



Network Tomography

Aiko Pras

*Design and Analysis of Communication Systems (DACS) University
of Twente*

the Netherlands

April 4, 2008



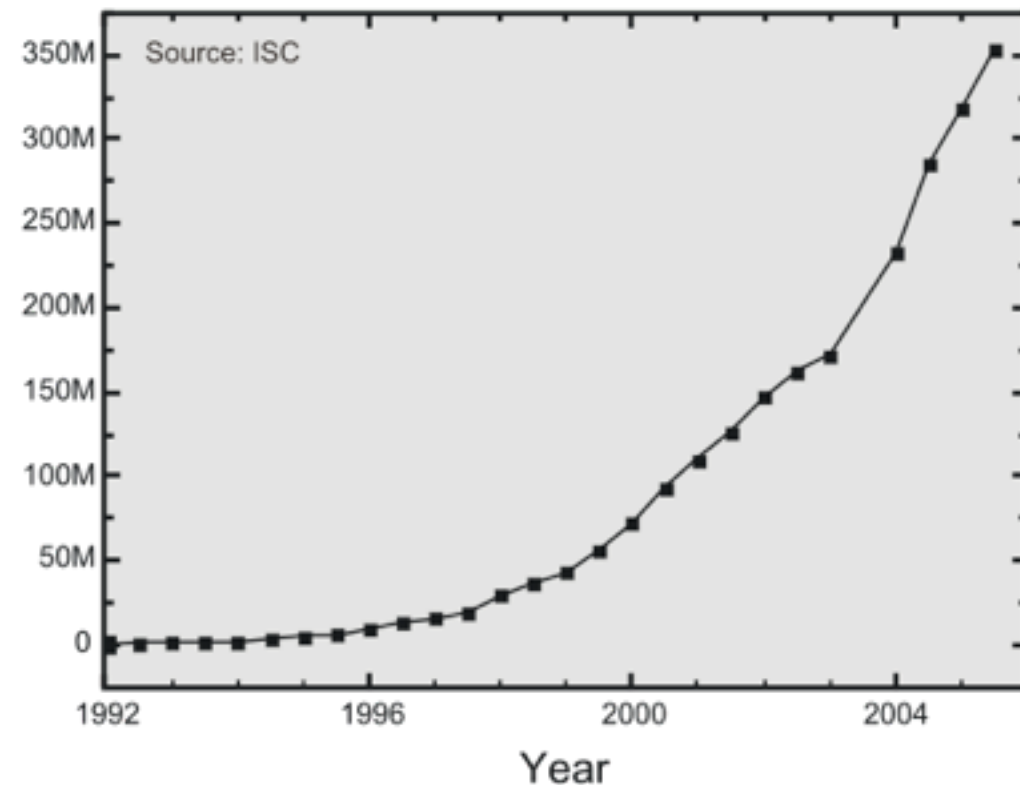
Some (private) history

- 1980
- Bachelor “Adaptive routing”
 - Masters “Connectionless network protocols”
 - UT connected to UNET (300 bps)
 - Desktops connected via Ethernet
- 1990
- UT connected to Internet
 - End of OSI “dream”
 - First web experiences
 - IP seemed to become dominant protocol
- 2000
- End of ATM
-
- A vertical black arrow pointing downwards, indicating the progression of time from 1980 to 2000.

Observation

- Internet has been growing very fast
- Internet has become very complex
- Too complex to predict its behavior, even if behavior of constituent parts is known
- **Design of core IP layer is done!**

Number of Internet Hosts




Observation

- Resembles a living organism
- Measure its health
- Network tomography

United States
China
Germany
United Kingdom
France
Brazil

Japan
Philippines
Russian Federation
Malaysia
Hong Kong





Internet health measurements



- How is the Internet growing?
 - How many users?
 - How much traffic?
 - What applications are being used?
 - Does capacity match demand?
- Is the Internet still healthy?
 - Is Quality of Service OK?
 - Is it under attack?



How to measure?

- Passive:
 - Netflow
 - TCPdump
- Active
 - Ping
 - Within TCP
 - Special probes


Netflow



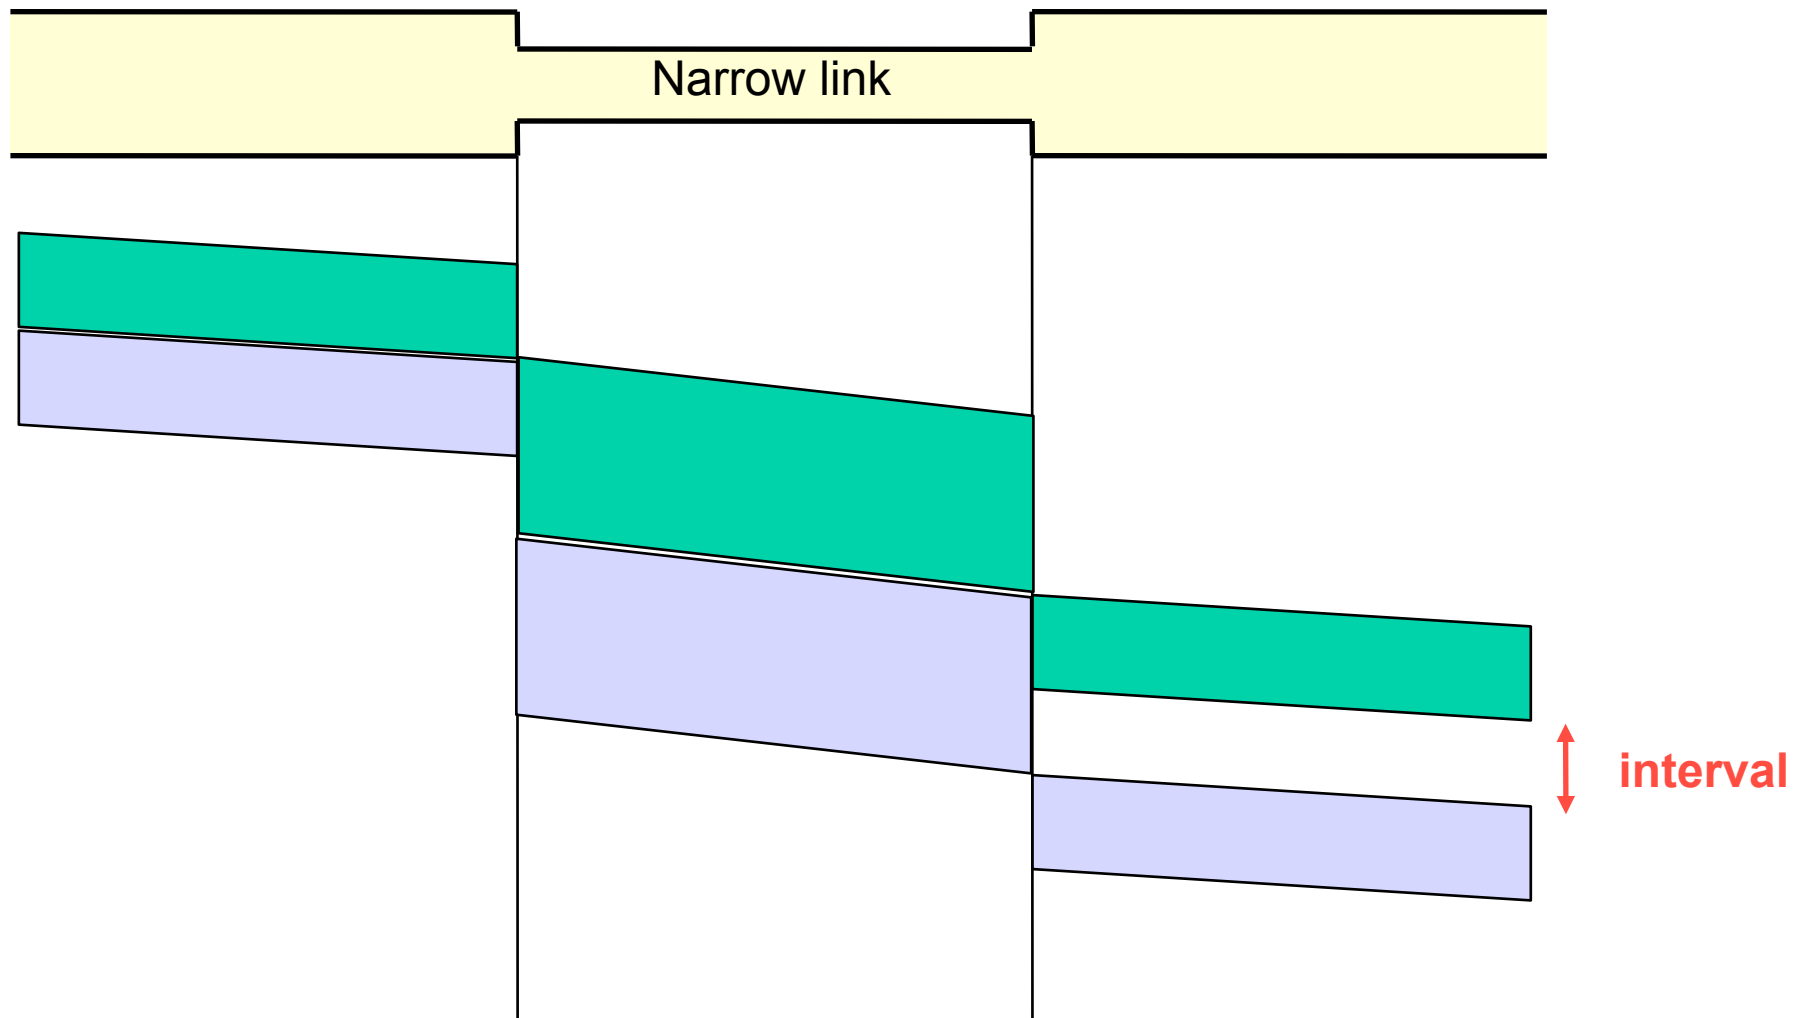
CAPTION

Measurement hardware

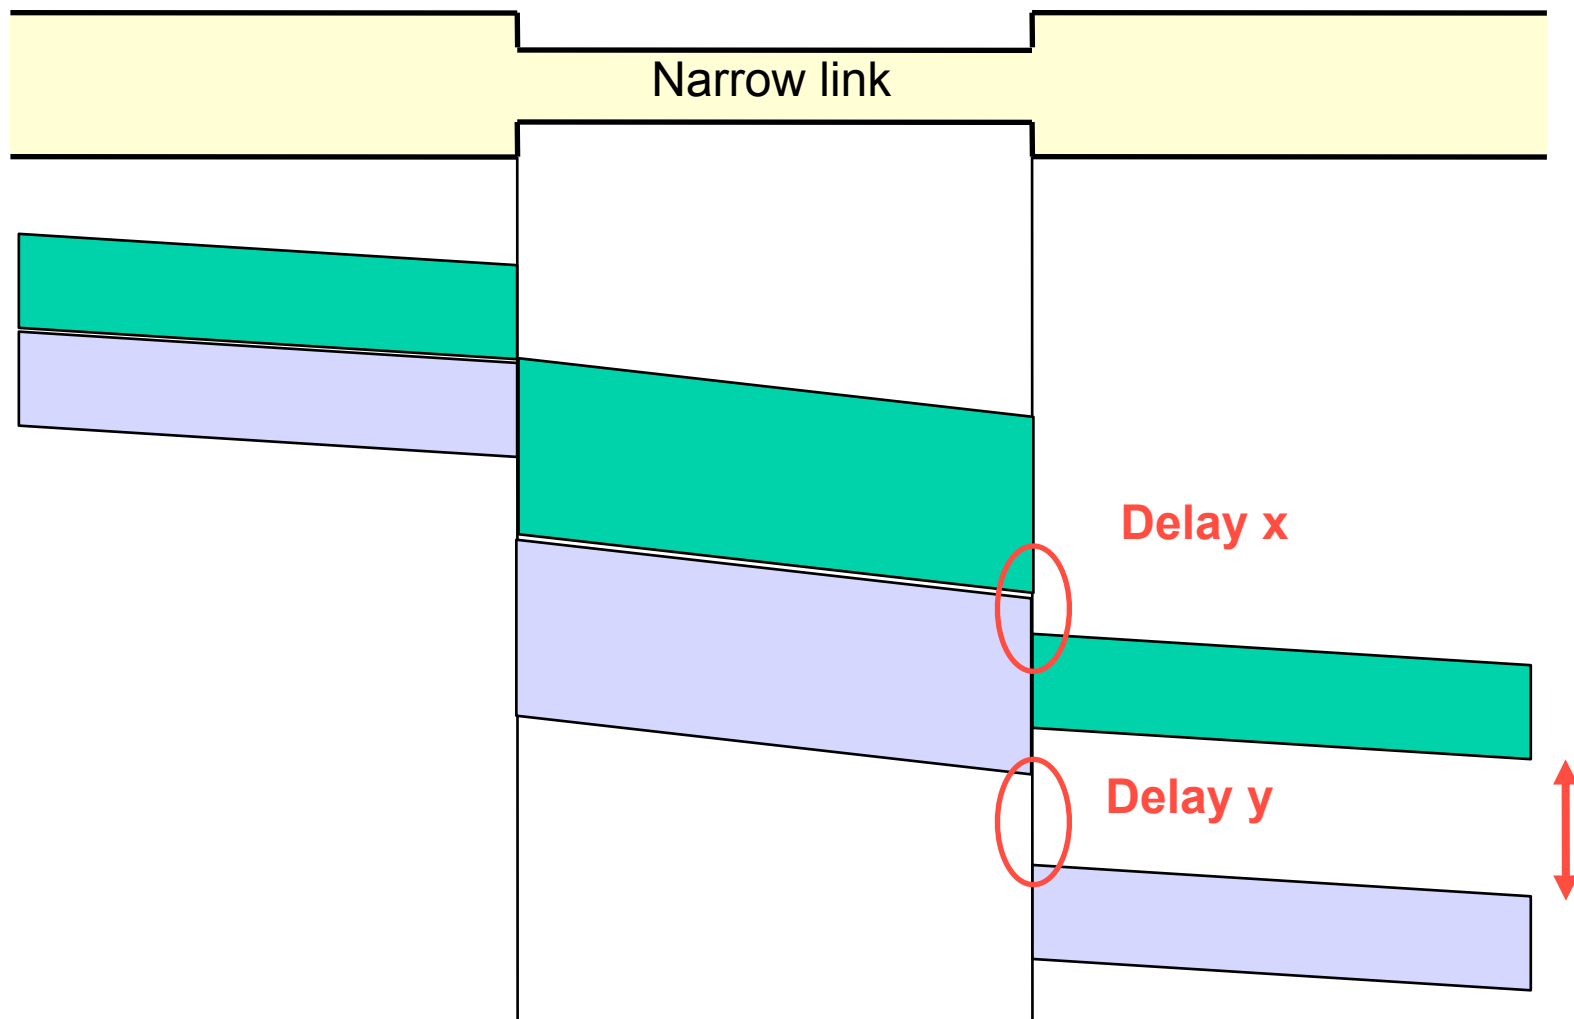
- From simple PC
- To expensive cards



Bandwidth estimation



Bandwidth estimation





ANSCHUTZ

Score sheet with columns for names and scores.





Example: SURFnet6

- Has been installed in previous years
- Access lines: 10 Gigabit (multiple times)
- Based on λ switching
- Direct “optical” connections between end users
- Backbone links: many 10 / 40 Gigabit optical connections



SURFnet 6

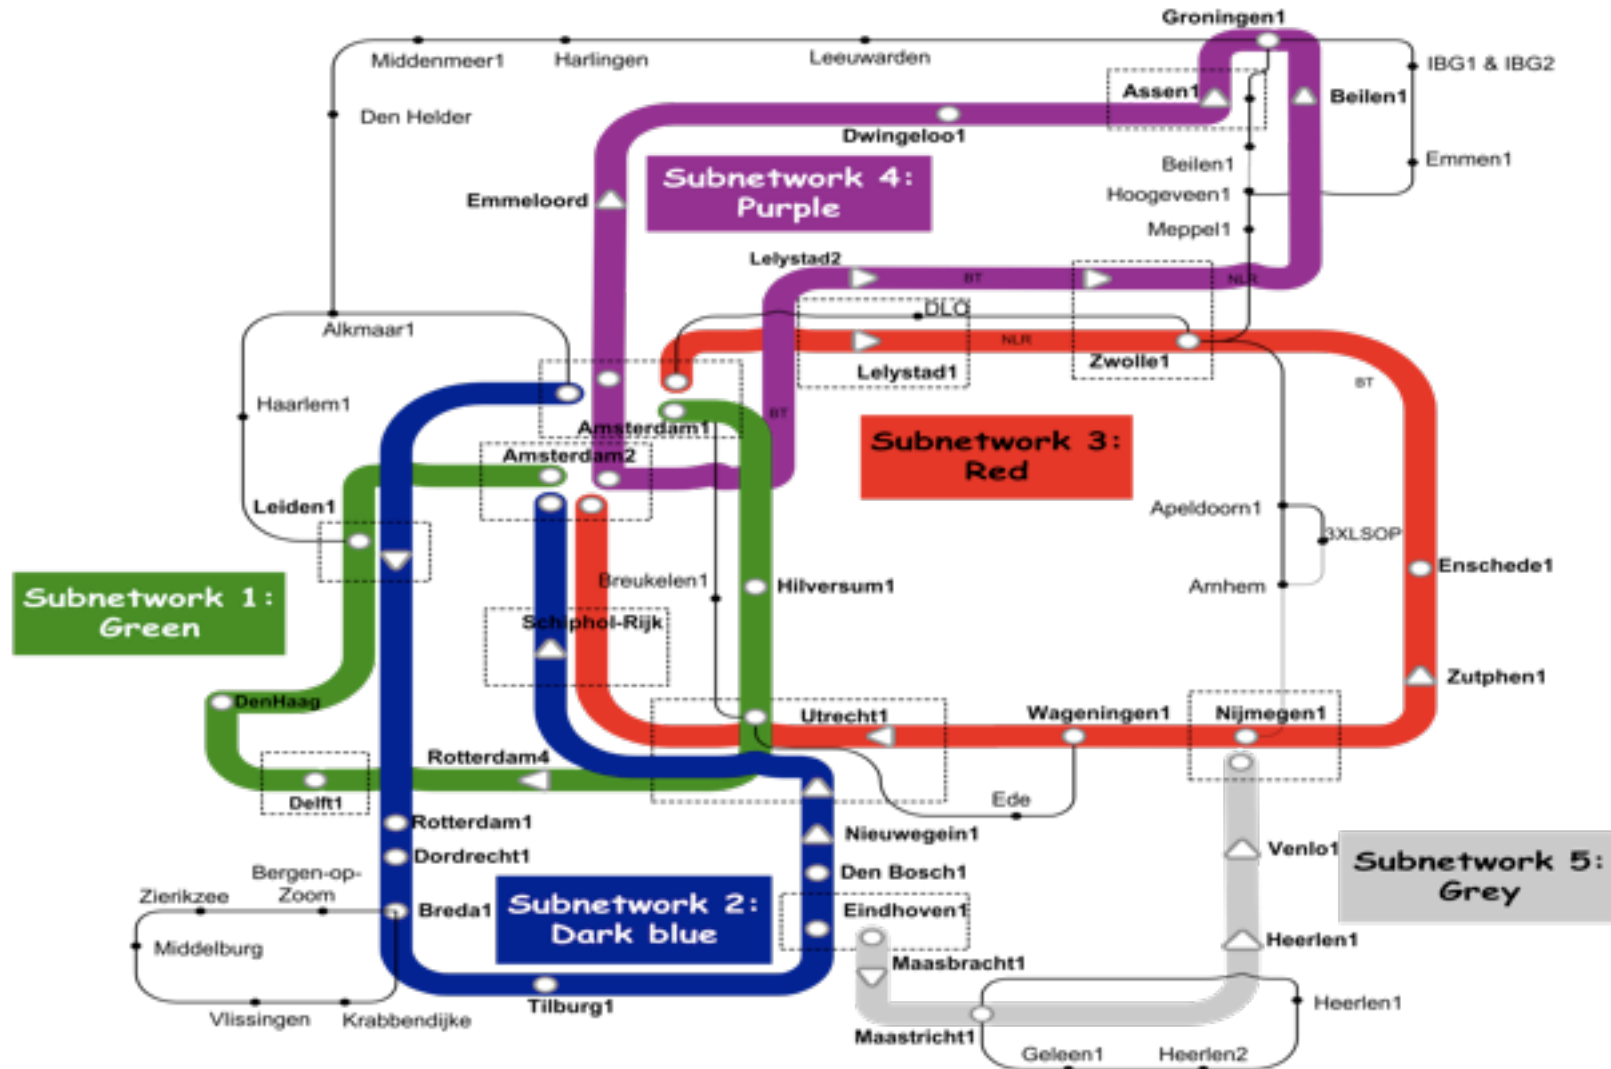


Picture: Yuri Demchenko





SURFnet 6





SURFnet6 Questions

Are such networks still useful?

Does quality, as perceived by the user, still improve?

Isn't it like increasing the maximum speed of a car from 200 km/h to 400 km/h?



SURFnet6 Questions

Is the network a performance bottleneck?

Or are End systems the bottleneck?


Or is there no bottleneck at all:

- Application has a streaming nature
- Application has an interactive nature



IF APPLICATION IS “BOTTLENECK”

Streaming application



Fourier analysis



IF NETWORK IS BOTTLENECK

CONGESTION:

- PACKET LOSS
- DELAY / JITTER

IN WHICH PART OF THE NETWORK?



SURFnet delay and jitter

Last hour IPv6 maximum RTT SURFnet backbone

Green: RTT < 12 Yellow: 12 <= RTT < 24 Red: RTT >= 24 Red: NA = 100% packet loss Gray: No RTT available

last hour IPv6 maximum RTT in ms	D01	D02	D03	D04	D05	D06	D07	D08	D09	D10	D11	D12	D13	D14	D15	D16
S01 ms1.amsterdam1.surf.net	1.3	17.1	4.5	4.6	10.5	6.2	5.0	32.0	17.2	34.3	15.1	15.7	18.3	23.1	30.4	22.0
S02 ms1.delft1.surf.net	11.9	0.5	7.4	28.6	15.8	29.3	10.3	39.2	29.3	33.6	29.6	7.4	6.8	29.7	34.7	27.4
S03 ms1.den Haag1.surf.net	4.6	10.3	0.5	6.7	18.8	18.2	17.4	35.6	9.3	31.3	15.8	14.4	8.8	28.2	23.7	28.3
S04 ms1.eindhoven1.surf.net	11.4	28.3	6.1	0.5	13.7	9.5	28.1	32.4	22.5	10.4	13.7	13.7	14.9	26.3	32.3	28.3
S05 ms1.enschede1.surf.net	6.5	8.7	20.6	8.8	0.4	20.6	14.6	41.9	12.0	30.9	6.1	14.9	9.7	22.4	43.4	23.2
S06 ms1.groningen1.surf.net	5.5	29.4	15.7	8.6	20.0	0.5	33.9	42.4	24.3	40.2	11.1	8.1	13.8	25.4	43.4	34.4
S07 ms1.hilversum1.surf.net	12.2	7.4	17.4	10.6	8.8	31.1	0.7	41.7	33.0	30.2	30.7	31.1	7.0	27.6	22.4	30.2
S08 ms1.leiden1.surf.net	7.8	19.3	14.6	41.6	41.2	41.5	41.1	0.5	32.8	16.5	42.9	41.9	30.1	41.3	43.7	38.9
S09 ms1.maastricht1.surf.net	7.1	16.9	10.7	10.5	11.0	15.1	17.0	20.6	0.5	35.1	20.9	10.7	18.1	20.4	25.5	30.5
S10 ms1.nijmegen1.surf.net	24.1	37.0	6.4	34.3	18.6	9.6	28.3	43.3	10.6	0.6	9.1	15.5	36.4	32.1	38.6	28.0
S11 ms1.rotterdam1.surf.net	5.7	11.9	8.4	14.3	6.3	8.3	34.9	43.0	9.3	19.9	0.5	10.4	17.6	23.4	39.6	31.3
S12 ms1.tilburg1.surf.net	5.7	7.7	16.9	8.2	7.5	16.1	11.6	41.4	13.0	38.3	13.7	0.5	5.6	36.1	31.3	27.1
S13 ms1.utrecht1.surf.net	10.9	10.9	21.1	31.3	23.6	30.5	14.9	35.3	10.2	36.6	30.2	30.6	0.6	31.1	37.8	30.3
S14 ms1.wageningen1.surf.net	4.9	11.4	22.0	29.6	32.6	32.4	8.0	38.2	33.1	38.0	32.0	33.4	21.0	0.4	41.1	28.4
S15 ms1.zwolle1.surf.net	41.0	43.9	37.9	16.6	43.8	15.2	43.7	38.5	47.5	44.0	6.4	13.3	41.8	30.4	0.6	31.9
S16 ms2.amsterdam1.surf.net	27.5	5.1	20.2	29.7	12.3	31.0	12.2	34.7	12.3	19.6	31.9	32.4	10.5	30.3	38.1	0.5

Measurement interval: Fri May 12 13:00:00 2006 until Fri May 12 14:00:00 2006

Last day IPv4 jitter SURFnet backbone

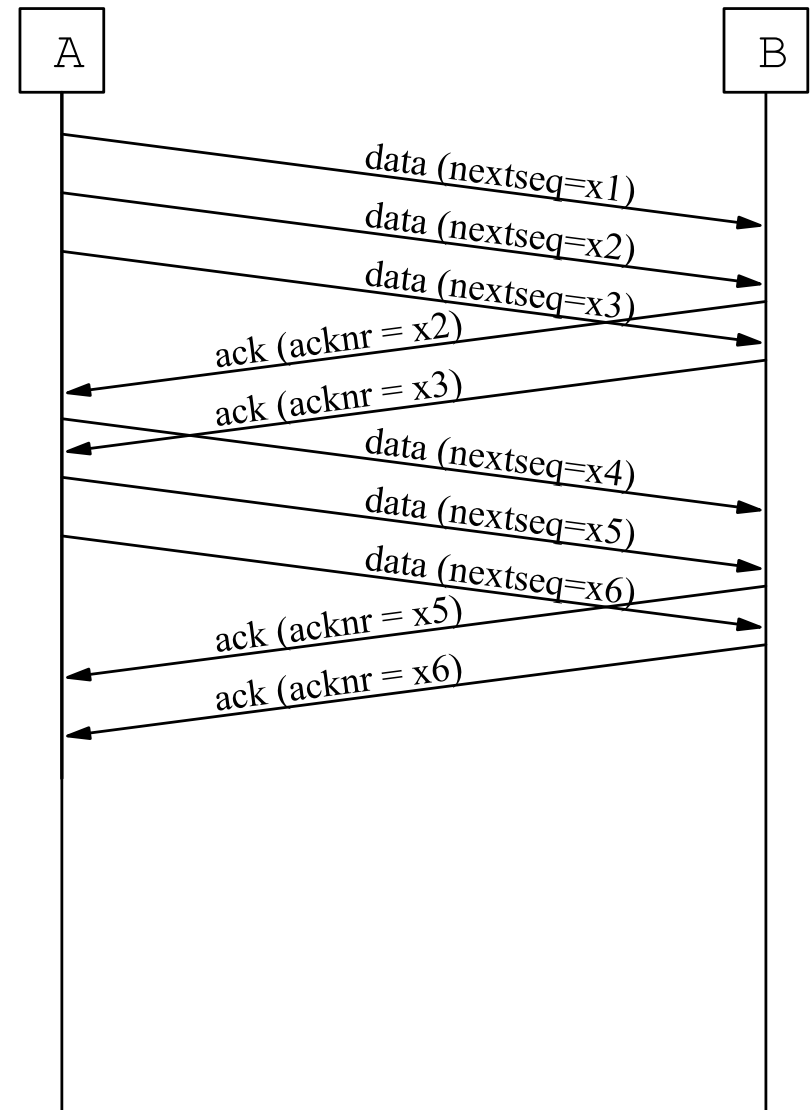
Green: jitter < 0.5 ms Yellow: 0.5 ms <= jitter < 1 ms Red: jitter > 1 ms Red: NA = 100% packet loss Gray: No jitter available

last day IPv4 jitter RTT in ms	D01	D02	D03	D04	D05	D06	D07	D08	D09	D10	D11	D12	D13	D14	D15	D16
S01 ms1.amsterdam1.surf.net	0.2	0.4	0.1	1.8	0.1	0.1	0.1	0.1	0.1	0.3	0.2	0.2	0.2	0.2	0.2	0.3
S02 ms1.delft1.surf.net	1.3	0.0	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.2	0.1	0.1	0.2
S03 ms1.den Haag1.surf.net	0.2	0.1	0.0	0.2	0.1	0.1	0.1	0.2	0.1	0.1	0.1	0.1	0.1	0.2	0.1	0.3
S04 ms1.eindhoven1.surf.net	0.2	0.1	0.1	0.0	0.1	0.1	0.1	0.1	0.1	0.1	0.2	0.1	0.1	0.1	0.1	0.1
S05 ms1.enschede1.surf.net	1.9	0.1	0.1	0.1	0.0	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
S06 ms1.groningen1.surf.net	1.3	0.1	0.1	1.3	0.1	0.0	0.1	0.1	0.2	0.1	0.4	0.1	0.1	0.1	0.1	0.1
S07 ms1.hilversum1.surf.net	0.8	0.4	0.1	0.1	0.1	0.1	0.0	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.2
S08 ms1.leiden1.surf.net	0.2	0.1	0.5	0.1	0.1	0.1	0.3	0.0	0.1	0.1	0.1	0.1	0.2	0.2	0.1	0.2
S09 ms1.maastricht1.surf.net	0.4	0.1	0.1	2.1	0.1	0.6	0.3	0.1	0.0	0.1	0.1	0.1	0.1	0.1	0.1	0.1
S10 ms1.nijmegen1.surf.net	0.2	0.1	0.1	0.9	0.1	0.1	0.1	0.1	0.1	0.0	0.1	0.1	0.1	0.1	0.1	0.2
S11 ms1.rotterdam1.surf.net	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.5	0.0	0.1	0.2	0.1	0.2	0.2
S12 ms1.tilburg1.surf.net	0.2	0.1	0.2	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.0	0.1	0.1	0.1	0.2
S13 ms1.utrecht1.surf.net	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.0	0.1	0.1	0.2
S14 ms1.wageningen1.surf.net	0.2	0.1	0.1	0.1	0.1	0.1	0.2	0.2	0.1	0.1	0.4	0.1	0.1	0.0	0.1	0.8
S15 ms1.zwolle1.surf.net	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.6	0.1	0.1	0.1	0.0	0.1
S16 ms2.amsterdam1.surf.net	0.2	0.1	1.1	0.1	0.2	0.1	0.3	0.1	0.1	0.2	0.2	0.1	0.2	0.2	0.2	0.0

Measurement interval: Thu May 11 00:00:00 2006 until Fri May 12 00:00:00 2006

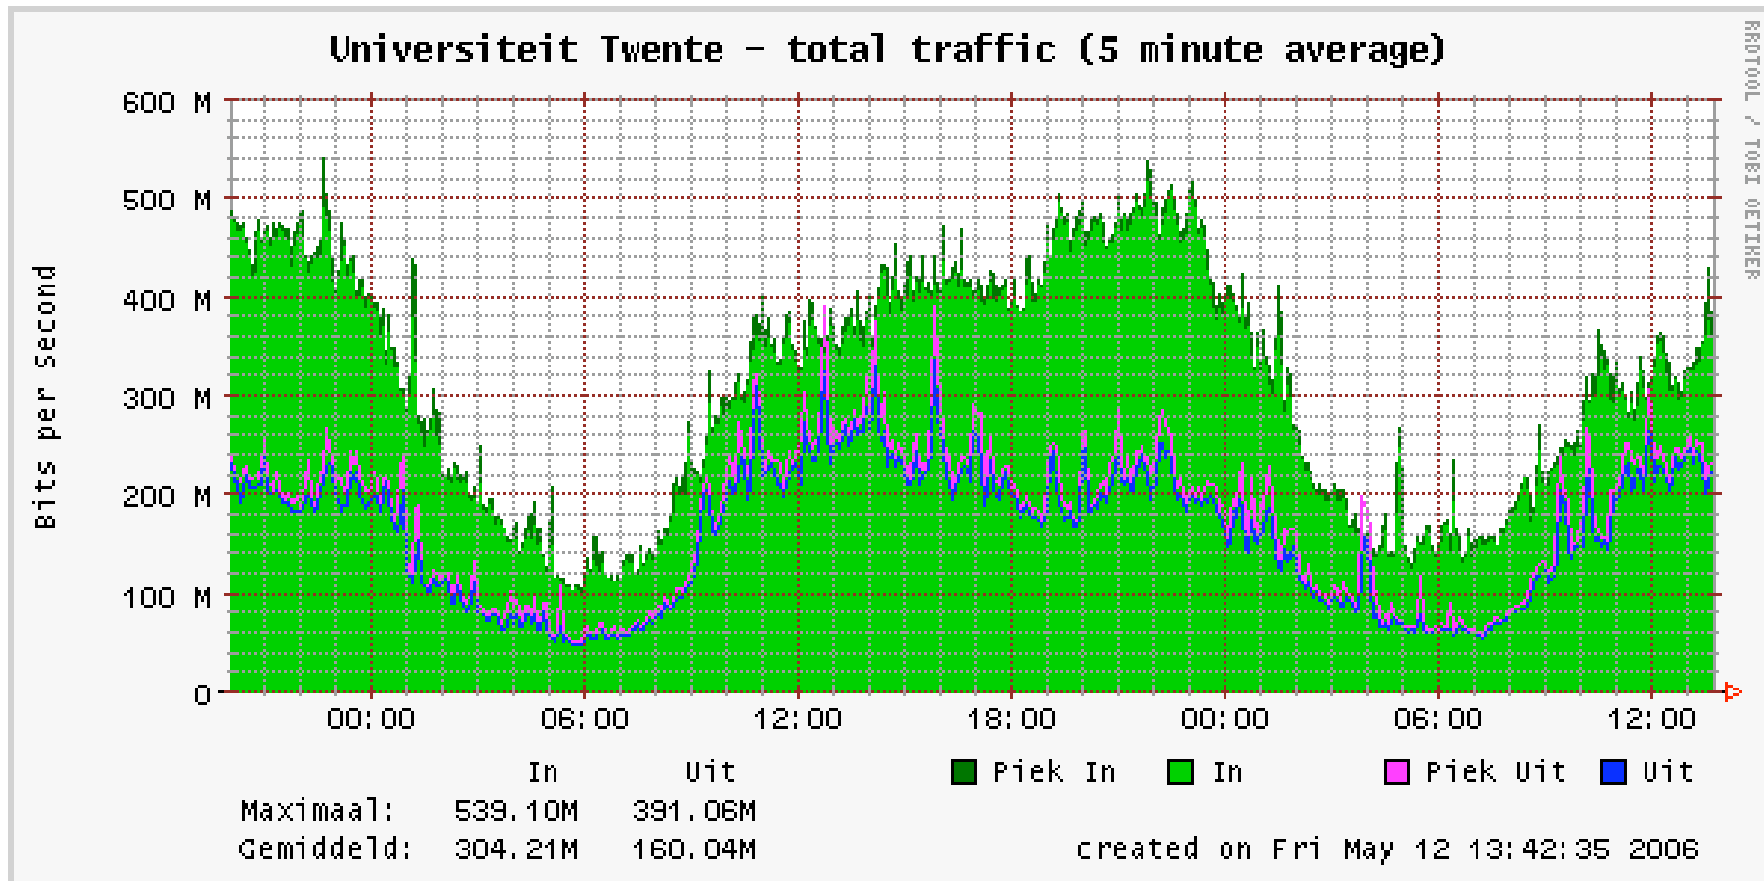
IF END SYSTEM IS BOTTLENECK

- The receiver cannot process the data faster
- The receiver cannot advertise a larger receive window
- With better OS settings bottleneck may disappear







How much capacity would be needed?



Other possibilities

- Geo-location of hosts
- Measure background radiation
- Detect viruses and worms
- Detect IP spoofing
- Etc. Etc. Etc.





CONCLUSIONS

- Design of the core Internet is ready
- Consider the Internet to be a living organism
- Measure the health of this “organism”
 - Delay
 - Loss
 - Bandwidth
 - Etc. etc. etc.