Bandwidth Studies

or Network Delay Quantification

Mark Santcroos

RIPE NCC New Projects Group

RIPE 46, Amsterdam, September 4, 2003

Overview

- *Bandwidth*
- Methods and tools
- Tools results
- DAG setup
- Network delay quantification
- Other factors
- Status summary
- Plans
- Questions and discussion

Why measure bandwidth?

- Plan capacity upgrades
- How much data can we get from here to there?
- Use as input parameter for TCP self tuning
- Application layer self tuning
- You can probably think of more …

Bandwidth: 2 (3) parameters

- "C": Capacity
  - Number of bytes that can be sent over a link per unit of time
- "A": Available Bandwidth
  - Number of bytes that be sent over a link considering the current cross traffic
- "C-A": Utilized Bandwidth
  - Number of bytes that is sent over the link right now

Overview

- Bandwidth
- *Methods and tools*
- Tools results
- DAG setup
- Network delay quantification
- Other factors
- Status summary
- Plans
- Questions and discussion

Variable Packet Size method

- Per-Hop capacity “C” estimation
- Use of IP TTL field
- Receive ICMP Time Exceeded
- Problems:
  - Level 2 store-and-forward devices
  - Variation in ICMP generation delays
  - More …

Pathrate

- Measures Capacity “C”
- Packet Pair/Train Dispersion
- Look at dispersion/interspacing at receiver
- Capacity = PacketSize / Dispersion
- More tools that use this technique, difference is in the statistical model

Pathload

- Measures Available Bandwidth “A”
- Self Loading Periodic Streams (SLoPS)
- Sending periodic packet streams
- Stream will show delay if rate is higher than available bandwidth
- Cross traffic is variable, need to be careful

Overview

- Bandwidth
- Methods and tools
- *Tools results*
- DAG setup
- Network delay quantification
- Other factors
- Status summary
- Plans
- Questions and discussion

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

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 tools – Capacity

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

New tools – Available BW

(on an unloaded network)

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)

Overview

- Bandwidth
- Methods and tools
- Tools results
- *DAG setup*
- Network delay quantification
- Other factors
- Status summary
- Plans
- Questions and discussion

DAG setup

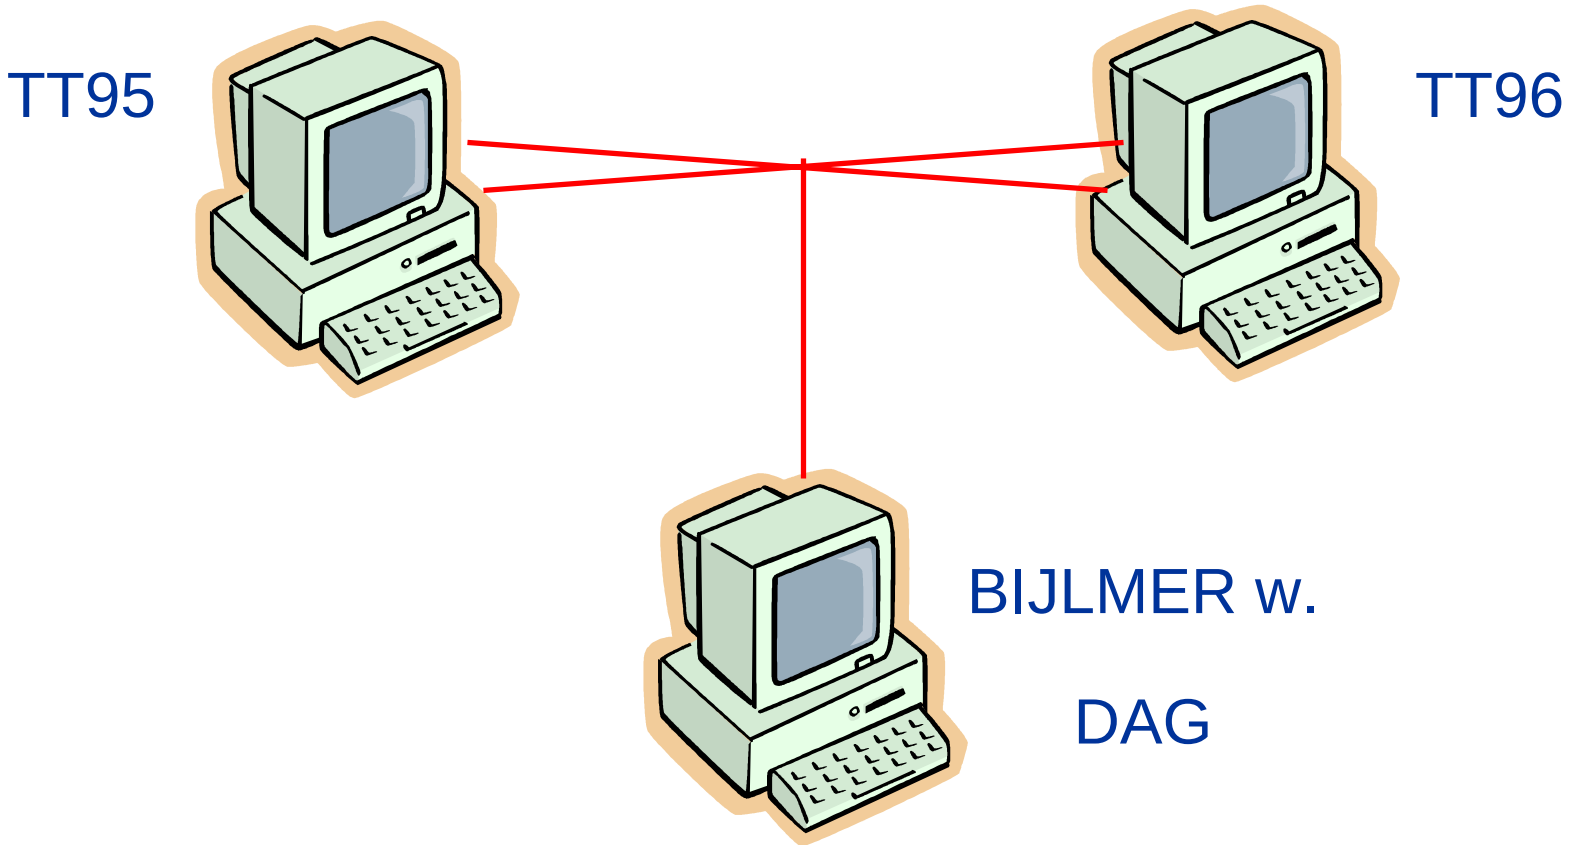

DAG Card


Dag v3.2E Waikato University
Design by Jed Martens
and Cameron Arthur 25/11/99
Dag 3.2E
Serial # 730
Not FCC or UL tested
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

Overview

- Bandwidth
- Methods and tools
- Tools results
- DAG setup
- *Network delay quantification*
- Other factors
- Status summary
- Plans
- Questions and discussion

Timestamps: where?

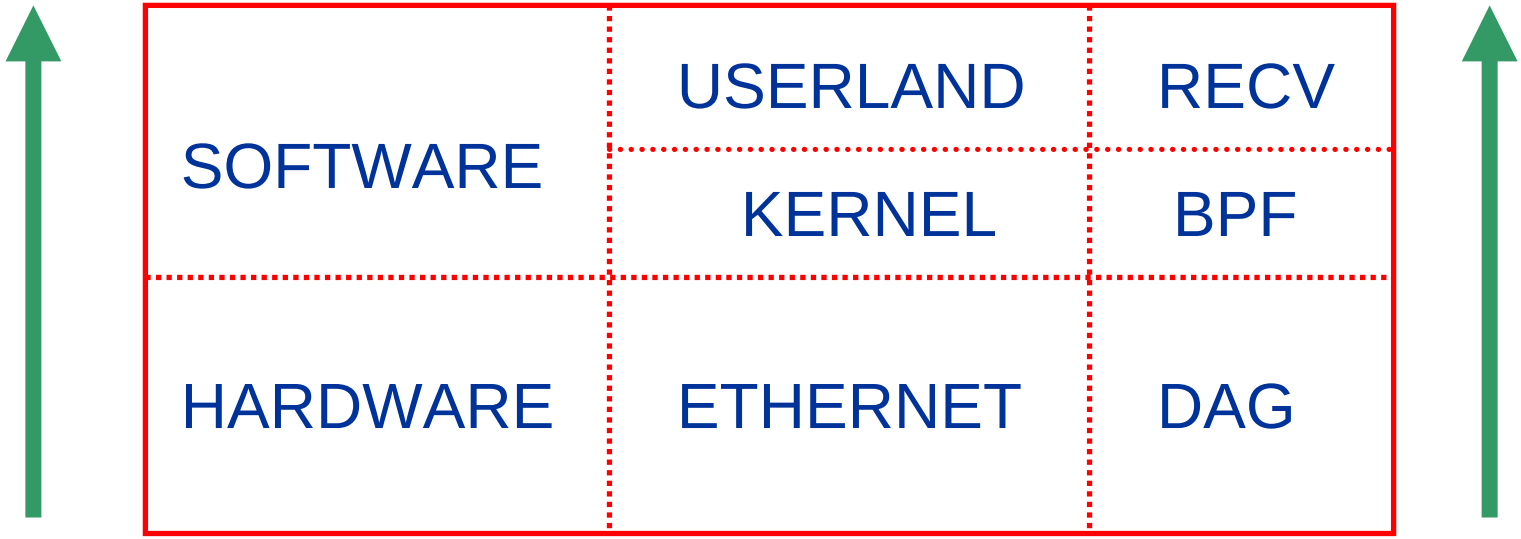

	Mean	Std dev	Median	95%	Max
DAG	95 us	22 us	89 us	148 us	225 us
BPF	194 us	23 us	190 us	249 us	332 us
RECV	405 us	43 us	401 us	468 us	911 us

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

Overview

- Bandwidth
- Methods and tools
- Tools results
- DAG setup
- Network delay quantification
- *Other factors*
- Status summary
- Plans
- Questions and discussion

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

Overview

- Bandwidth
- Methods and tools
- Tools results
- DAG setup
- Network delay quantification
- Other factors
- *Status summary*
- *Plans*
- *Questions and discussion*

Status summary

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

Plans

- Continue this study
- Write up in a paper
- At the same time start looking at bandwidth measurements for IPv6

Questions and discussion