ATM Virtual Path Tunneling Experiences

A. Schindler, J. Venema, C.T.A.M. de Laat and V. Reijs

University of Geneva

Faculty of Physics and Astronomy

SURFnet bv

Universiteit Utrecht

SVC tunnel principle

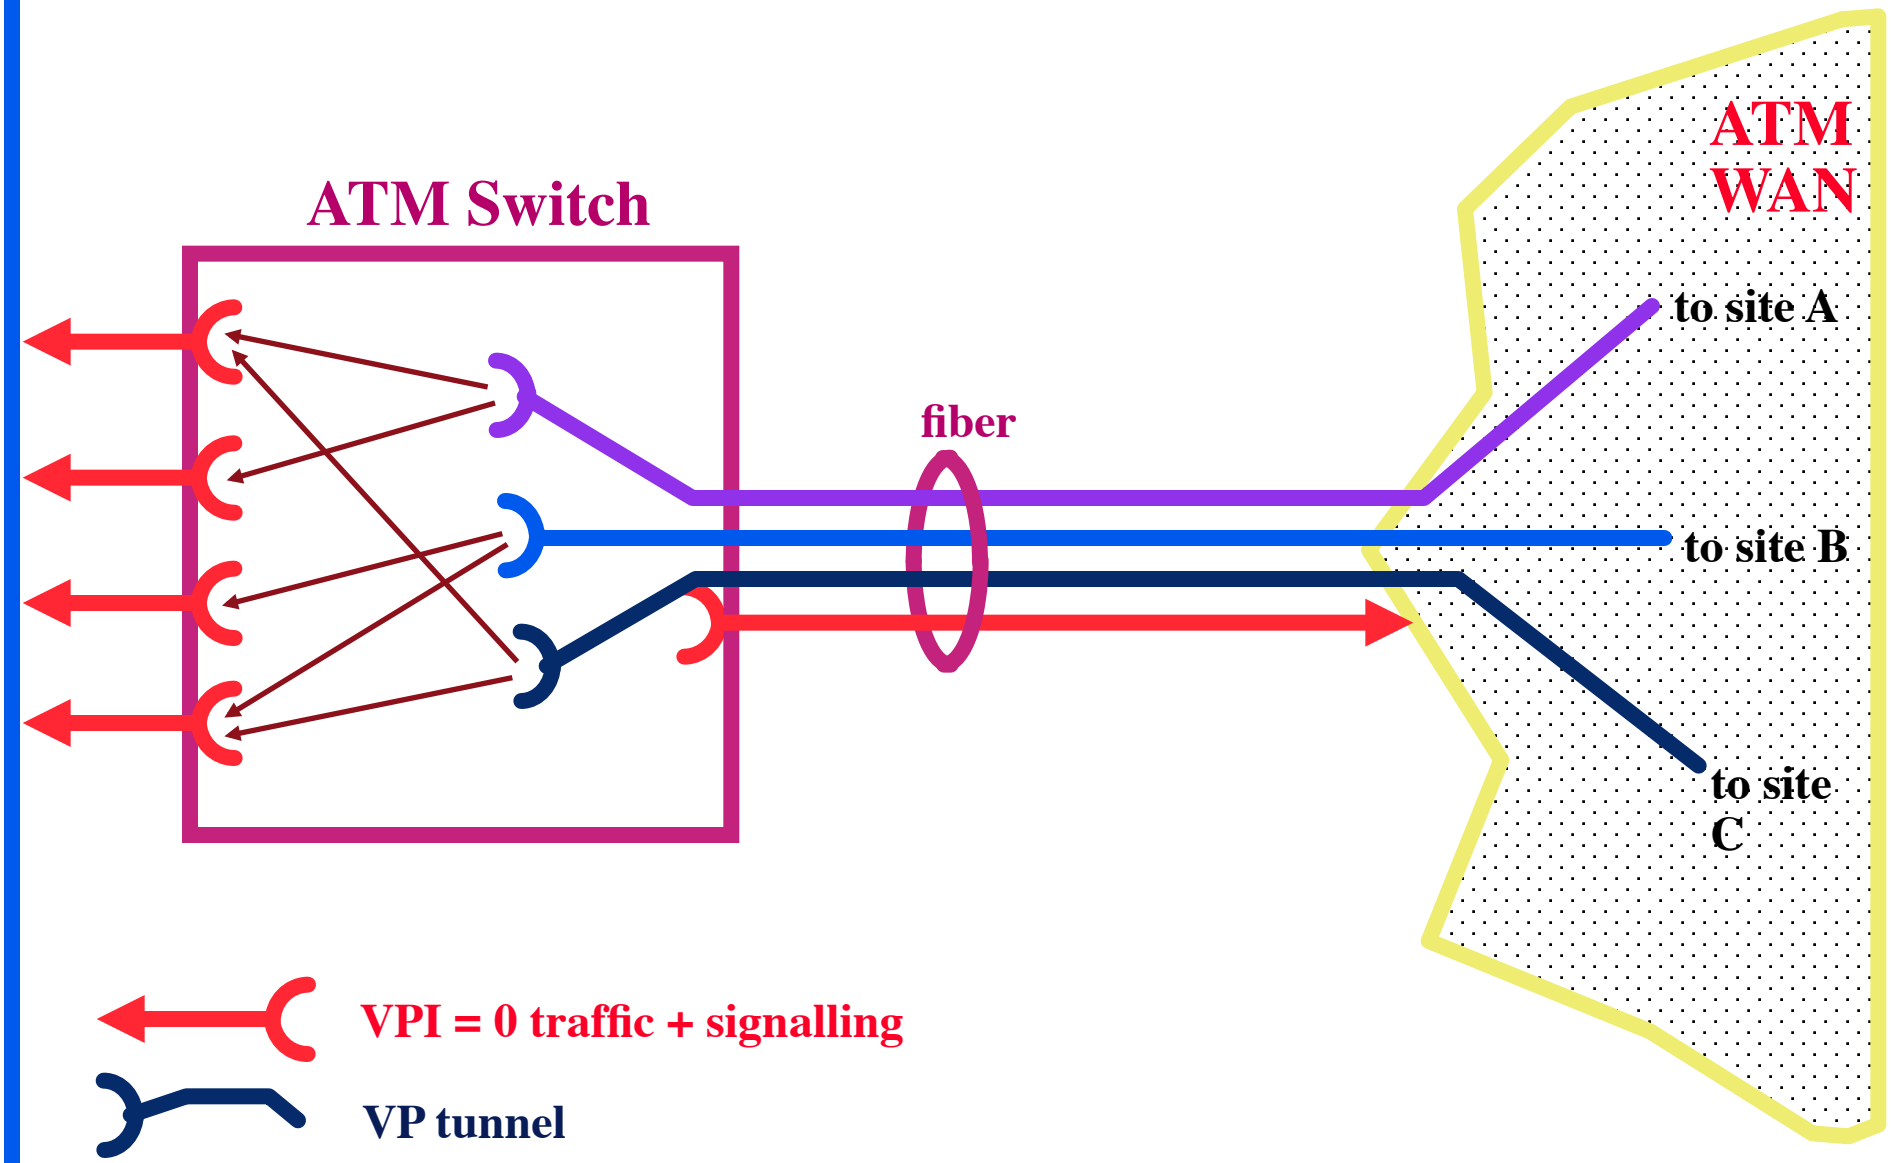

Traffic shaping

- We use 2 * Digital GIGAswitchATM
- VP level shaping
  - One linecard can handle:
    - » VP 0 with signalling
    - » 3 shaped VP's (VPI = [1..63])
      - for best effort VC's through CBR or VBR VP
    - » 6 unshaped VP's (in a later upgrade)
      - best used for CBR VC's to have optimal CDV
  - Shaper ensures that cells are sent out equally spaced
  - Shaper ensures not to exceed traffic contract for entire VP
- VCI level shaping
  - a tunnel can contain CBR, VBR, ABR, UBR VC's (VCI = [1..4095])
  - a tunnel can handle point-to-multipoint connections and LANE
  - for VBR, ABR, UBR VC's local loop's can use :
    - » Credit based flow control
    - » EPD and EFCI

Setup

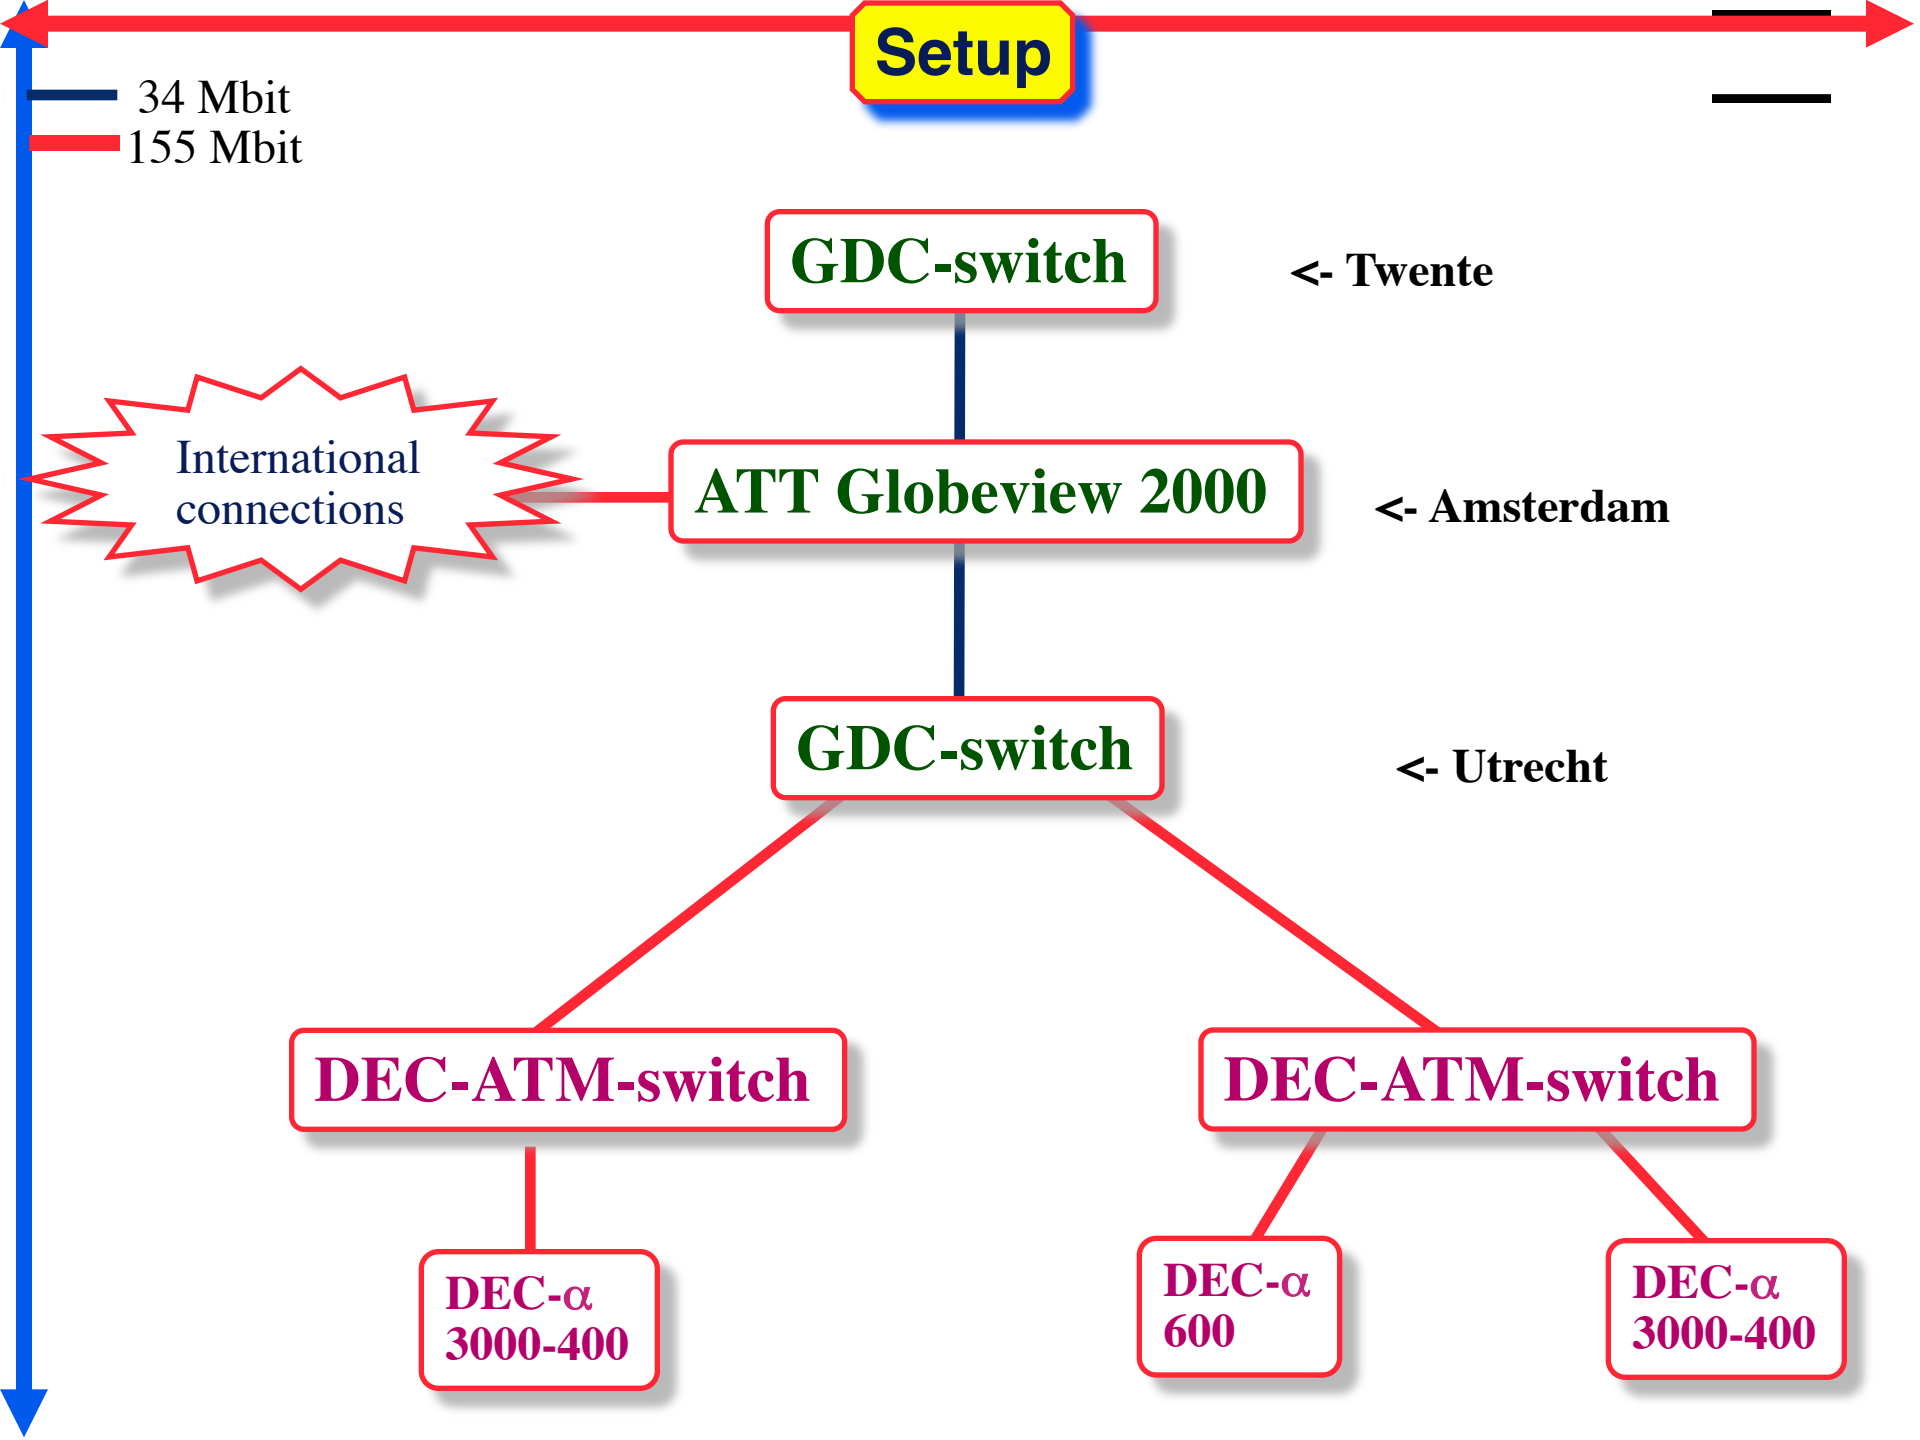
34 Mbit
155 Mbit
GDC-switch
<- Twente
International connections
ATT Globeview 2000
<- Amsterdam
GDC-switch
<- Utrecht
DEC-ATM-switch
DEC-ATM-switch
DEC-α 3000-400
DEC-α 600
DEC-α 3000-400

Setup

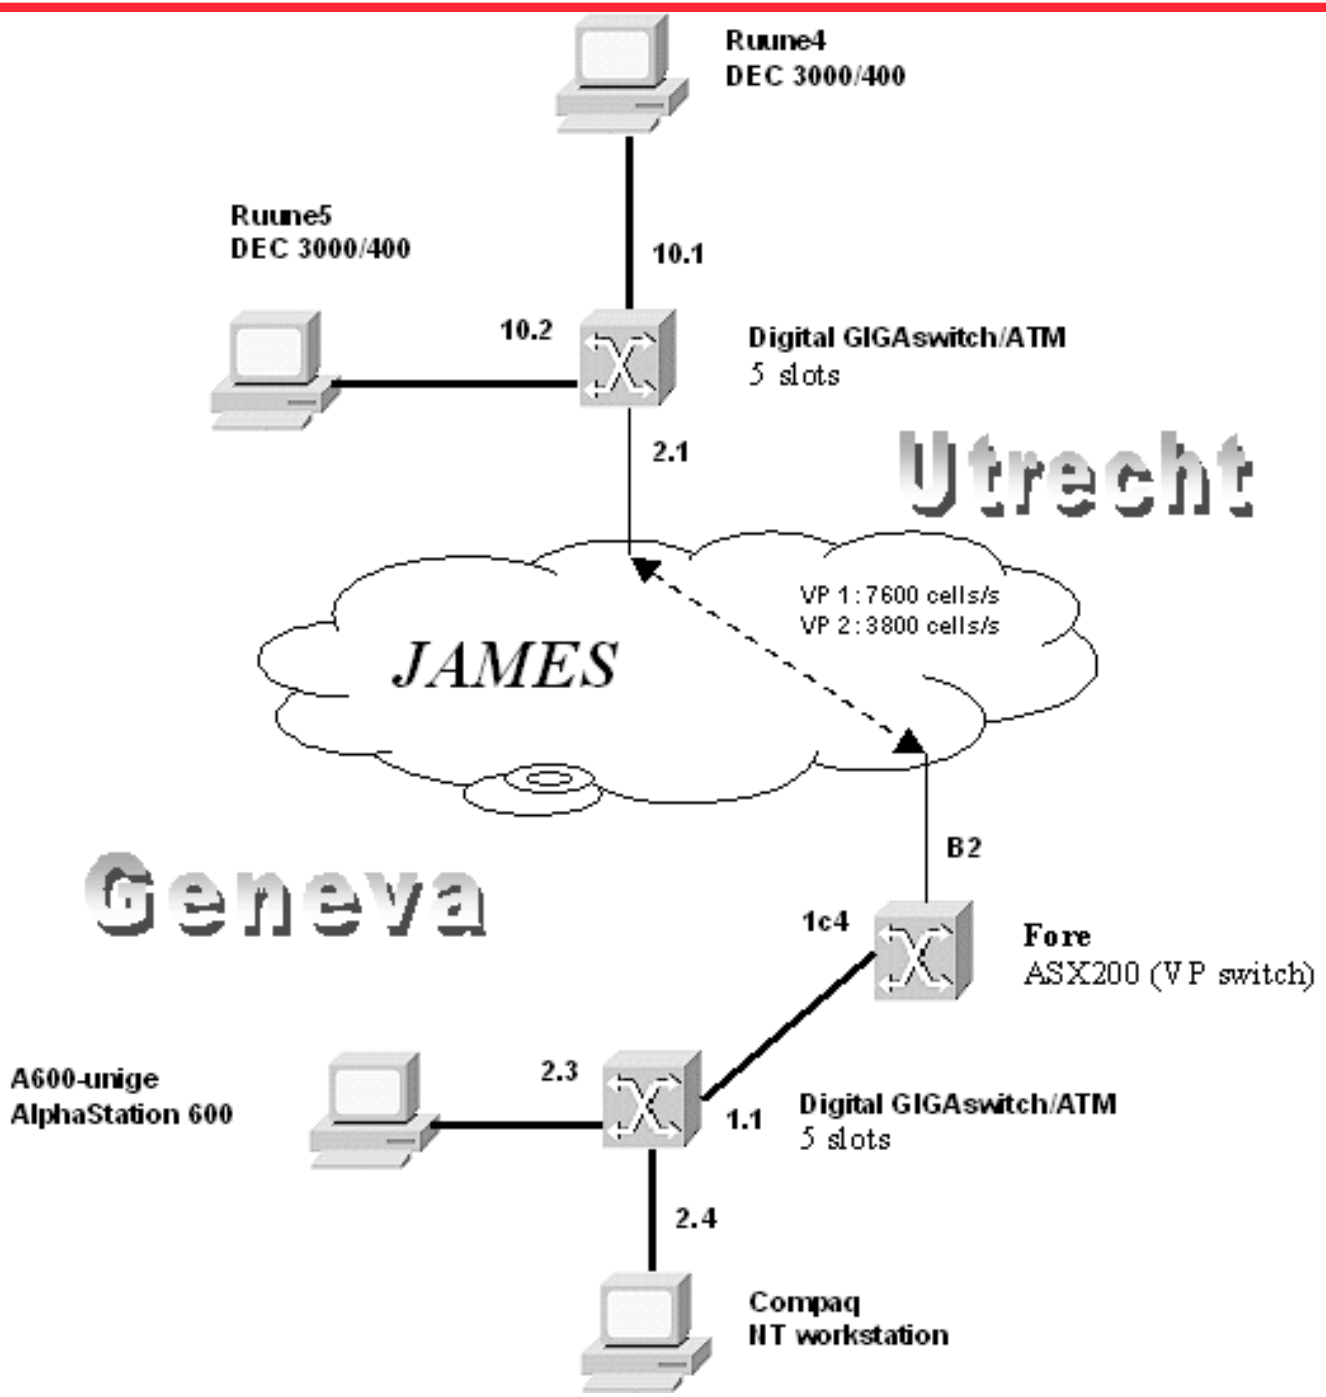
Ruune4
DEC 3000/400
Ruune5
DEC 3000/400
10.1
10.2
Digital GIGAswitch/ATM
5 slots
2.1
Utrecht
VP 1: 7600 cells/s
VP 2: 3800 cells/s
JAMES
B2
Geneva
1c4
Fore
ASX200 (VP switch)
A600-unige
AlphaStation 600
2.3
1.1
Digital GIGAswitch/ATM
5 slots
2.4
Compaq
NT workstation

Eperiments and results local

- GIGAswitches back to back
  - works as advertized after some minor field test and bug fixing
  - tested with up to 3 VP's up to link speed and down to 1 Mbit/s
  - no cell loss when using Flowmaster
  - down to a few % traffic when using no Flowmaster
- GIGAswitches connected via GDC switch
  - worked as advertized
  - GDC does no policing (wrong setting on GDC?)
  - throughput up to 110 Mbit/s depending on tunnel settings in shaper
- VP loop on ATT Globeview via GDC switch
  - worked as advertized
  - throughput close to bandwidth setting of tunnel
  - VPI = 10 <=======> VPI = 11 works so VPI does not need to be the same on both ends!

Eperiments and results UU - UT

- GIGAswitches via SURFnet
  - works as advertized, rtt ≈ 4 ms
  - dynamical routing of DEC GIGAswitch did not work
  - used static routes instead
  - CBR + ABR at same time works, CBR idle -> ABR
- GIGAswitch connected to UB GeoSwitch 155
  - Newbridge does not accept vci > 99, DEC goes up to 4096
  - after each reboot only a certain number of connections
  - signalling problem GeoSwitch <-> Digital
  - GeoSwitch can not shape traffic
  - TCP throughput Utrecht to Twente at most 2 Kbit/s
  - From Twente to Utrecht there was no TCP traffic at all.
  - UDP throughput from Utrecht to Twente close to tunnel bandwidth
  - from Twente to Utrecht no traffic was possible (policing due to no shaping)

Eperiments and results UU - Ge

- GIGAswitches via JAMES
  - VP gnve-ut/atm5, bandwidth = 7600 cells/s.
  - VPI=21 in Utrecht and VPI=3 in Geneva.
  - Digital switch shaper resolution: 1878 cells/s -->
  - Expect a maximum bandwidth of 4*1878 = 7512 cells/s = 2.75 Mbit/s.
  - VP gnve-ut/atm4, bandwidth = 3800 cells/s.
  - VPI=20 in Utrecht and VPI=2 in Geneva.
  - bandwidth of 3756 cells/s = 1.38 Mbit/s.
- Test results:
  - local credit based flow control, CBR VP on WAN
  - rtt = 17.9 ms (≈ 3580 km fiber)
  - CBR + ABR --> ABR gets dynamically idle cells of CBR
  - SVC setup times measurement when static routes t ARP server are set
  - SVC setup times seem to be around 300 ms.
  - Needs further investigation

round trip times

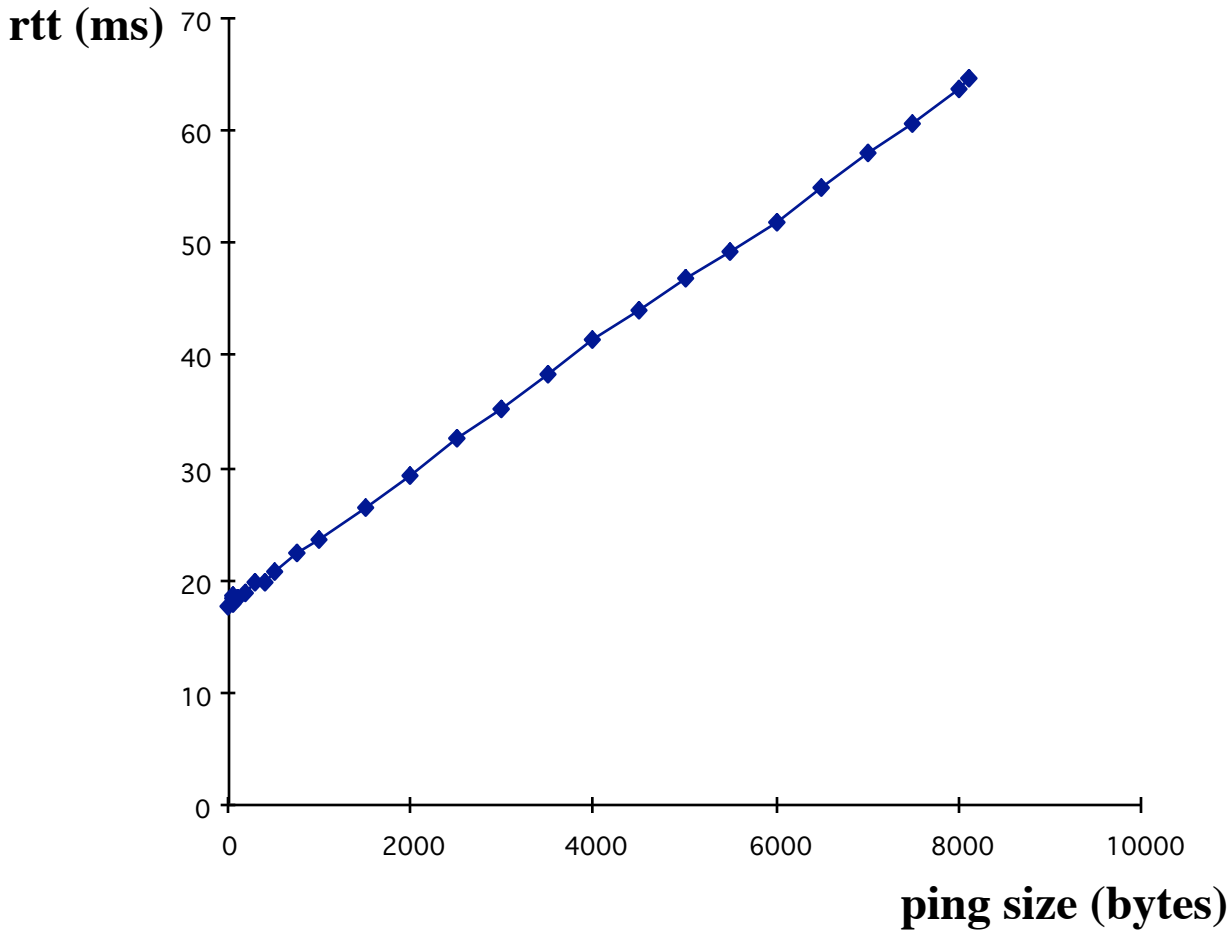

rtt = size * 5.746069E-03 + 17.7768 -->
bw = 2 * 8 / (5.746069E-03 * 1.0E-03) = 2.78 Mbit/s

LANE results

•Data transfer rates with LANE:

Sender	Receiver	Kbytes/s	Mbits/s
MTU = 1516			
4 VME 68k + Sun + Pc	Alpha 3000-400	3500	28
11 Dec+HP Workstations	Alpha Station 600	11500	94
Alpha Station 600	Alpha Station 600	6250	50
2 Alpha Station 3000-400	Alpha Station 600	8625	69
MTU = 9234			
Alpha Station 600	Alpha Station 600	16000	128
2 Alpha Station 3000-400	Alpha Station 600	16400	131

CLIP and LANE about equal heavy for CPU when same MTU used.

User and System load

User and system load in instructions per byte for the data transmitting and data receiving computers. The transmitting computers rate 142 MIPS (V2.1), the receiving computers rate 459 MIPS.

transmit protocol	MTU	Rate MByt/s	user %	sys %	user i/Byt	sys i/Byt
Ethernet	1516	0.47	1.7	6.6	5.1	20
LAN Emulation	1516	4.4	13	87	4.3	28
LAN Emulation	9234	8.2	27	73	4.6	13
Classical IP on ATM	9200	8.1	24	63	4.3	11

receive protocol	MTU	Rate MByt/s	user %	sys %	user i/Byt	sys i/Byt
Ethernet	1516	0.94	0.8	5.6	3.9	27
LAN Emulation	1516	8.7	6.5	60	3.4	23
LAN Emulation	9234	16.4	6.3	43	1.8	12
Classical IP on ATM	9200	16.2	6.1	42	1.7	12

Discussion

- Traffic contracts
- shaping
- local flow control versus end to end
- requirements for edge switches ?
  - 1) UNI 3.1 support, migration to 4.0 support
  - 2) UBR, CBR, VBR en ABR support (ABR relevant with flow control/UNI 4.0)
  - 3) EPD and PPD
  - 4) PNNI (IISP)
  - 5) VC switching
  - 6) VP switching
  - 7) VP tunneling support
  - 8) Shaping on VP and VC (also in VP)
  - 9) LANE support (?)
  - 10) Accounting (?)
- Suggestions are welcome!

Videoconferencing

- No ISDN, no leased lines for financial and political reasons
- Mbone over QoS cirquits (see MERCI)
- native ATM based (FORE Nemesys)
- Need 4 - 6 Mbit/s for broadcast quality
- Need one to one and one to many
- We started using Nemesys boxes for meetings between University Twente and University Utrecht
- Doing MBone experiments over SURFnet
  - **connection setup and teardown**
  - **hardware compression**

- **What is needed**
  - tabel with along axes:
    - » Type of equipment of sender
    - » Type of equipment of receiver
    - » quality required, available network, one to on or many
  - Result:
    - » Recommended hardware
    - » Recommended software
    - » Recommended settings, network bandwidth
  - Doing this in a SURFnet project.