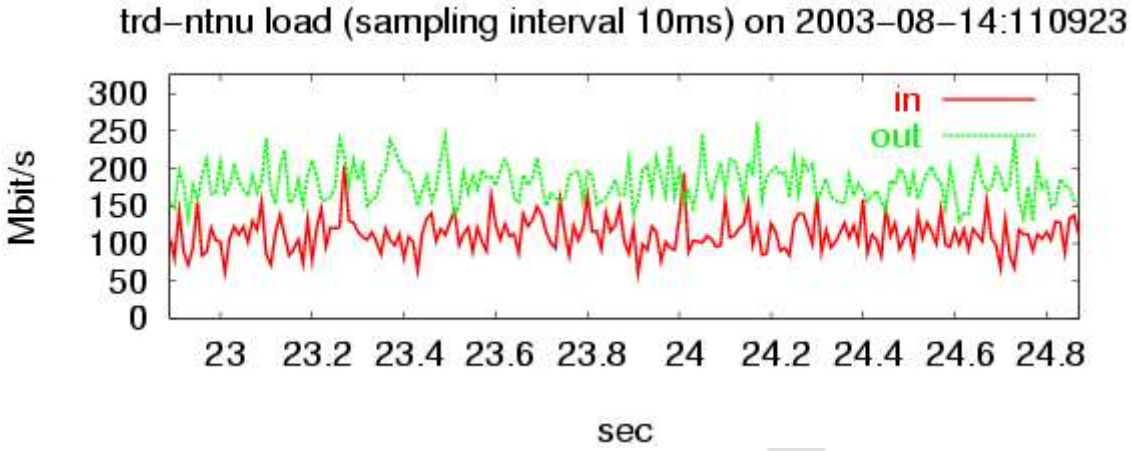

UNINETT

Are you bored yet?

10 ms

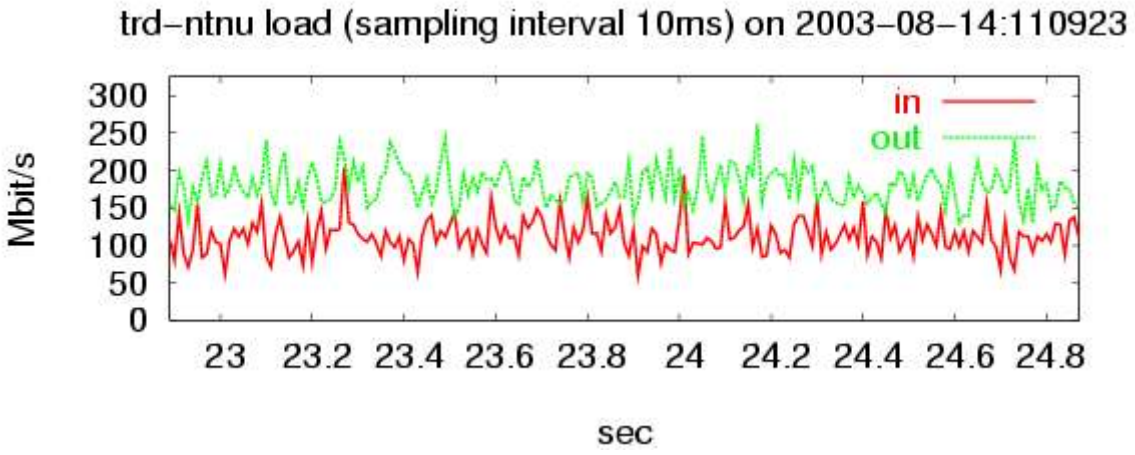

1 ms

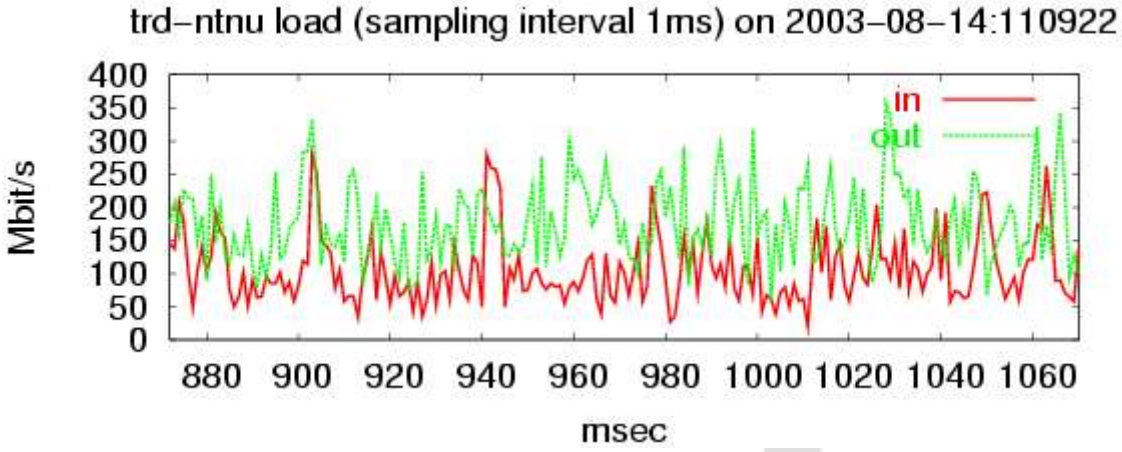

UNINETT

As short as it makes sense to go

1 ms

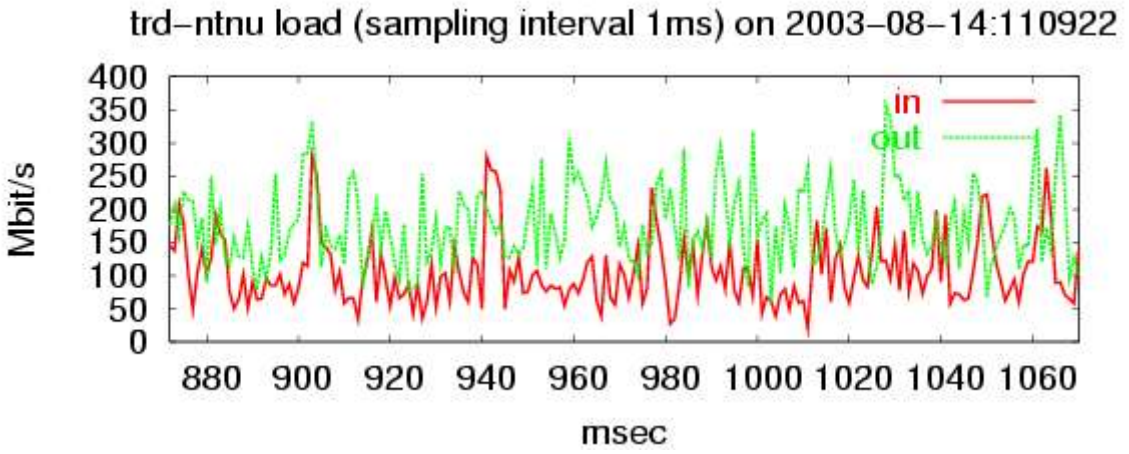

100 µs

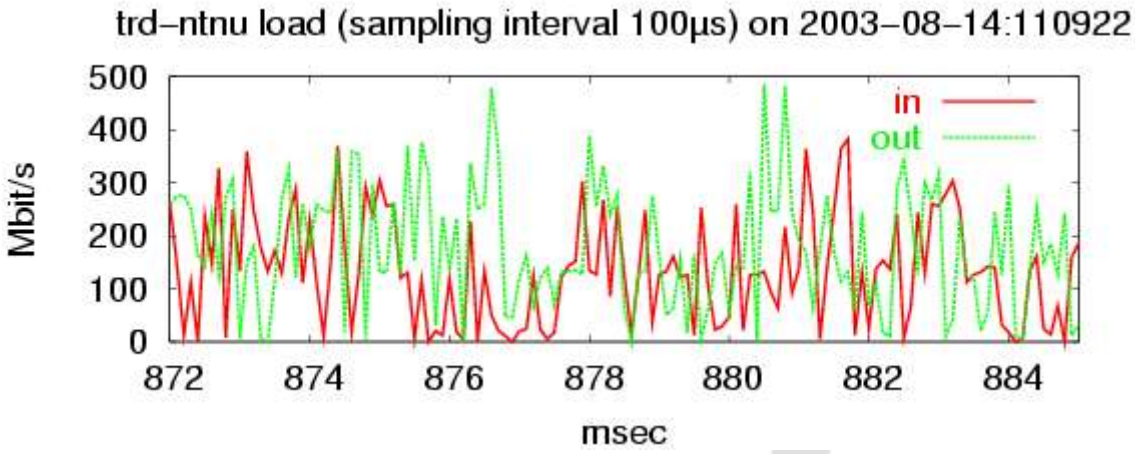

UNINETT

Is this traffic self-similar?

- That's how it looks
- But optical illusion is possible
  - Each plot shows the same number of samples
  - Consecutive samples are different
  - Upward and downward change equally likely
- However that may be, it is certainly bursty

It is indeed self-similar!

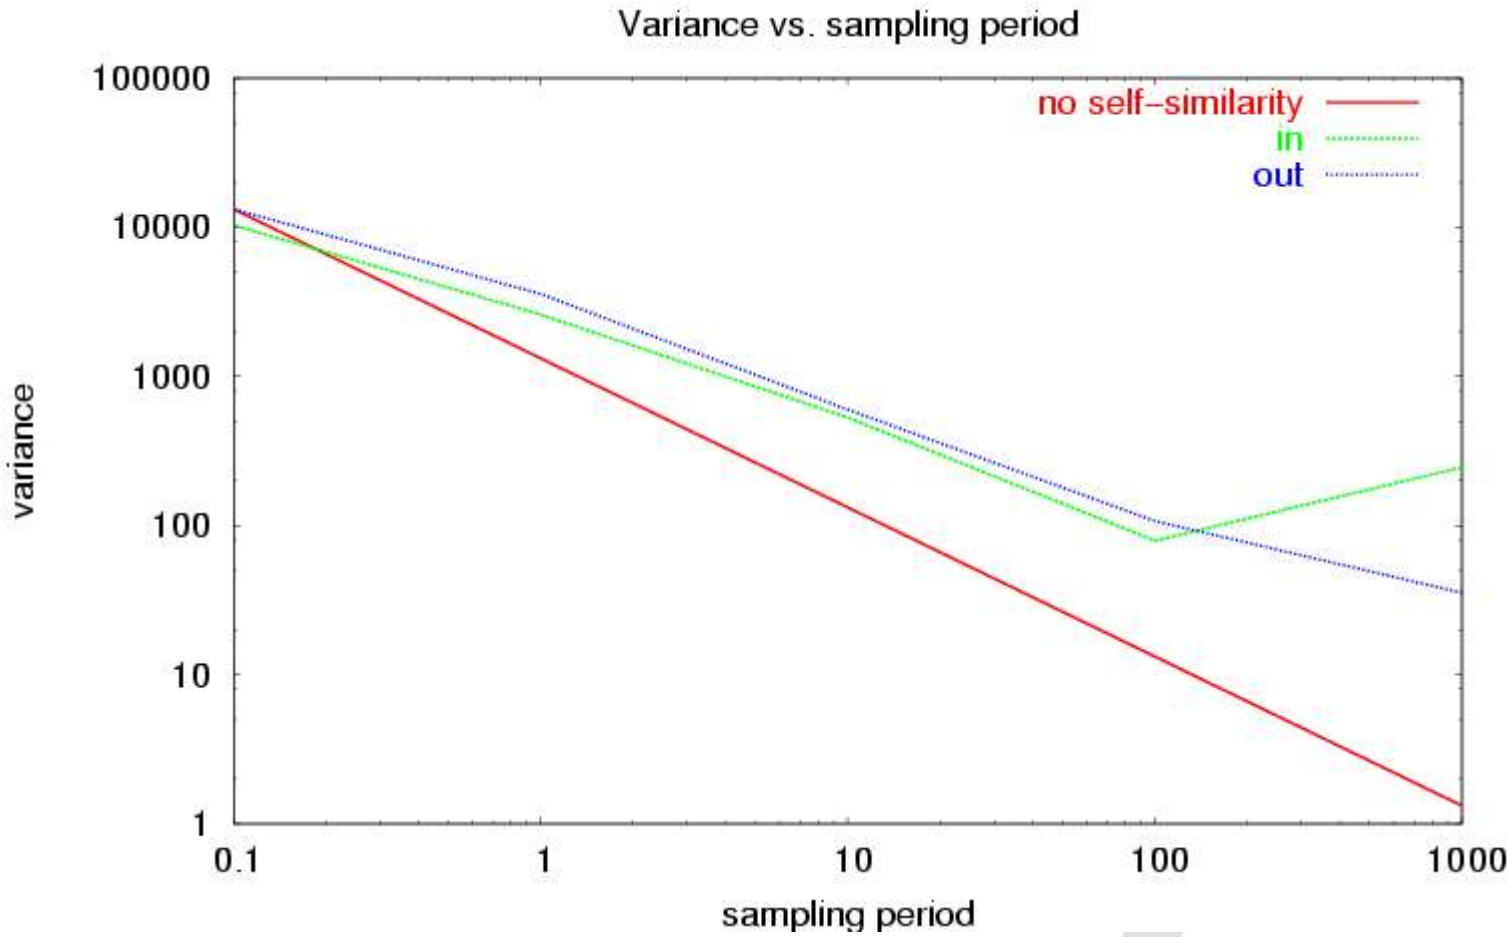

UNINETT

Inter-arrival times

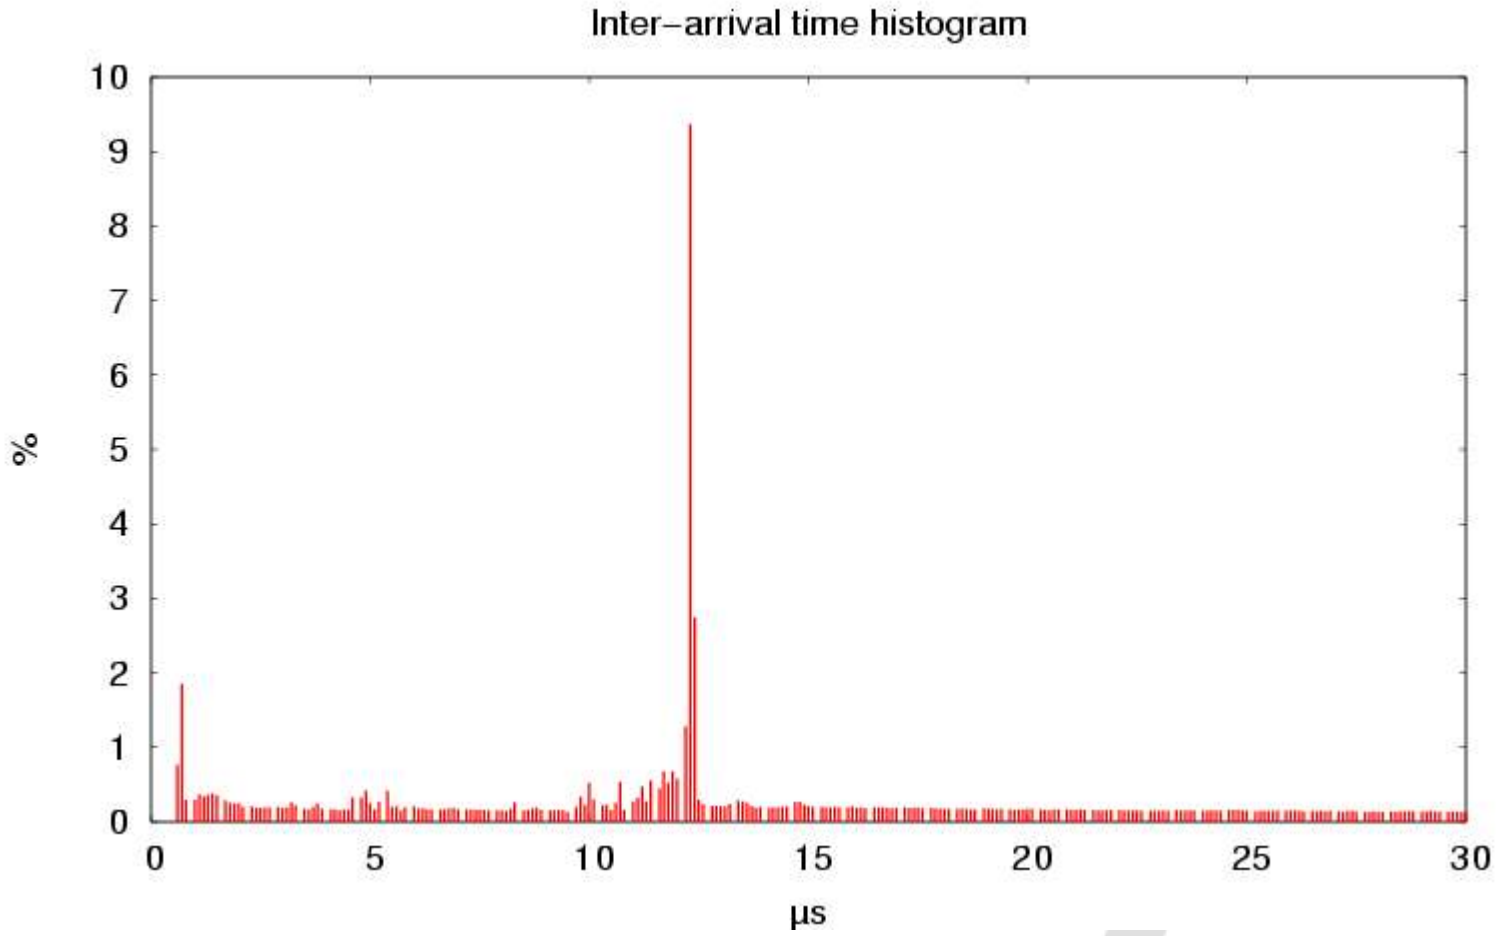

Inter-arrival times

- Two peaks, max and min frame length
  - 13.8% ≈ 12.3 µs
  - 13.0% 0.7-0.9 µs
- Min or max length frames sent just after another frame

UNINETT

Frame lengths

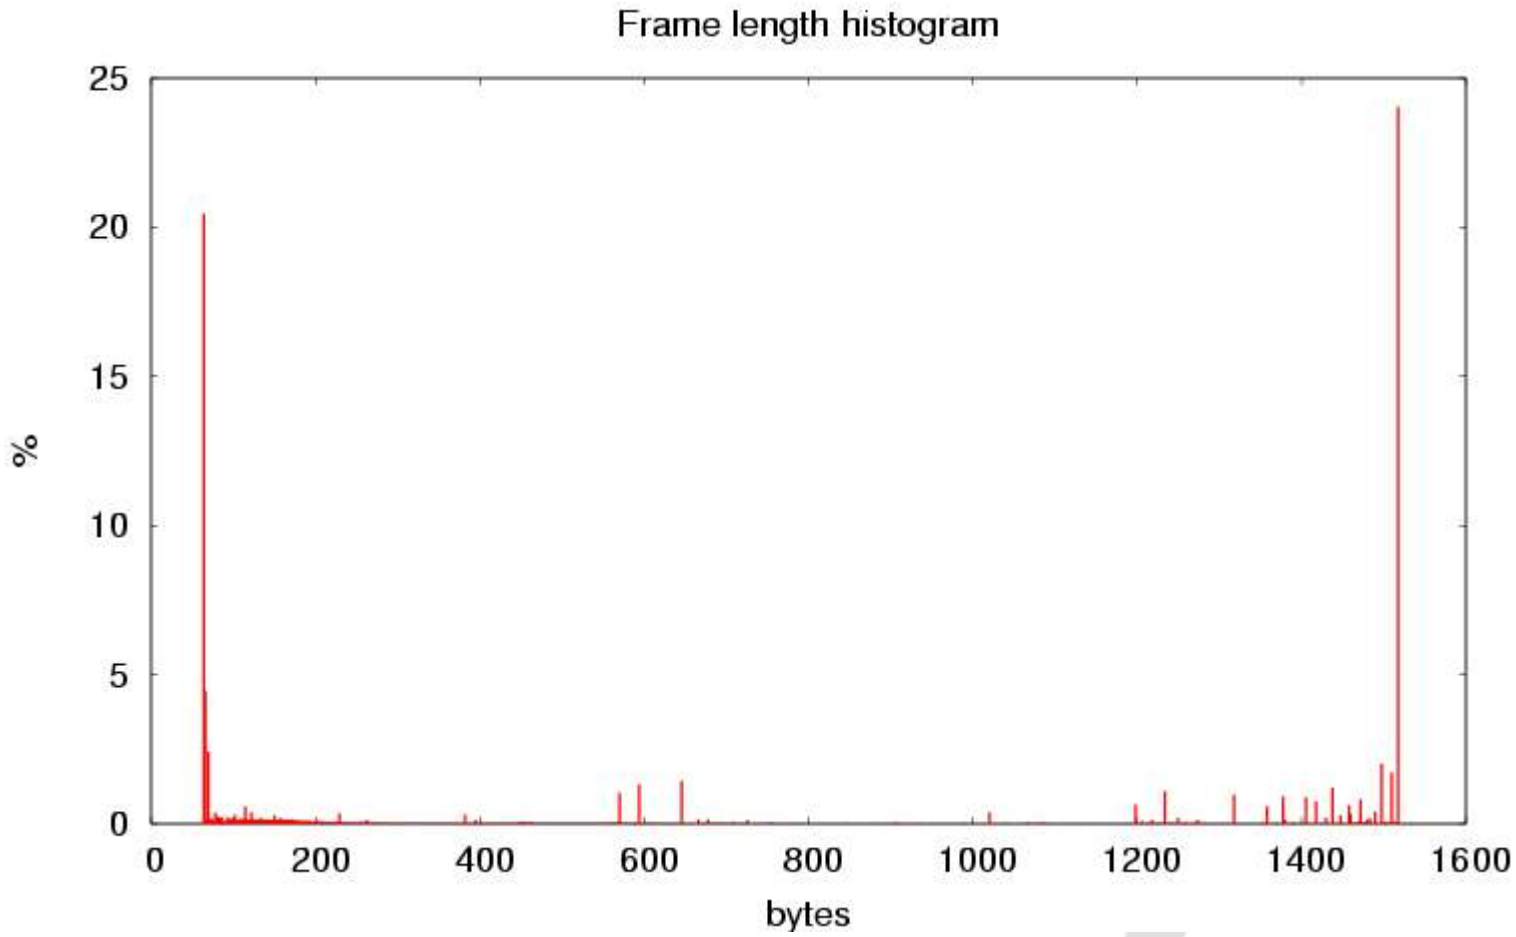

Spikes

- 24% of frames are 1518 bytes
- 27% are 64 - 70 bytes
- Load average: 18.2%
- If arrivals were independent, spike for 12.3 µs inter-arrival time would be 24% x 18% = 4.3%
- Observed: 13.8%

Line idle time

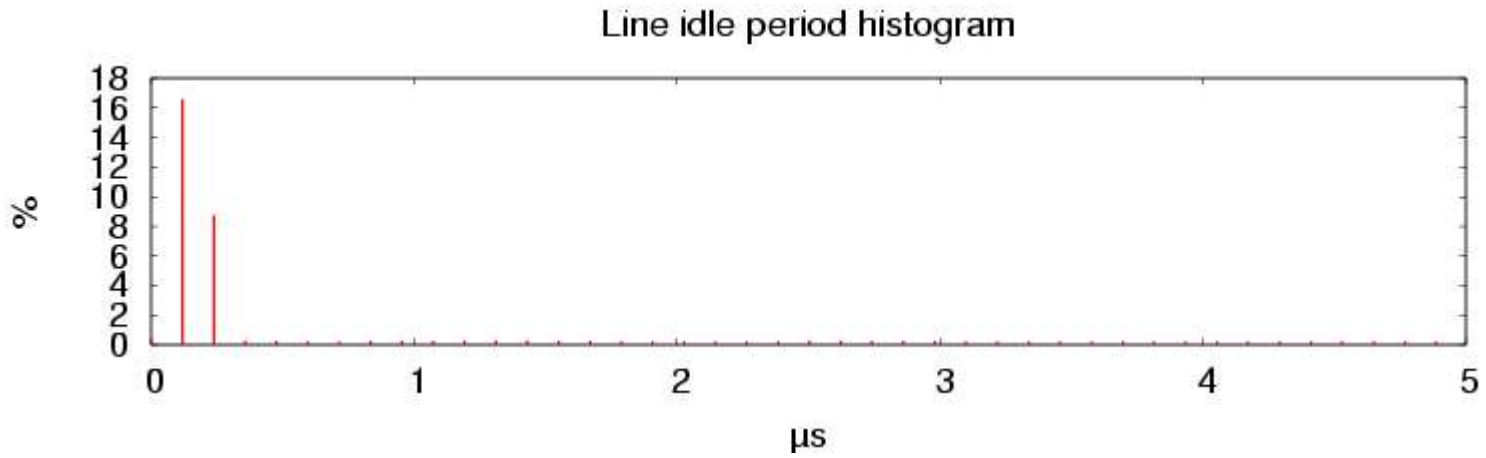

- Very frequent values < 0.3 µs must be line card sending frames as close to each other as it can.

UNINETT

Runs of packets (log scale)

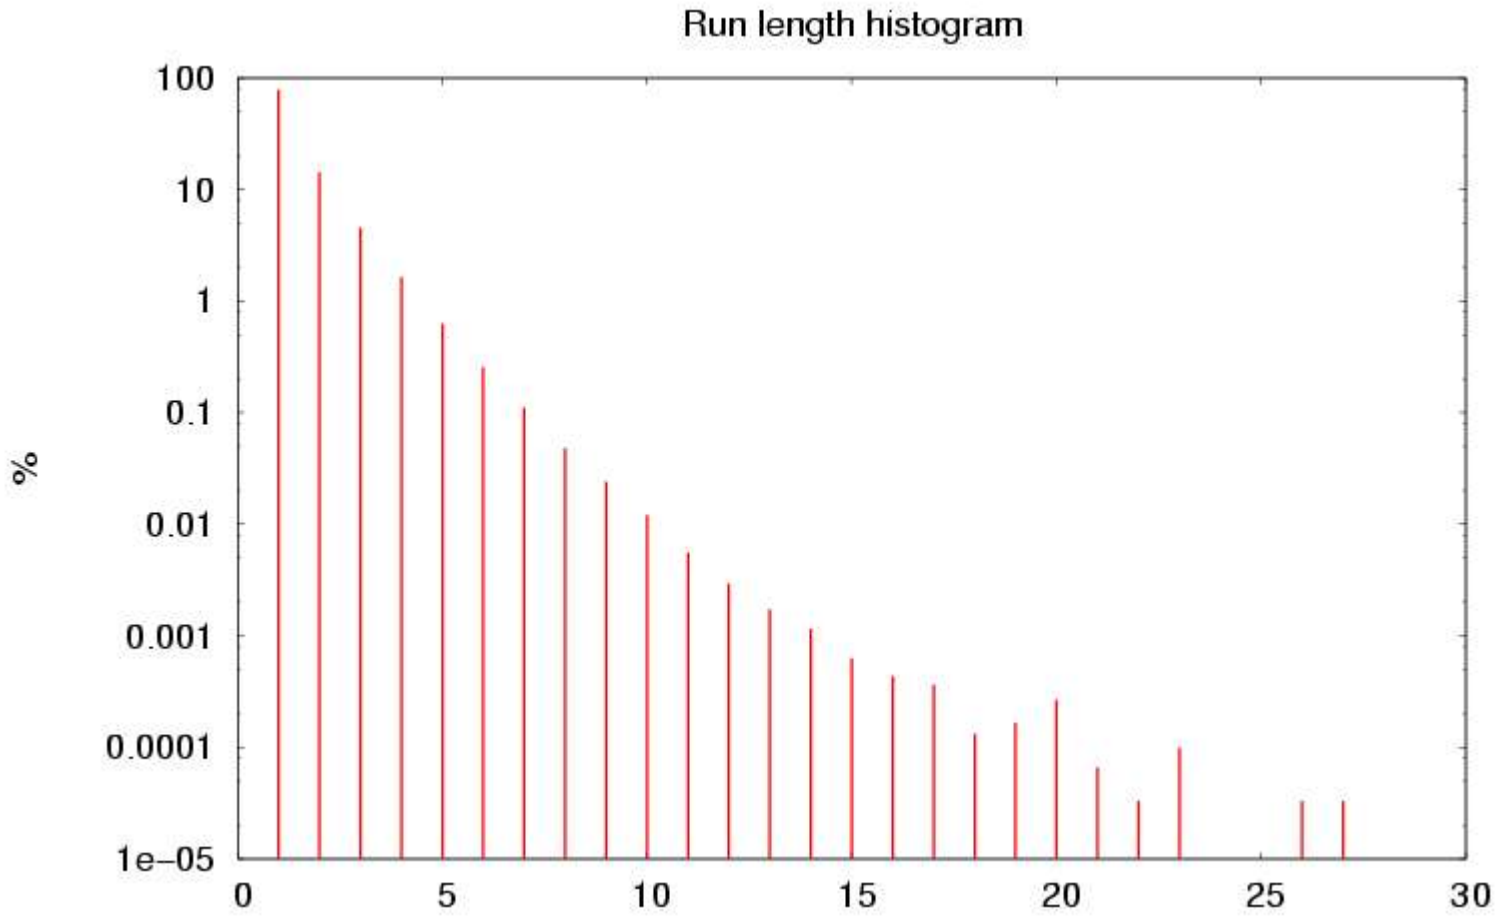

Runs of packets

- May indicate queuing
- 25% of frames are not first in a run
- Longest run in 2 minute period:

  27 frames

  34033 bytes

  0.27 ms
- 99.99% of runs shorter than 10 frames

UNINETT

Run durations (log scale)

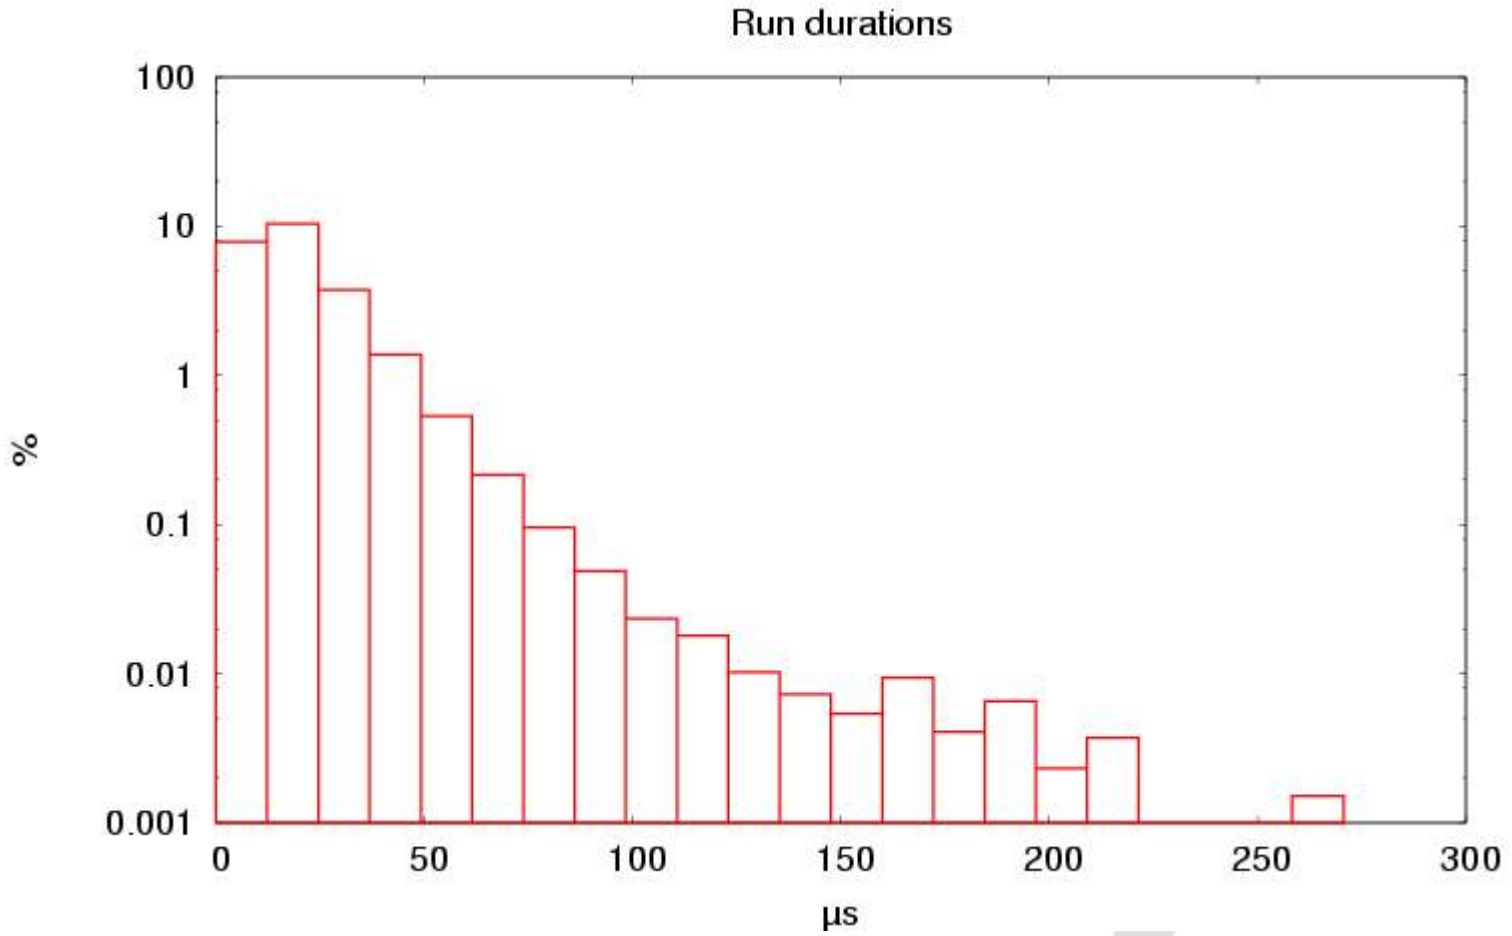

Conclusions

- Hardware probes reveal a lot
- Monitoring run lengths may be particularly useful.
  - Runs/queues may be long enough that we notice, but not yet so long that users notice

Thank you for your attention

- http://domen.uninett.no/~jk/micromeasurements/

- Email: Jon.Kare.Hellan@uninett.no