Bandwidth estimations

for Test Traffic Measurements project

Henk Uijterwaal
Mark Santcroos

RIPE NCC New Projects Group
BEst2003, December 9

RIPE NCC

- Regional Internet Registry for Europe
- Based in Amsterdam, The Netherlands
- New Projects group

Test Traffic Measurements

- Worldwide network of measurement nodes (however focussed on Europe)
- Active end-to-end measurements
- Measuring performance of Internet infrastructure:
  - Packet Delay
  - Packet Loss
  - Jitter (Packet Delay Variation)
  - Routing vectors (traceroute)
  - DNS servers (root, gtld's & cctld's)
- GPS synchronized

Test Traffic Measurements (2)

- Typical hosts of Testbox:
  - ISP's
  - Internet Exchanges
  - Universities
  - Network research institutions
- Not doing measurements for the sake of doing measurements
- Provide (end) users of testboxes with understandable data
- TTM -> Information Services

Bandwidth measurements

- Want to add bandwidth measurements to our list of measurements:
  - TB hosts can check what they are paying for
  - TB hosts can plan network upgrades
- Want to measure:
  - Capacity bandwidth
  - Available bandwidth
- Use (and / or improve) existing tools, no need/time to write another one

BW requirements for TTM

- Non-intrusive for the network
- Reliable and consistent
- Not interfere with other measurements on the Testboxes
- Relatively fast

Traffic light

- Creating a green, orange and red light is simple of course
- Determining when to show what color is the problem
- For delay's we have “alarms”, for bandwidth we could do the same

What we did ...

- First idea (some time ago):
  - Write some wrappers around existing tools
  - Ran it in production network
  - Too many problems

- Second idea (more recently):
  - Test the tools in a controlled testbed
  - Very simple parameters
  - Results still not satisfying

Result

- Only more questions, but I'll come to that later ...

Test environment

- Two machines connected with cross-cable
- 100 Mbps full-duplex
- tt95: PII 333 MHz, 64 MB, 3Com NIC, FreeBSD 4.6
- tt96: PII 466 MHz, 256 MB, 3Com NIC, FreeBSD 4.6

Test environment (2)

- Leaping ahead of conclusions:
  - probably too slow

Throughput tools

- Netperf and Iperf
- Saturate the link with UDP and TCP

TYPE	FROM	TO	Payload (Mbps)	Ethernet (Mbps)
UDP	tt95	tt96	92.1	94.6
UDP	tt96	tt95	87.4	89.9
TCP	tt95	tt96	86.6	89.8
TCP	tt96	tt95	83.6	86.7

New tool – Capacity

Tool	From	To	Capacity(Mbps)
pathrate	tt95	tt96	92/93
pathrate	tt96	tt95	88/89

New tools – Available BW

(no cross-traffic!)

Tool	From	To	Capacity (Mbps)
pathload	tt95	tt96	100/111
pathload	tt96	tt95	96/97
pathchirp	tt95	tt96	67
pathchirp	tt96	tt95	85

CPU influence

- How does the CPU frequency have influence on the results?
- Results from all tools different in the two directions between tt95 and tt96
- Find out what is the real influencing factor!

Motivation

- Evaluation of tools showed there are more parameters than just bandwidth
- Doing Network Delay Quantification study
  - Try to find out how delay is build up on it's way
  - (Will also help understand network delay distribution for other TTM measurements)

DAG setup

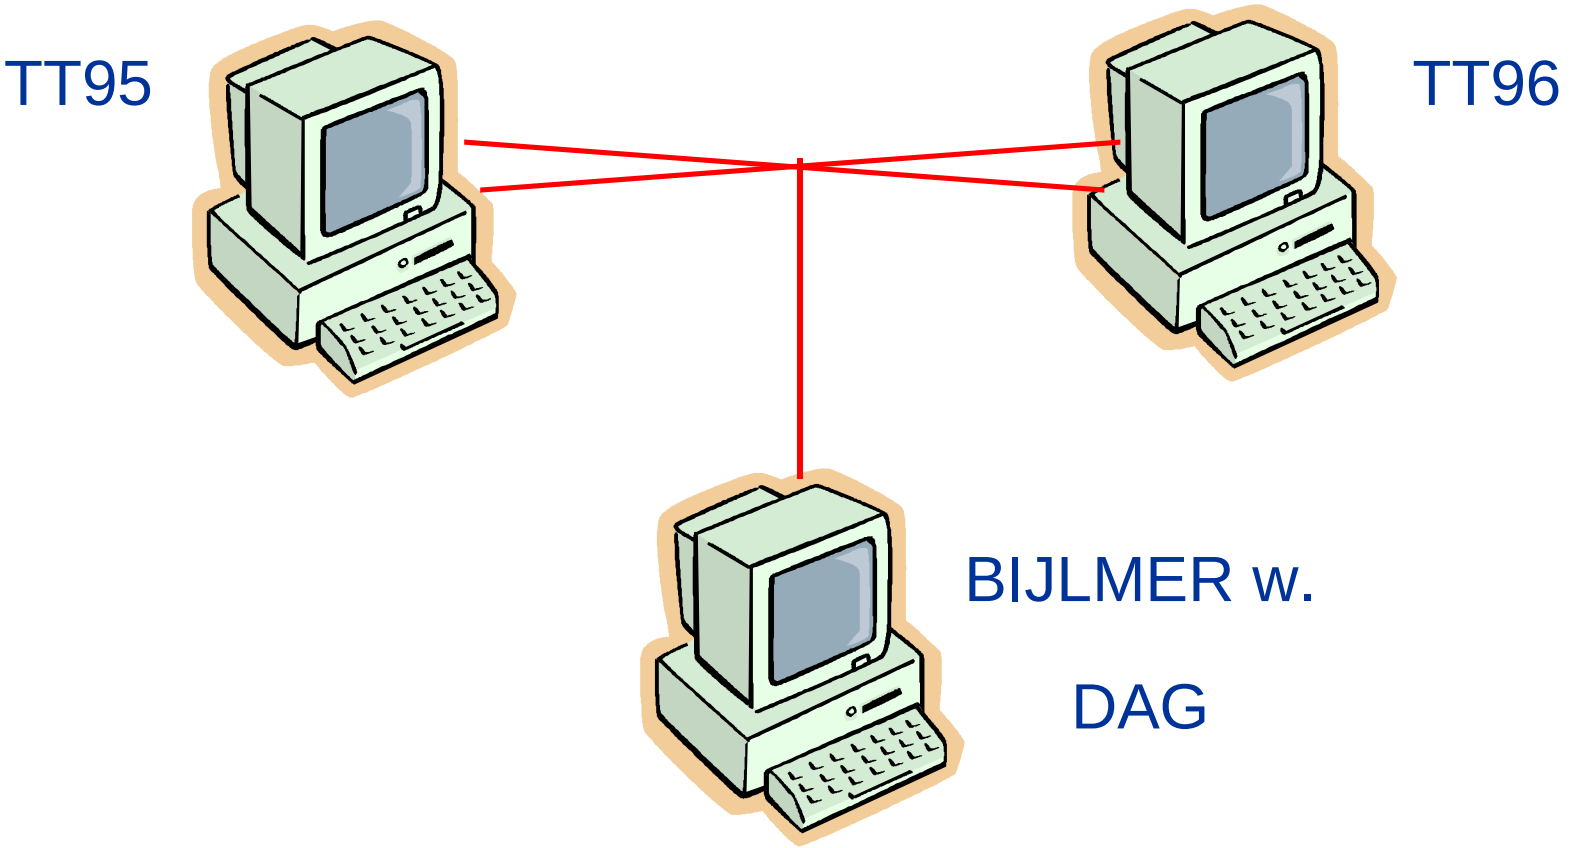

DAG Card


Dag v3.2E Waikato University
Design by Jed Martens
and Cameron Arthur 25/11/99
Dag 3.2E
Serial # 730
For use in laboratory
environment only
© University of Waikato 2000
GSI
GSI

DAG Card

- Xilinx FGPA equipped NIC
- “realtime” hardware timestamps
- GPS synchronized
- First 64 byte of frame

  http://dag.cs.waikato.ac.nz

Timestamps: where?

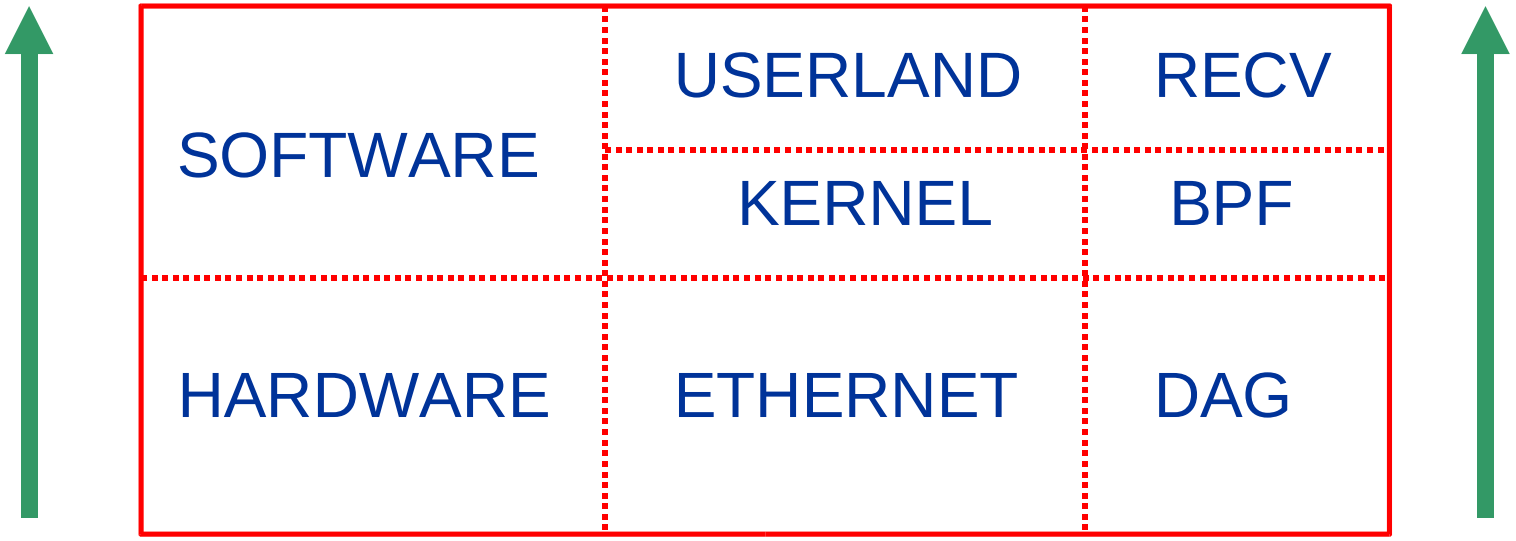

	Mean	Std dev	Median	95%	Max
DAG	95 us	22 us	89 us	148 us	225 us
BPF	194 us	23 us	190 us	249 us	332 us
RECV	405 us	43 us	401 us	468 us	911 us

Delay layout

- End-to-End delay is build up from several components
  - D1 = SEND <->DAG
  - D2 = DAG <-> BPF
  - D3 = BPF <-> RECV

	Mean	Std dev	Median	95%	Max
D1	95 us	22 us	89 us	148 us	225 us
D2	100 us	4 us	99 us	107 us	114 us
D3	211 us	29 us	207 us	249 us	600 us

Clock accuracy

- What is the error of various clocks involved?
- Trimble antenna accuracy: ~50ns to UTC
- System + NTP + GPS: ~1us to UTC
- DAG card + GPS: ~120ns to UTC
- Still small enough with today's delays

NIC influence

- How do NIC's work?
- Could be of influence
- Will try different NIC's

Status summary

- Promised this as a service a while ago
- Results still not satisfying enough
- We still want to deliver at some point
- Doing more research on all related matters

Plans

- More careful logging of all results
- Delay measurements at different rates
- Analyse real traffic to get (more) insight about available bandwidth properties
- Investigate ways to improve sending process

Questions and discussion