Measuring Network Traffic

Remco van de Meent

r.vandemeent@utwente.nl

University of Twente

Outline

- Introduction
- Contribution and Approach
- Measurement Setup
- Measurements
- Selection of Results
- Conclusions Future Plans

Introduction

Background Ph.D study combining telematics and mathematics; industry-funded Internet-NG and M2C projects; supervisors: Aiko Pras, Michel Mandjes, Hans van den Berg, Bart Nieuwenhuis

Introduction

Background Ph.D study combining telematics and mathematics; industry-funded Internet-NG and M2C projects; supervisors: Aiko Pras, Michel Mandjes, Hans van den Berg, Bart Nieuwenhuis

Rationale QoS in networks may be achieved by means of overprovisioning $\rightarrow$ need to understand traffic characteristics

Introduction

Background Ph.D study combining telematics and mathematics; industry-funded Internet-NG and M2C projects; supervisors: Aiko Pras, Michel Mandjes, Hans van den Berg, Bart Nieuwenhuis

Rationale QoS in networks may be achieved by means of overprovisioning $\rightarrow$ need to understand traffic characteristics

Goal Intelligent overprovisioning of network links

Introduction — Related Work

Numerous measurement “projects”, e.g.:

Research IRTF (imrg, e2e, aaaarch, nmrg?), COST (242, 257, 279), IST-SCAMPI

Introduction — Related Work

Numerous measurement “projects”, e.g.:

Research IRTF (imrg, e2e, aaaarch, nmrg?), COST (242, 257, 279), IST-SCAMPI

Operations RIPE (ttm), CAIDA, NLANR, MCI-Worldcom, Sprint ATL

Introduction — Related Work

Numerous measurement “projects”, e.g.:

Research IRTF (imrg, e2e, aaaarch, nmrg?), COST (242, 257, 279), IST-SCAMPI

Operations RIPE (ttm), CAIDA, NLANR, MCI-Worldcom, Sprint ATL

Standardization IETF (aaa, rtfm, ipfix, psamp, rmonmib, ippm, bmwg), not so much ISO/ITU

Introduction — Related Work

Numerous measurement “projects”, e.g.:

Research IRTF (imrg, e2e, aaaarch, nmrg?), COST (242, 257, 279), IST-SCAMPI

Operations RIPE (ttm), CAIDA, NLANR, MCI-Worldcom, Sprint ATL

Standardization IETF (aaa, rtfm, ipfix, psamp, rmonmib, ippm, bmwg), not so much ISO/ITU

European NoEs MAUI $\rightarrow$ E-NEXT

Contribution

intelligent overprovisioning

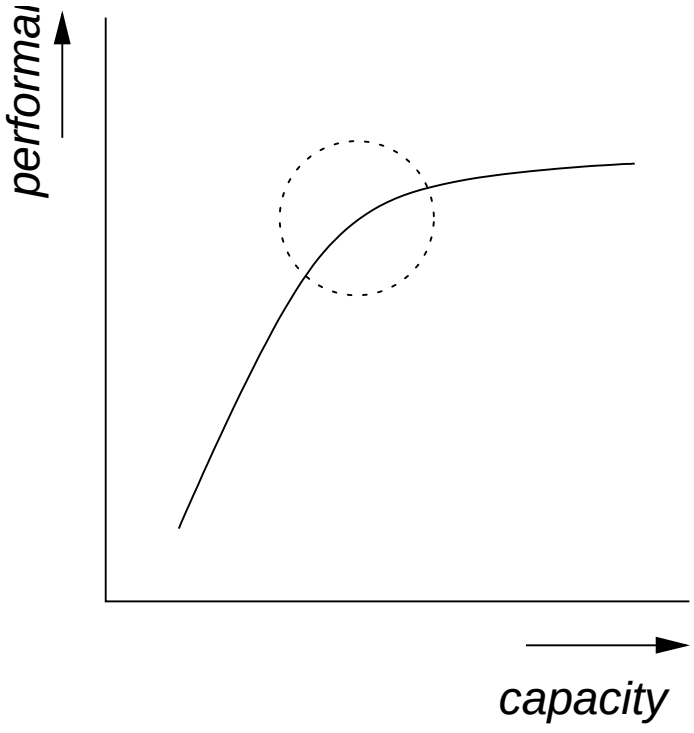

Contribution

intelligent overprovisioning

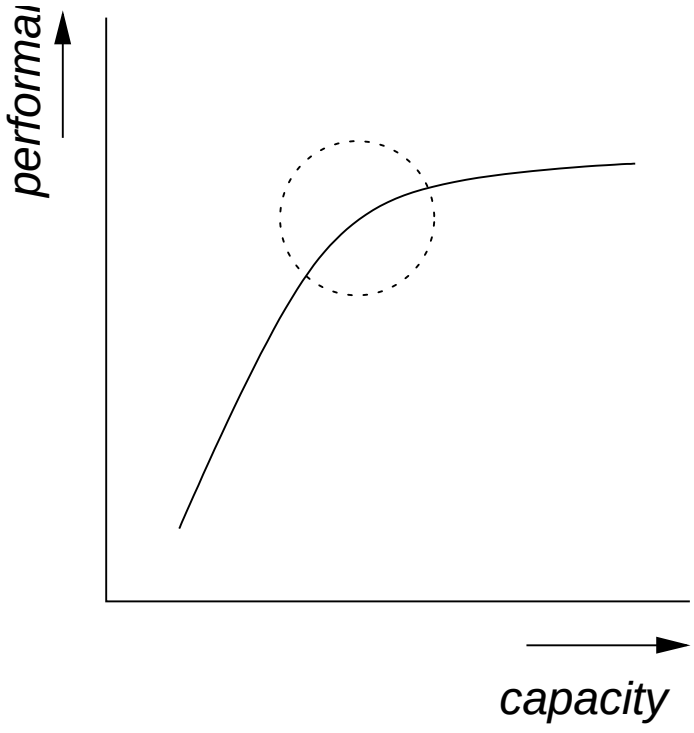

- course measurements
- what must be measured

Our Approach

- course measurements (e.g. MRTG)
  fine-grained measurements (tcpdump + custom analysis tools)

Our Approach

- course measurements (e.g. MRTG)
  fine-grained measurements (tcpdump + custom analysis tools)
- find relations between course- and fine-grained measurements

Our Approach

- course measurements (e.g. MRTG)
  fine-grained measurements (tcpdump + custom analysis tools)
- find relations between course- and fine-grained measurements
- verification using measurements on other networks

Our Approach

- course measurements (e.g. MRTG)
  fine-grained measurements (tcpdump + custom analysis tools)
- find relations between course- and fine-grained measurements
- verification using measurements on other networks
- mathematical models

Our Approach

- course measurements (e.g. MRTG)
  fine-grained measurements (tcpdump + custom analysis tools)
- find relations between course- and fine-grained measurements
- verification using measurements on other networks
- mathematical models
- “go back to Start”

Measurement Setup

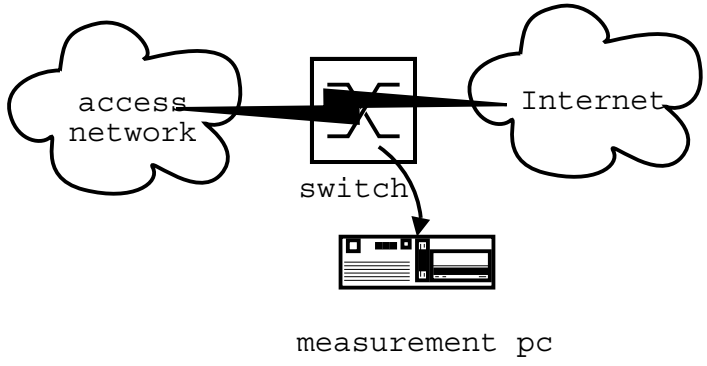

Measurement Setup

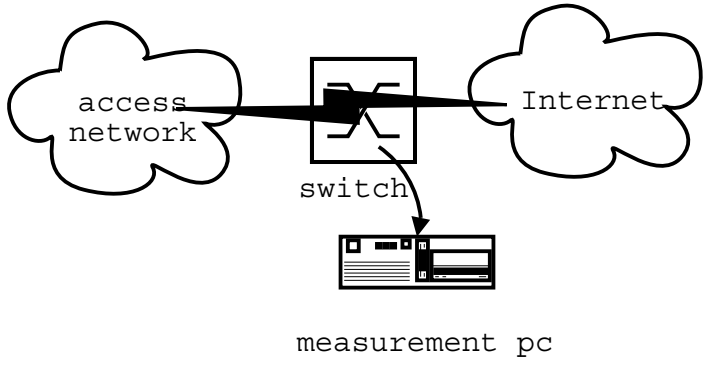

Measurement PC:

- Pentium-III, 1 GHz
- 512MB RAM
- 64-bit PCI
- standard Linux 2.4 kernel
- Gigabit networking.

Measurement Setup

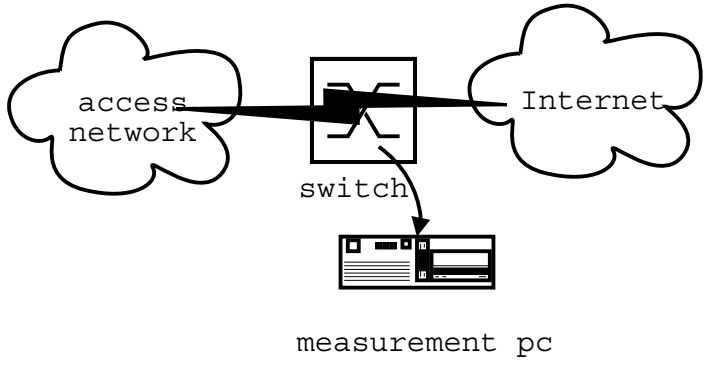

Measurement PC:

- Pentium-III, 1 GHz
- 512MB RAM
- 64-bit PCI
- standard Linux 2.4 kernel
- Gigabit networking.

captures packet headers using tcpdump, anonymization through tcpdpriv

Measurements (1)

- store packet header traces (first 64 octets, includes everything up to tcp/udp layer)
- 15 minutes each
- multiple times a day
- 7 days per week
- since 2002, different locations, SURFnet “backbone”

Measurements (2)

- ˜ 2000 users $\times$ 100 Mbit/s, 300 Mbit/s uplink
- ˜ 200 users $\times$ 100 Mbit/s, 1 Gbit/s uplink
- ˜ 1000 users $\times$ 10–100 Mbit/s, 1 Gbit/s uplink
- (not SURFnet) ADSL: hundreds of users, 0.5 – 1 Mbit/s, 155 Mbit/s uplink, multiple locations

Measurements (2)

- 10 – 500 packets per 10 milliseconds
- up to 4 GB of disk space per 15 minutes

So... what did we do?

Packet traces give detailed information on, e.g.,

- throughput on arbitrary time scales (tomorrow)
- burstiness and dimensioning (peak / mean) (submitted, sneak preview here)
- other characteristics that help to understand traffic (visualization tools)
- application recognition
- arrival process
- multi-level view on traffic (packets, flows, sessions)

What did we not do

In our research we do not perform measurements based on multiple, correlated metering points, so:

- no end-to-end information
- no available bandwidth estimation
- no information on delay
- no packet-loss information
- ...

Dimensioning Rules

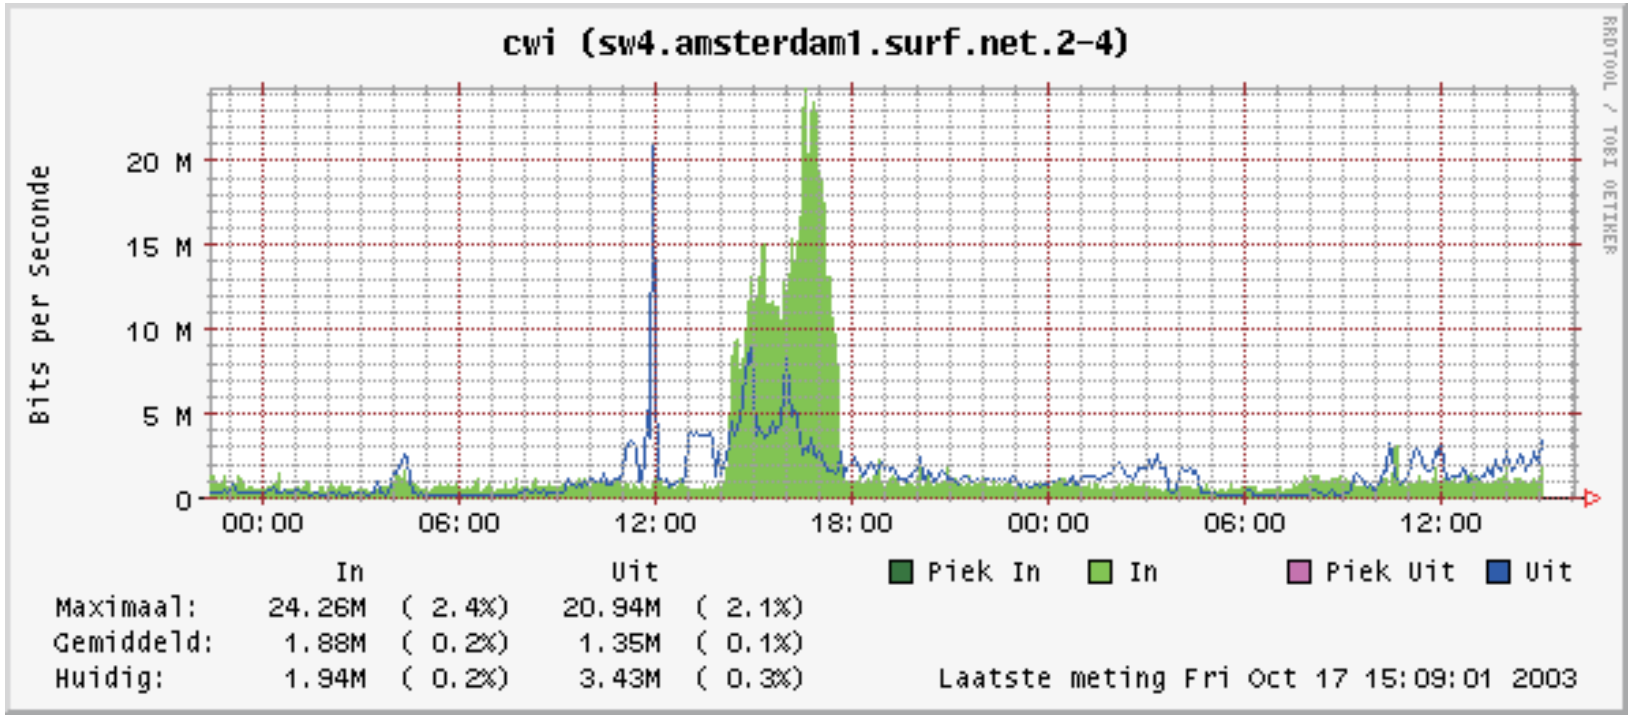

Dimensioning Rules

- 5 minute average throughput m
- 99th percentile of 1 second average p
- $\rightarrow$ hundreds of (m, p)-tuples

Dimensioning Rules

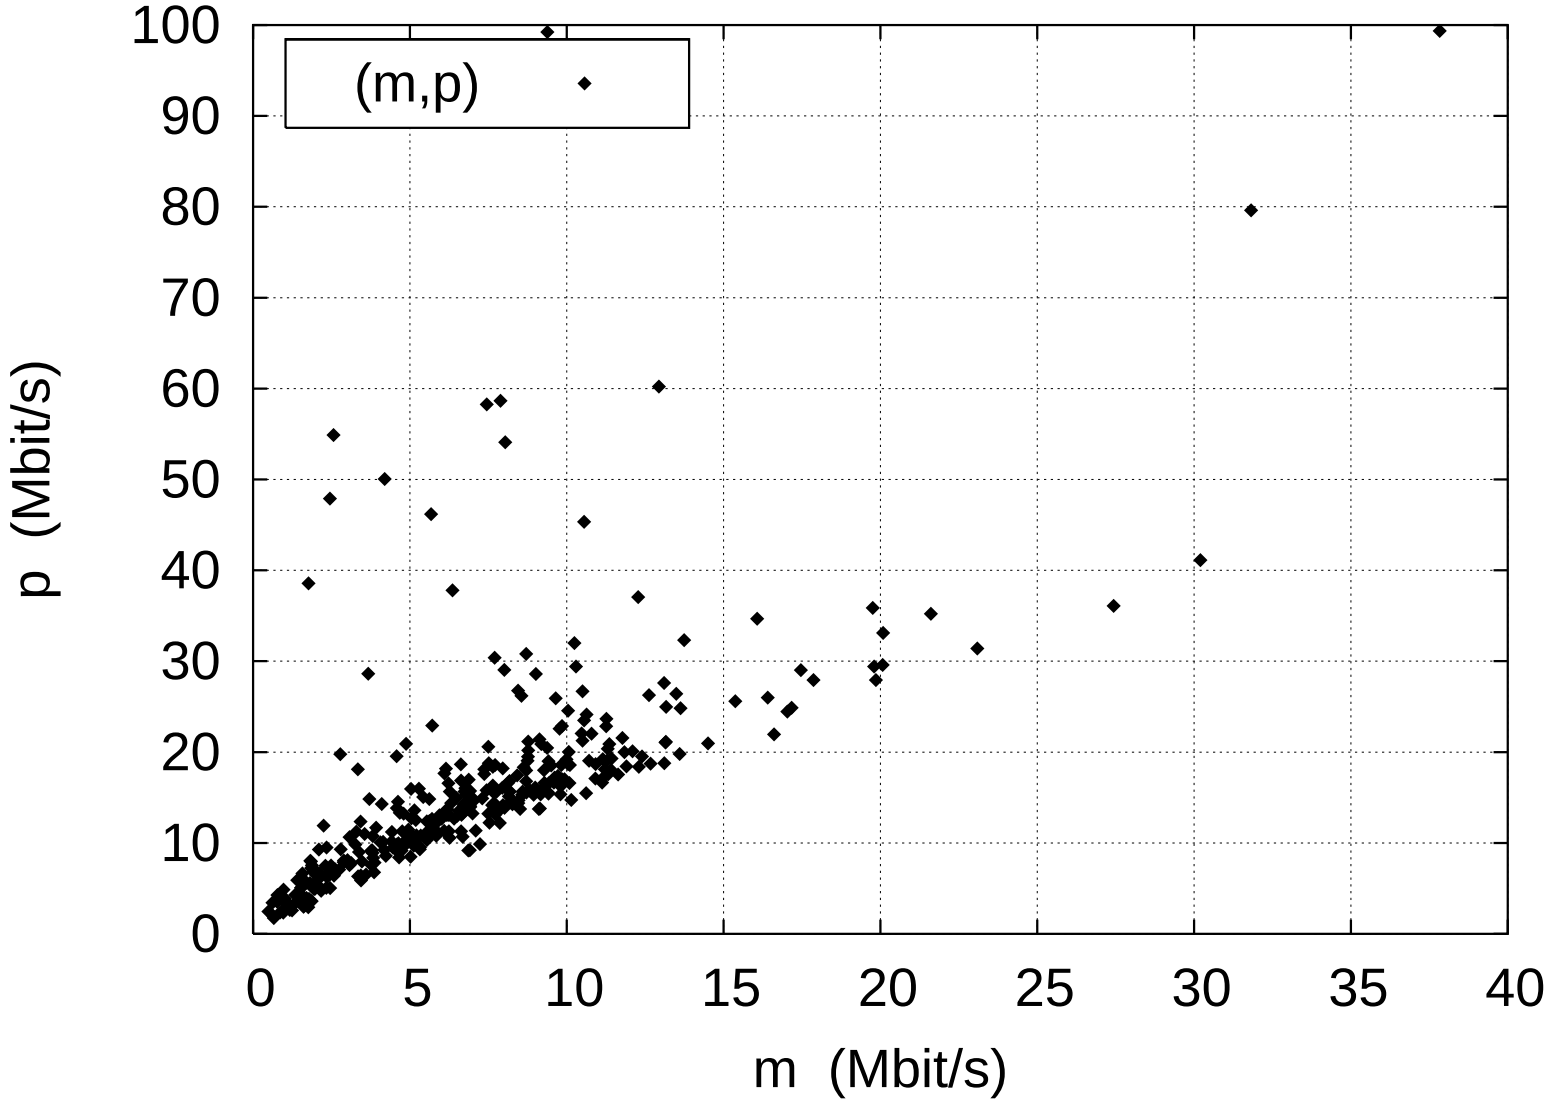

Dimensioning Rules

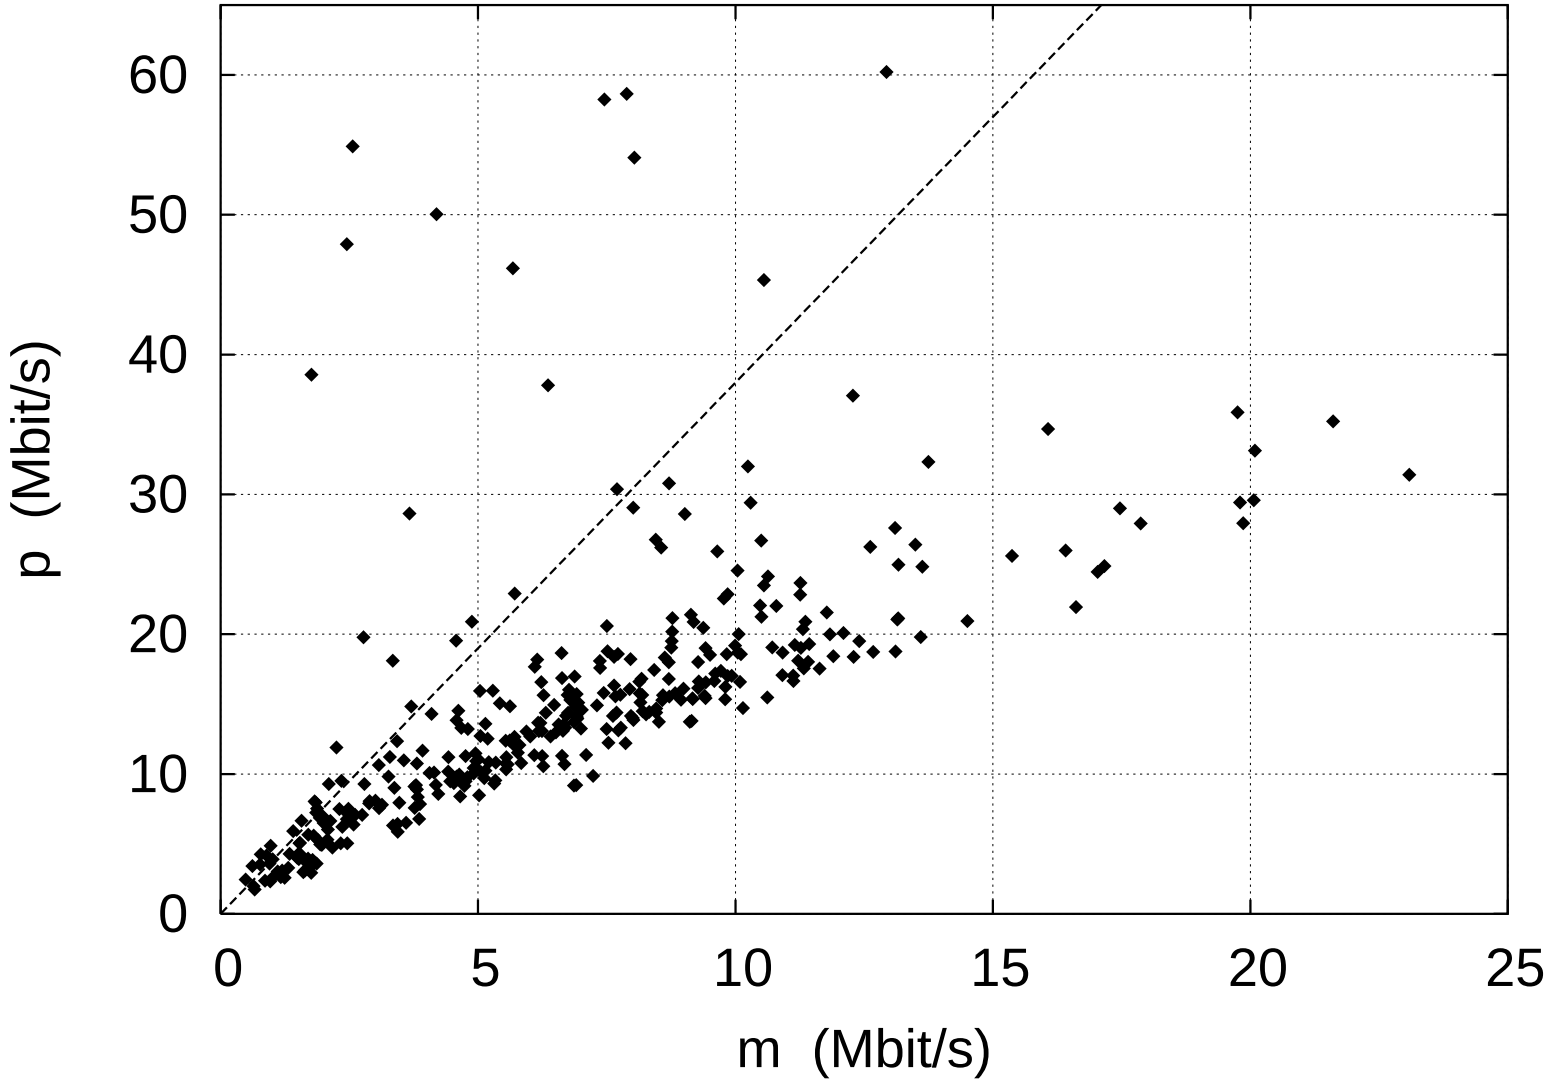

Dimensioning Rules

- we tried multiple rules (lines, curves), all give better (tens of percents) results than original “50% overprovisioning required guideline”
- based on course-grained measurements; fine-grained measurements every now and then, to finetune parameter settings

Visualisation Tools

- works on flow-level
  5-tuple: ip src/dst, proto, tp src/dst
- supports different types of statistics (extendible)
- `http://m2c-a.cs.utwente.nl/bsc-visual/`

Visualisation Tools

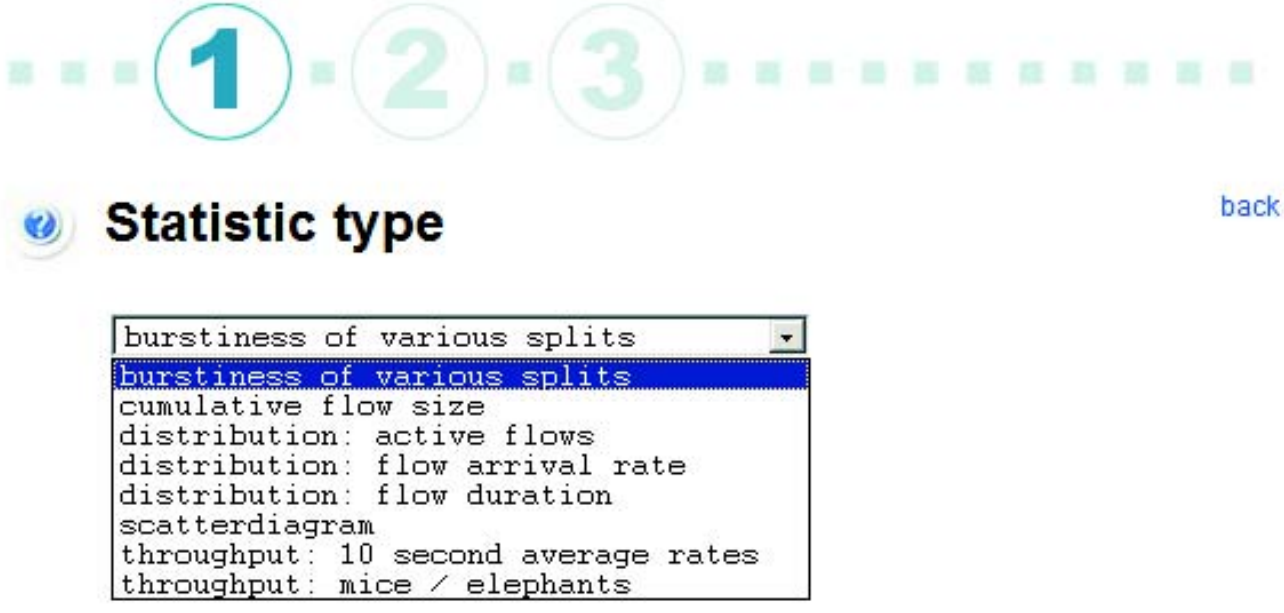

Visualisation Tools

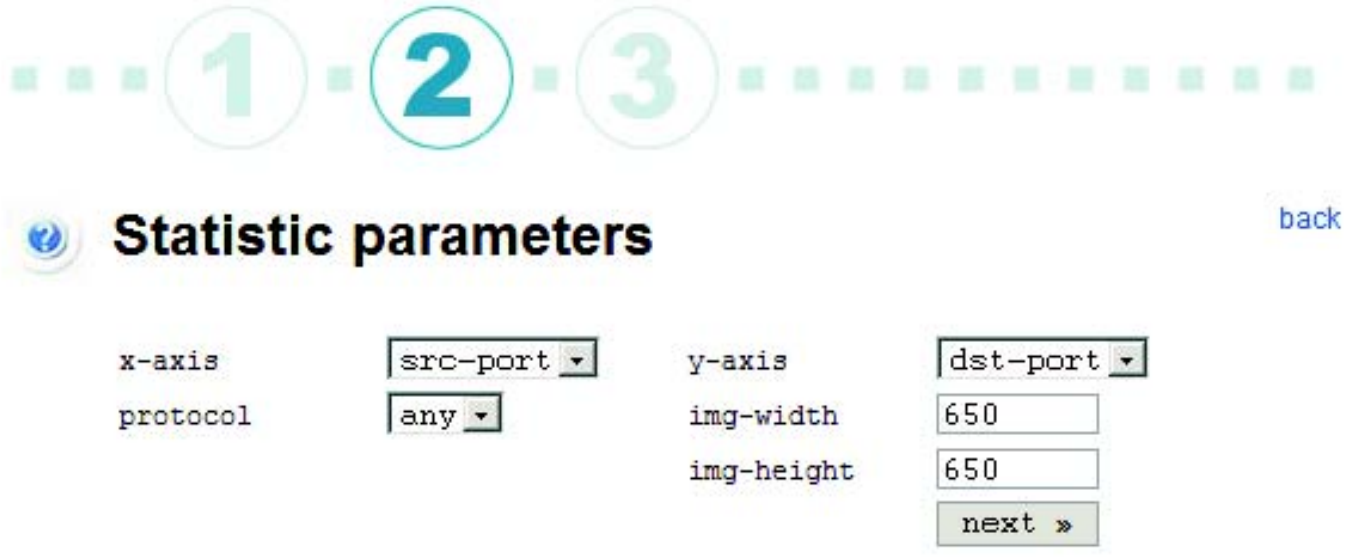

Visualisation Tools

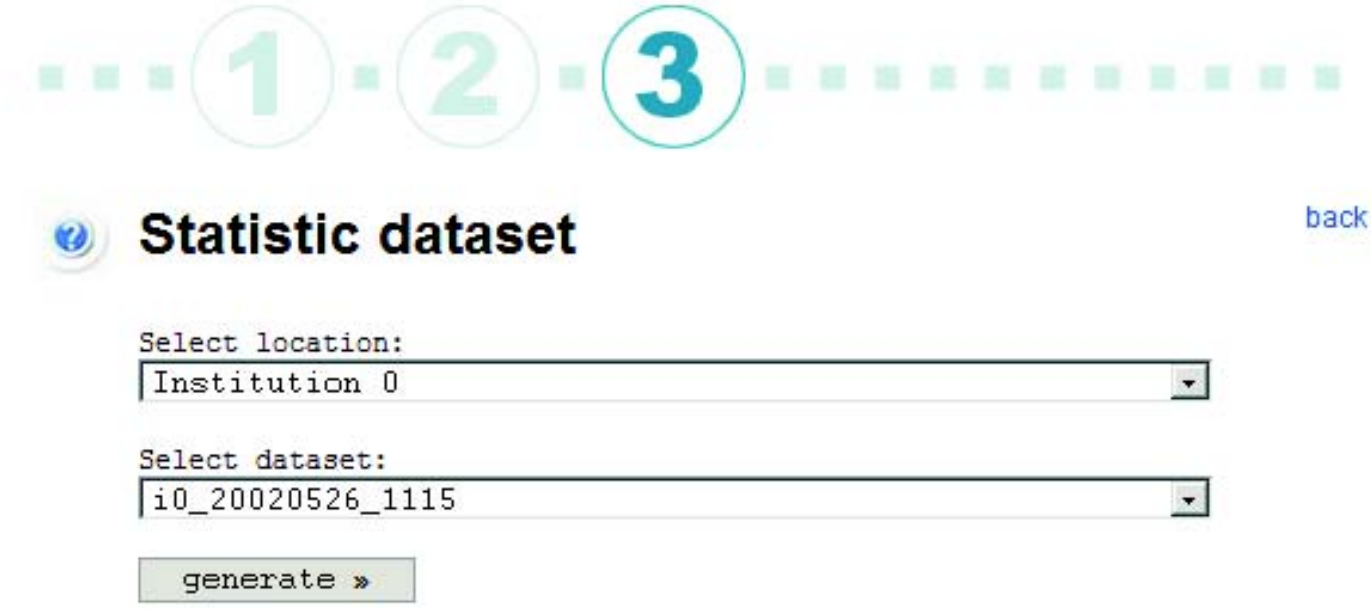

Visualisation Tools


```
Timing    : Image creation took 6.639 seconds...
Showing   : src_port vs dst_port
Dataset   : i0_20020526_1115
Condition: protocol any
```

Application Recognition

- port lists (IANA, other...)

Application Recognition

- port lists (IANA, other...)
- relating flows

Application Recognition

- port lists (IANA, other...)
- relating flows
- no payload inspection

Application Recognition

- port lists (IANA, other...)
- relating flows
- no payload inspection
- *Campusnet:*
  - > 90% of local traffic is windows networking, little email or web, hardly any p2p
  - traffic from/to Internet: still tens of percents unknown

Packets / Flows

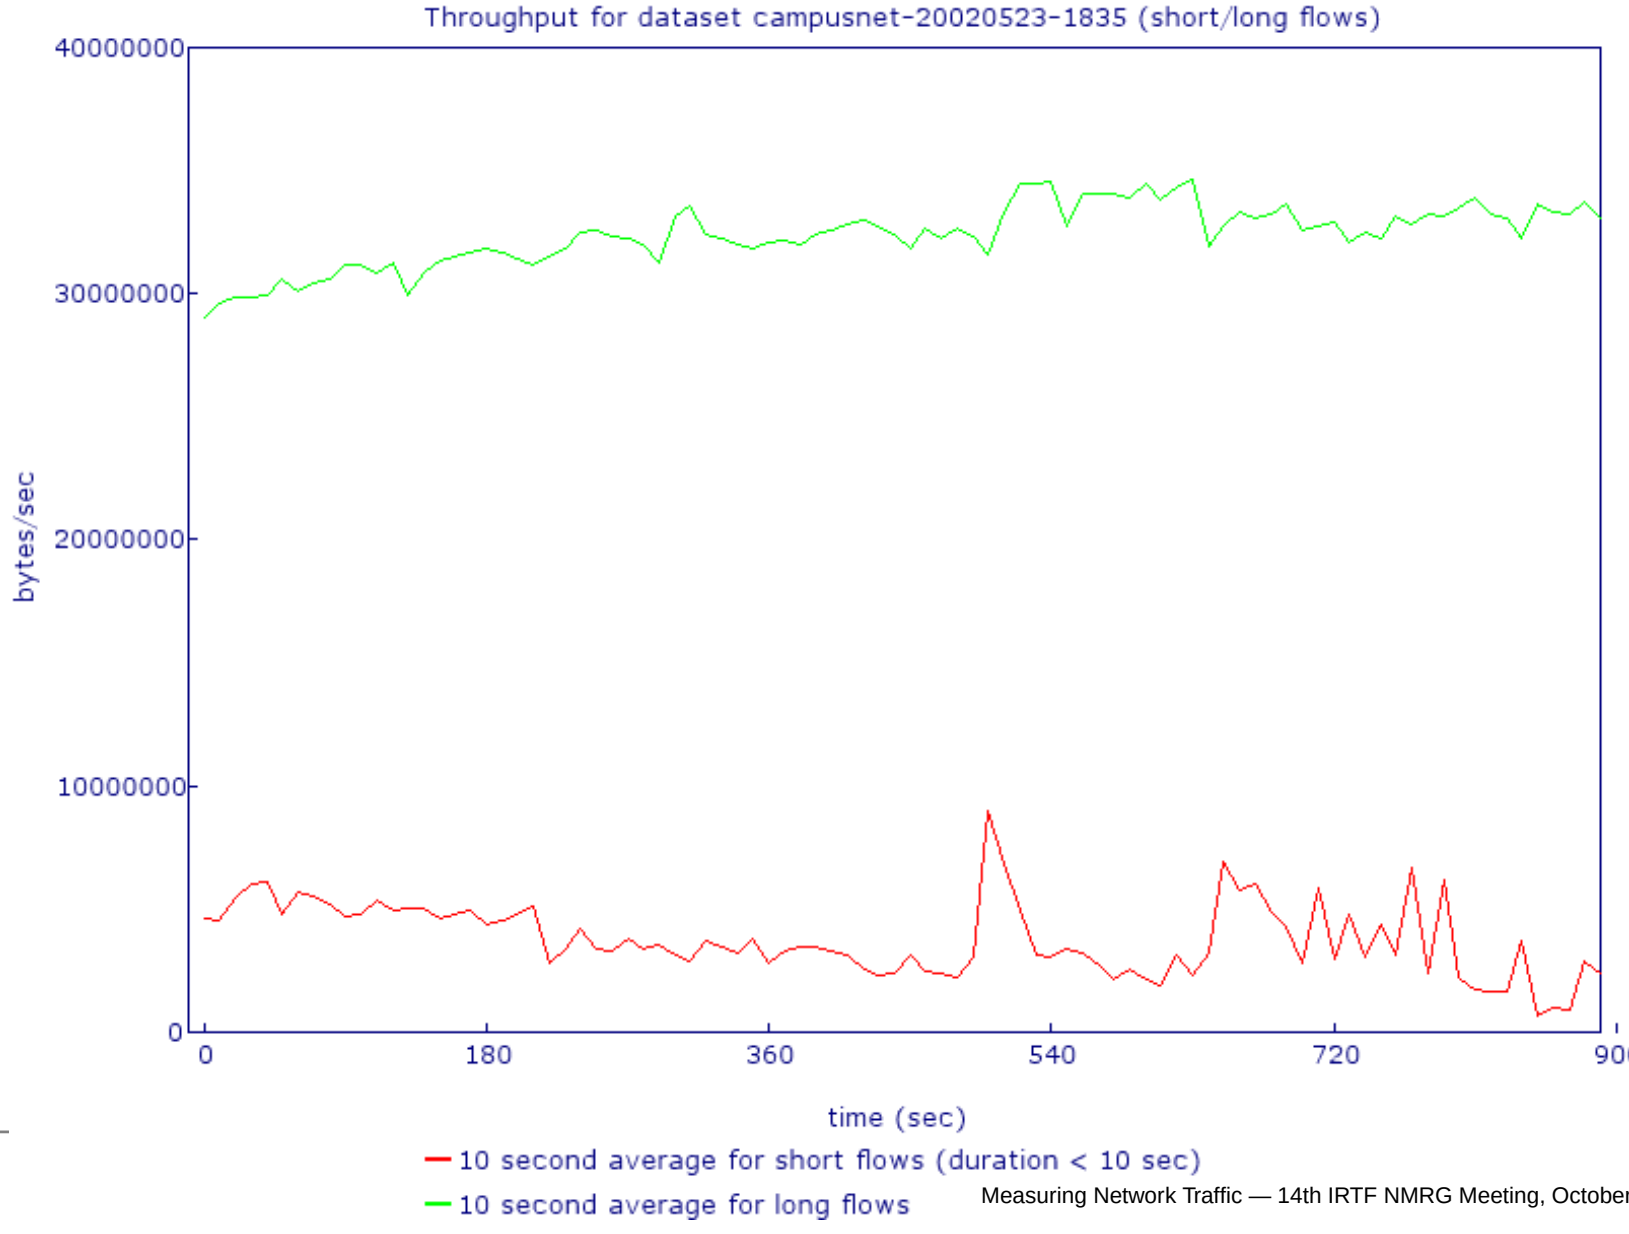

Conclusions

- detailed measurements with moderate hardware are feasible

Conclusions

- detailed measurements with moderate hardware are feasible
- course measurements together with occasional detailed measurements seem to give reasonable idea of peak load

Conclusions

- detailed measurements with moderate hardware are feasible
- course measurements together with occasional detailed measurements seem to give reasonable idea of peak load
- (increasing) amount of unidentified traffic is worrying

Conclusions

- detailed measurements with moderate hardware are feasible
- course measurements together with occasional detailed measurements seem to give reasonable idea of peak load
- (increasing) amount of unidentified traffic is worrying
- data and tools available for public use soon

Future plans

- more measurements, other networks

Future plans

- more measurements, other networks
- repository of measurement data
  - anonymous
  - possibly distributed
  - easy access via web

Future plans

- more measurements, other networks
- repository of measurement data
  - anonymous
  - possibly distributed
  - easy access via web
- network dimensioning rules based on course measurements and “out-of-band” information

Future plans

- more measurements, other networks
- repository of measurement data
  - anonymous
  - possibly distributed
  - easy access via web
- network dimensioning rules based on course measurements and “out-of-band” information
- simulations, mathematical foundations