Lambda's between reception and buffet

www.science.uva.nl/~delaat

Cees de Laat

University of Amsterdam

SARA
NCF

GigaPort EU

Contents of this talk

- This page is intentionally left blank
  - Ref: www.this-page-intentionally-left-blank.org

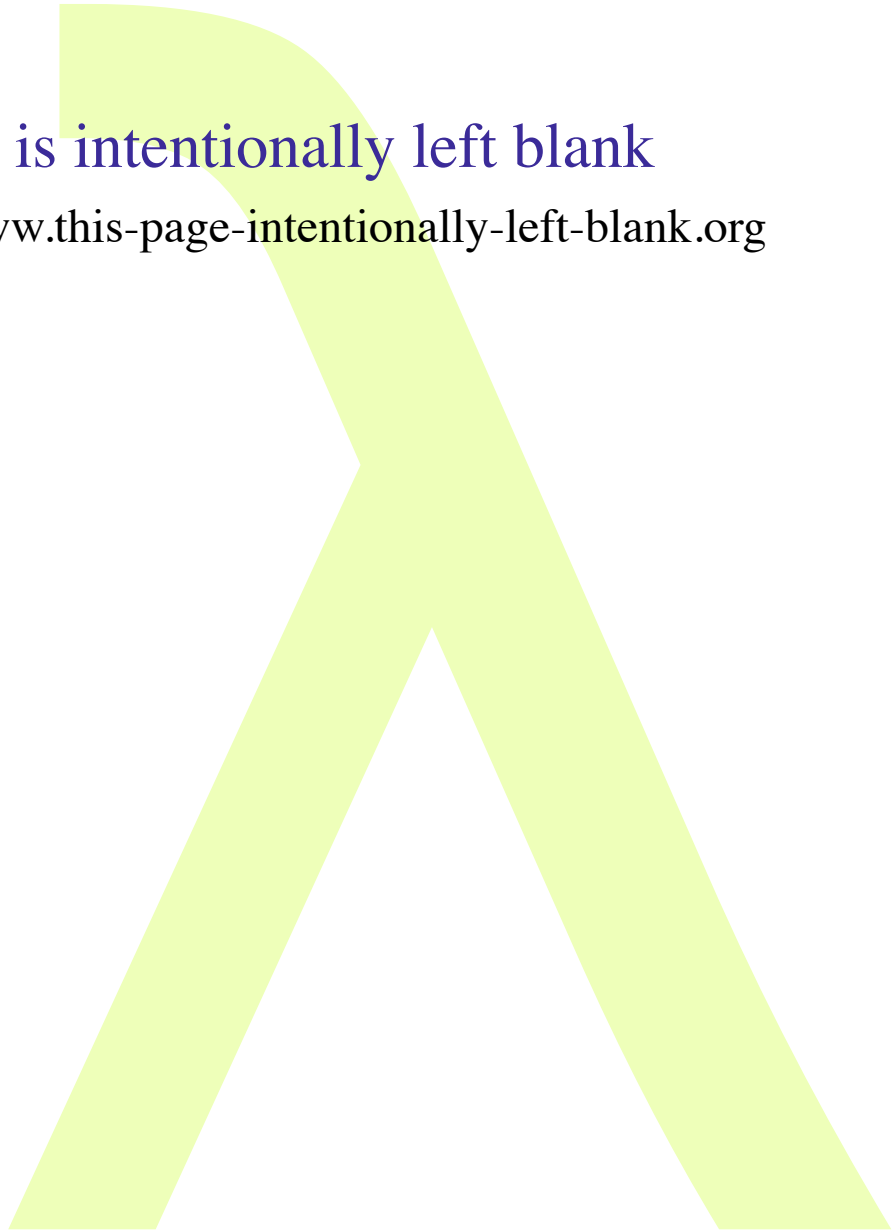

Four LHC Experiments: The Petabyte to Exabyte Challenge

- **ATLAS, CMS, ALICE, LHCB**

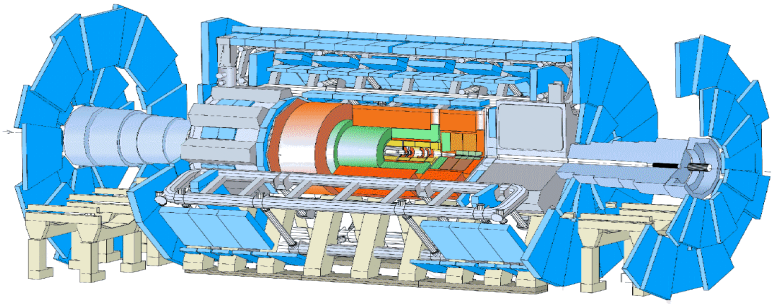

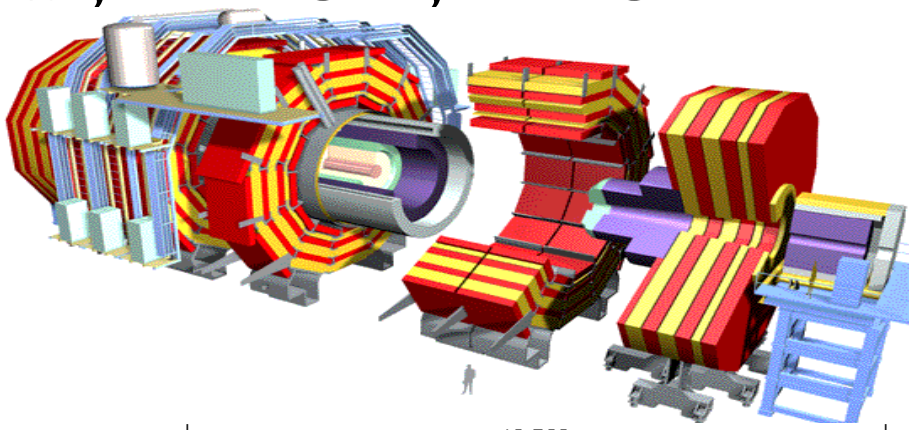

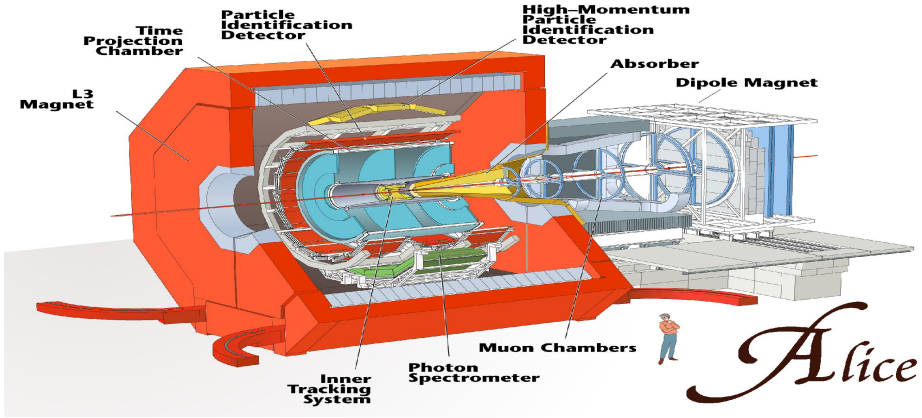

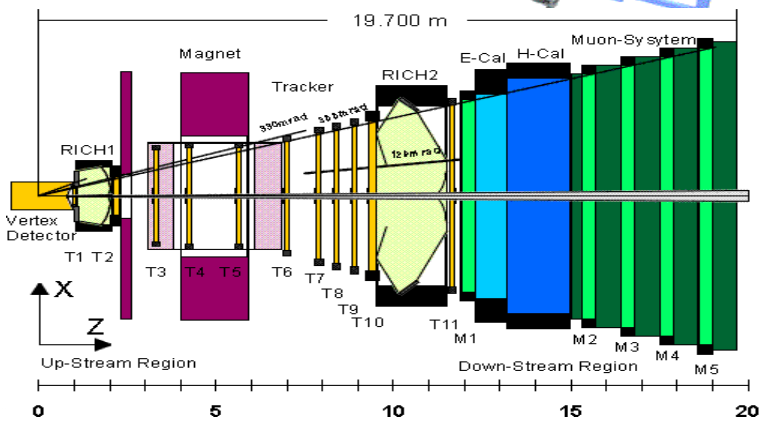

6000+ Physicists & Engineers; 60+ Countries; 250 Institutions

Tens of PB 2008; To 1 EB by ~2015

Hundreds of TFlops To PetaFlops

VLBI

ger term VLBI is easily capable of generating many Gb of data per
The sensitivity of the VLBI array scales w
(=data-rate) and there is a strong push to
. Rates of 8Gb/s or more are entirely feasibl
der development. It is expected that parall
orrelator will remain the most efficient approa
s distributed processing may have an appli
lti-gigabit data streams will aggregate into la
or and the capacity of the final link to the da
tor.

Westerbork Synthesis Radio Telescope - Netherlands

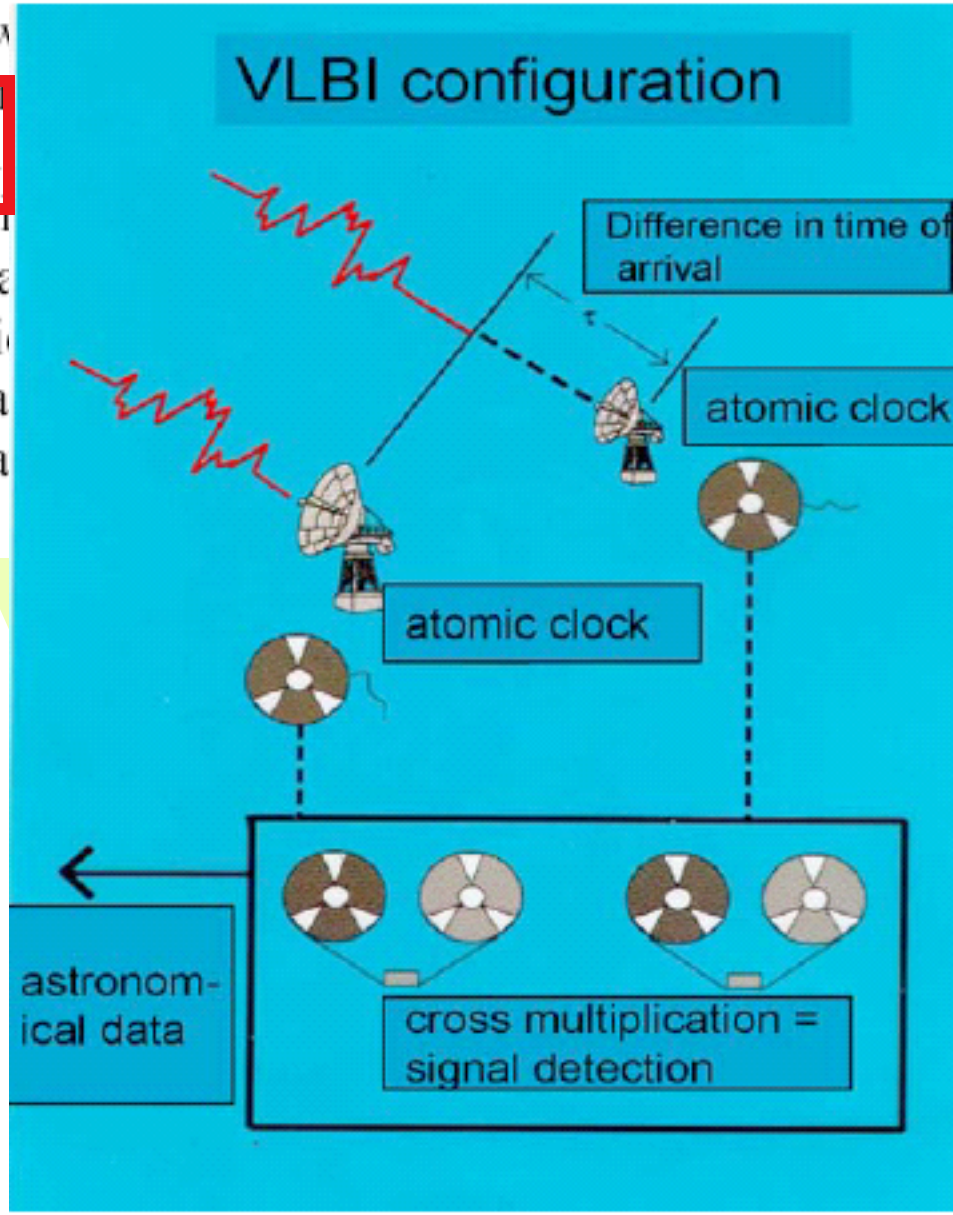

Lambdas as part of instruments

GigaPort

LOFAR

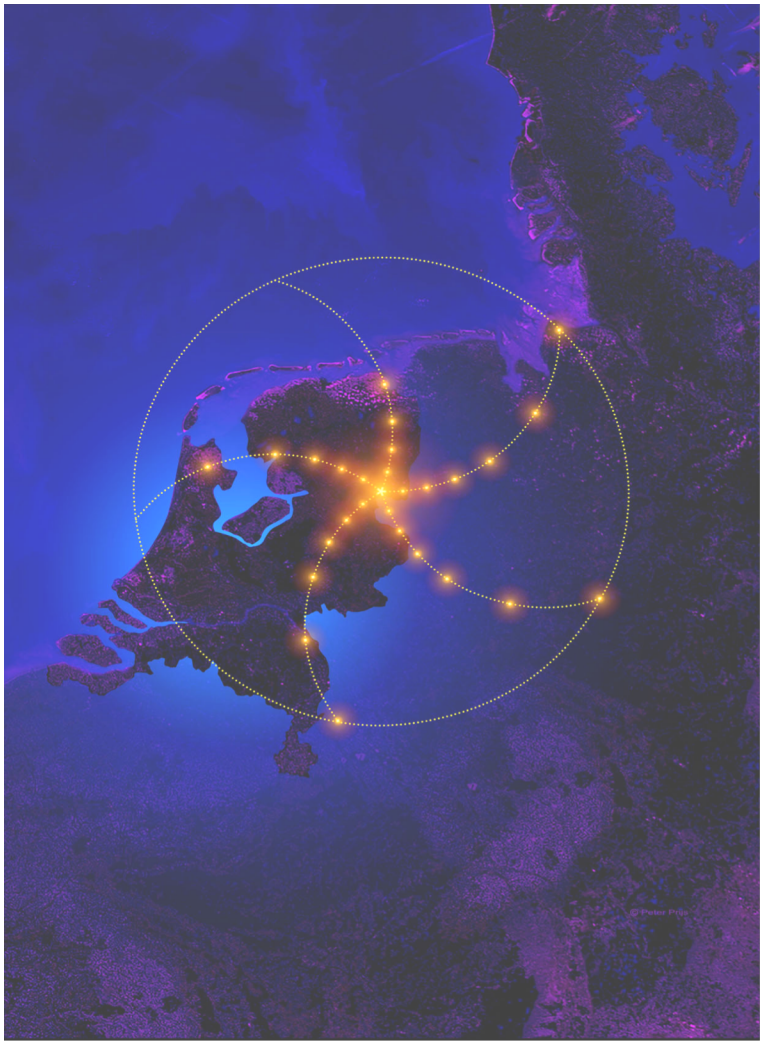

www.lofar.org

1 - 45 Tbit/s,

http://www.lofar.org/p/systems.htm

http://web.haystack.mit.edu/lofar/technical.html

SURFnet

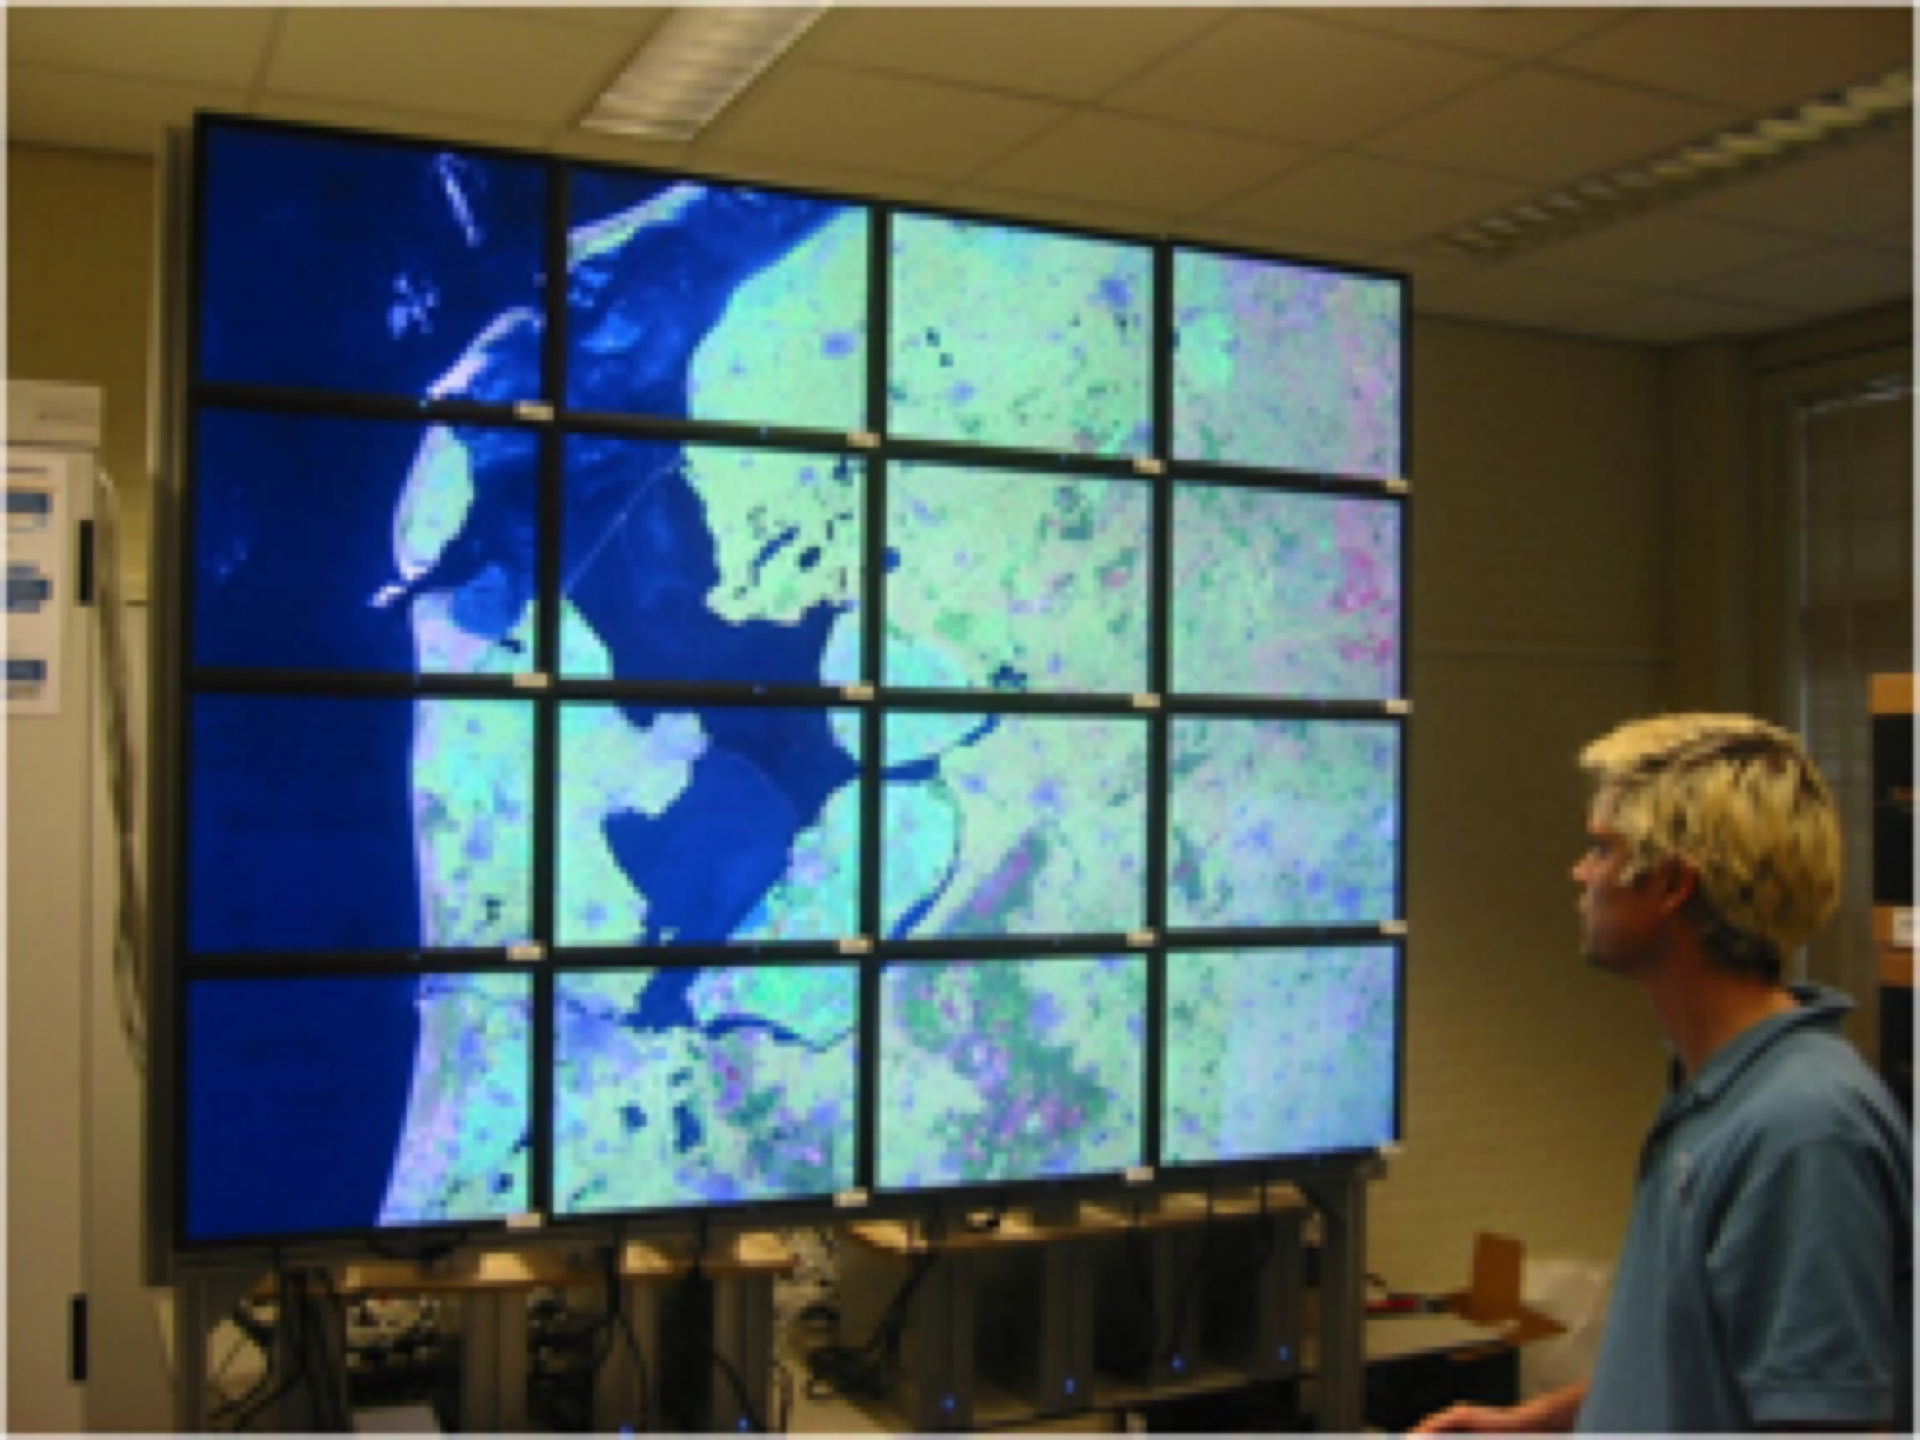

Grids

Showed you:

- **Computational Grids**
  - **HEP and LOFAR analysis requires massive CPU capacity**
- **Data Grid**
  - **Storing and moving HEP, Bio and Health data sets is major challenge**
- **Instrumentation Grids**
  - **Several massive data sources are coming online**
- **Visualization Grids**
  - **Data object (TByte sized) inspection, anywhere, anytime**

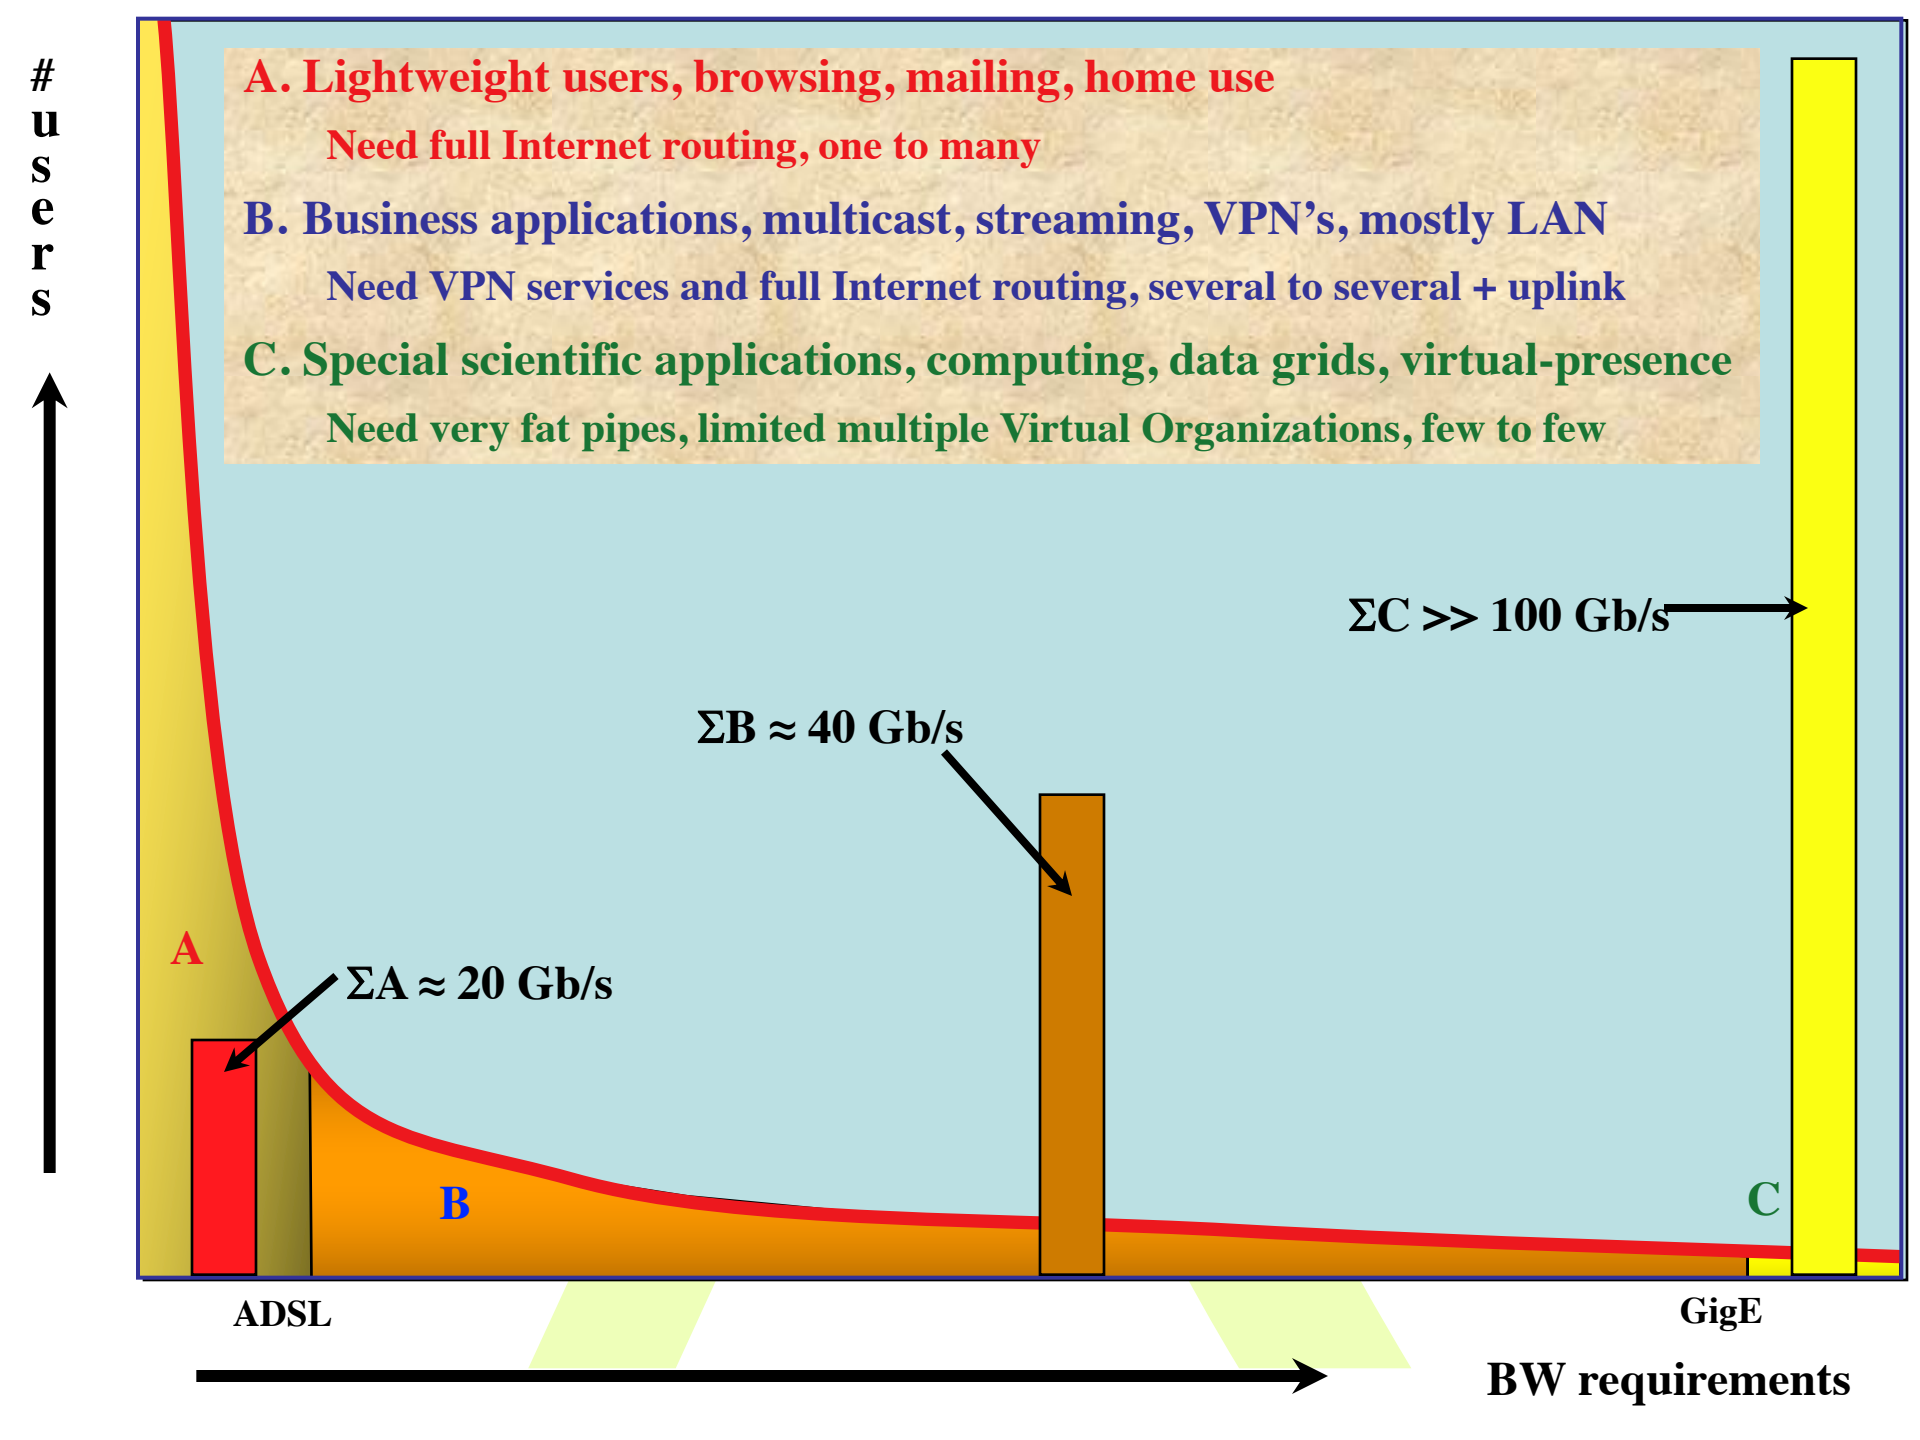
#users
A. Lightweight users, browsing, mailing, home use
Need full Internet routing, one to many
B. Business applications, multicast, streaming, VPN's, mostly LAN
Need VPN services and full Internet routing, several to several + uplink
C. Special scientific applications, computing, data grids, virtual-presence
Need very fat pipes, limited multiple Virtual Organizations, few to few
ΣC >> 100 Gb/s
ΣB ≈ 40 Gb/s
A
ΣA ≈ 20 Gb/s
B
C
ADSL
GigE
BW requirements

AMS-IX

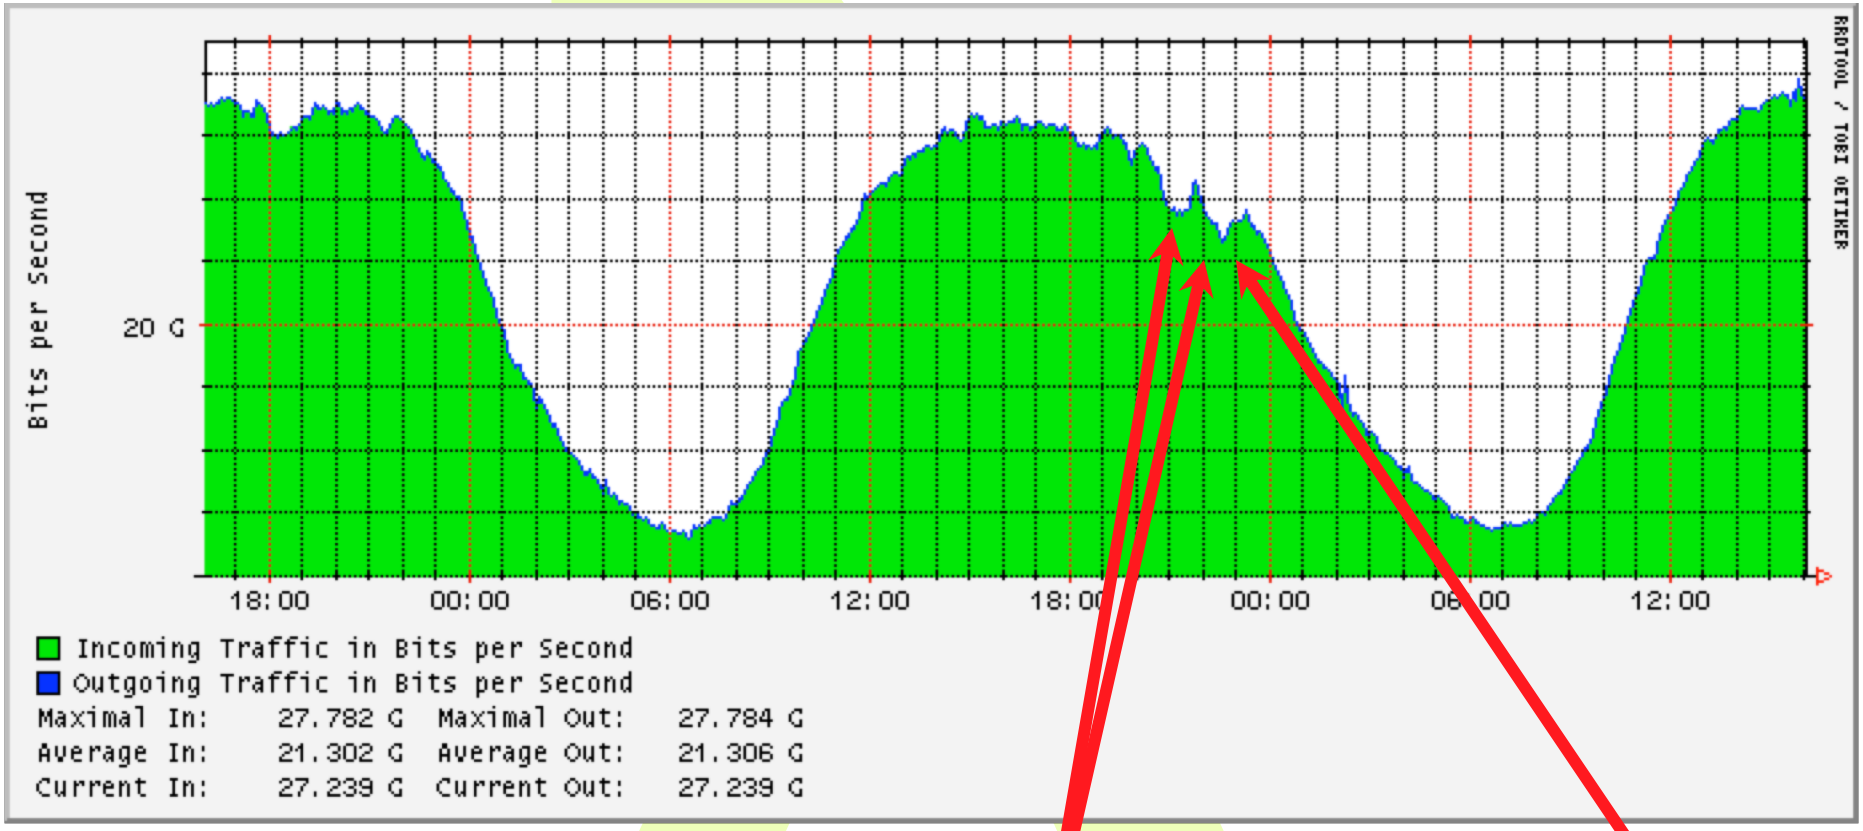

June 19th 2004

Lost :-(

European championship football Holland -- Czech Republic

So what?

- Costs of optical equipment 10% of switching 10 % of full routing equipment for same throughput
  - 10G routerblade -> 100-300 k$, 10G switch port -> 10-20 k$, MEMS port -> 0.7 k$
  - DWDM lasers for long reach expensive, 10-50k$ (???)
- Bottom line: look for a hybrid architecture which serves all classes in a cost effective way (A -> L3 , B -> L2 , C -> L1)
- Give each packet in the network the service it needs, but no more

L1 - 0.7 k$/port

L2 - 10-20 k$/port

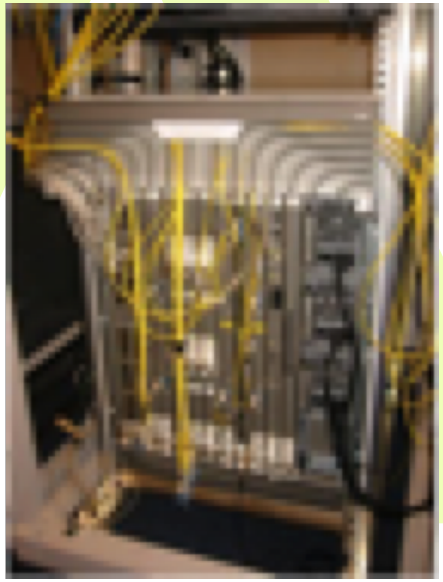

L3 - 100-300 k$/port

How low can you go?

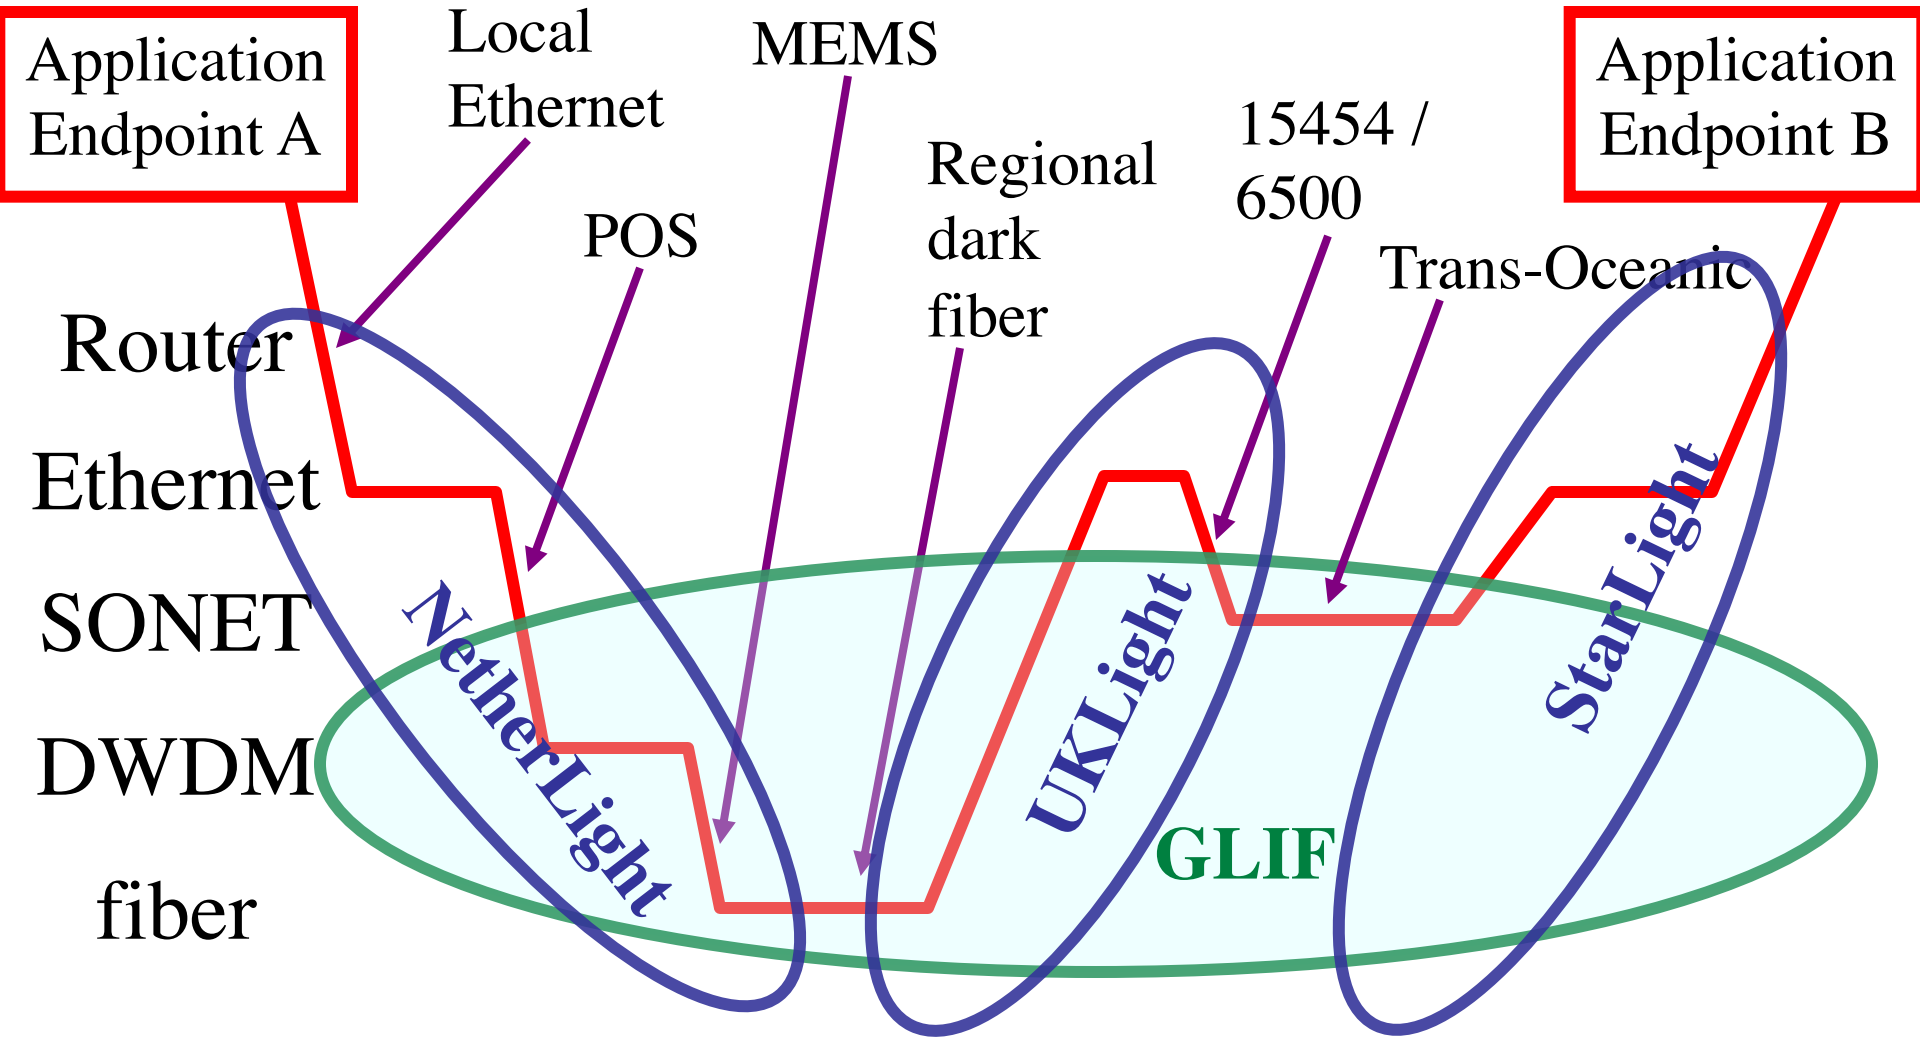

Optical Exchange as Black Box

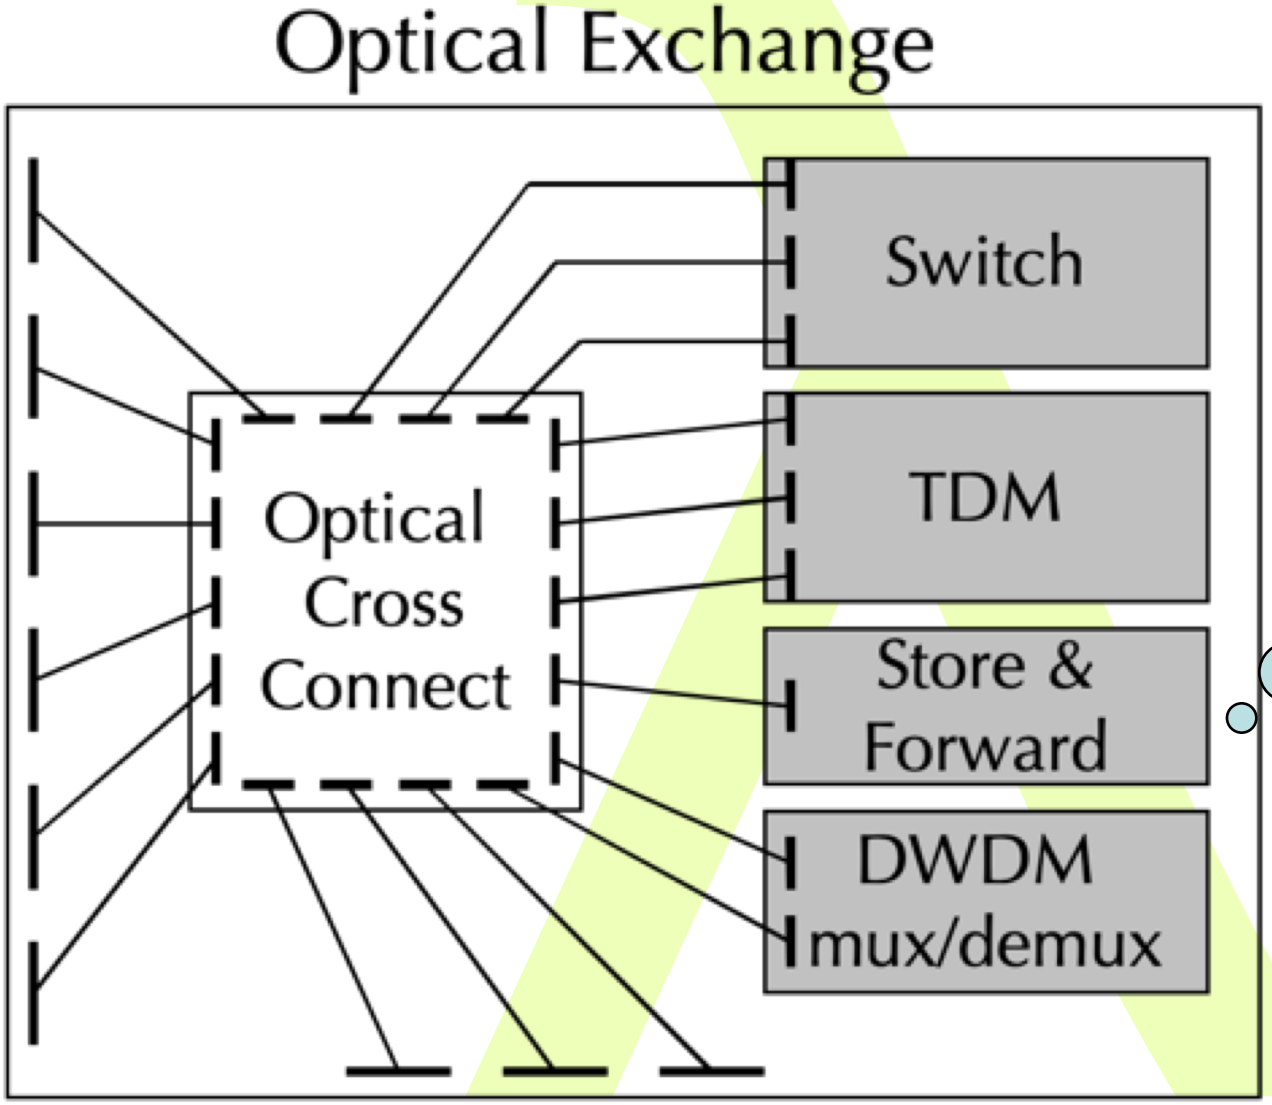

TeraByte
Email
Service

GLIF: Global Lambda Integrated Facility

- Established at the 3rd Lambda Grid Workshop, August 2003 in Reykjavik, Iceland
- Collaborative initiative among worldwide NRENs, institutions and their users
- A world-scale Lambda-based Laboratory for application and middleware development

GLIF vision: To build a new grid-computing paradigm, in which the central architectural element is optical networks, not computers, to support this decade's most demanding e-science applications.

Coordinated by UvA, SURFnet and UIC

GLIF Q3 2004

Visualization courtesy of
Bob Patterson, NCSA.

SURFnet fibers

(pict outdated anytime ;-)

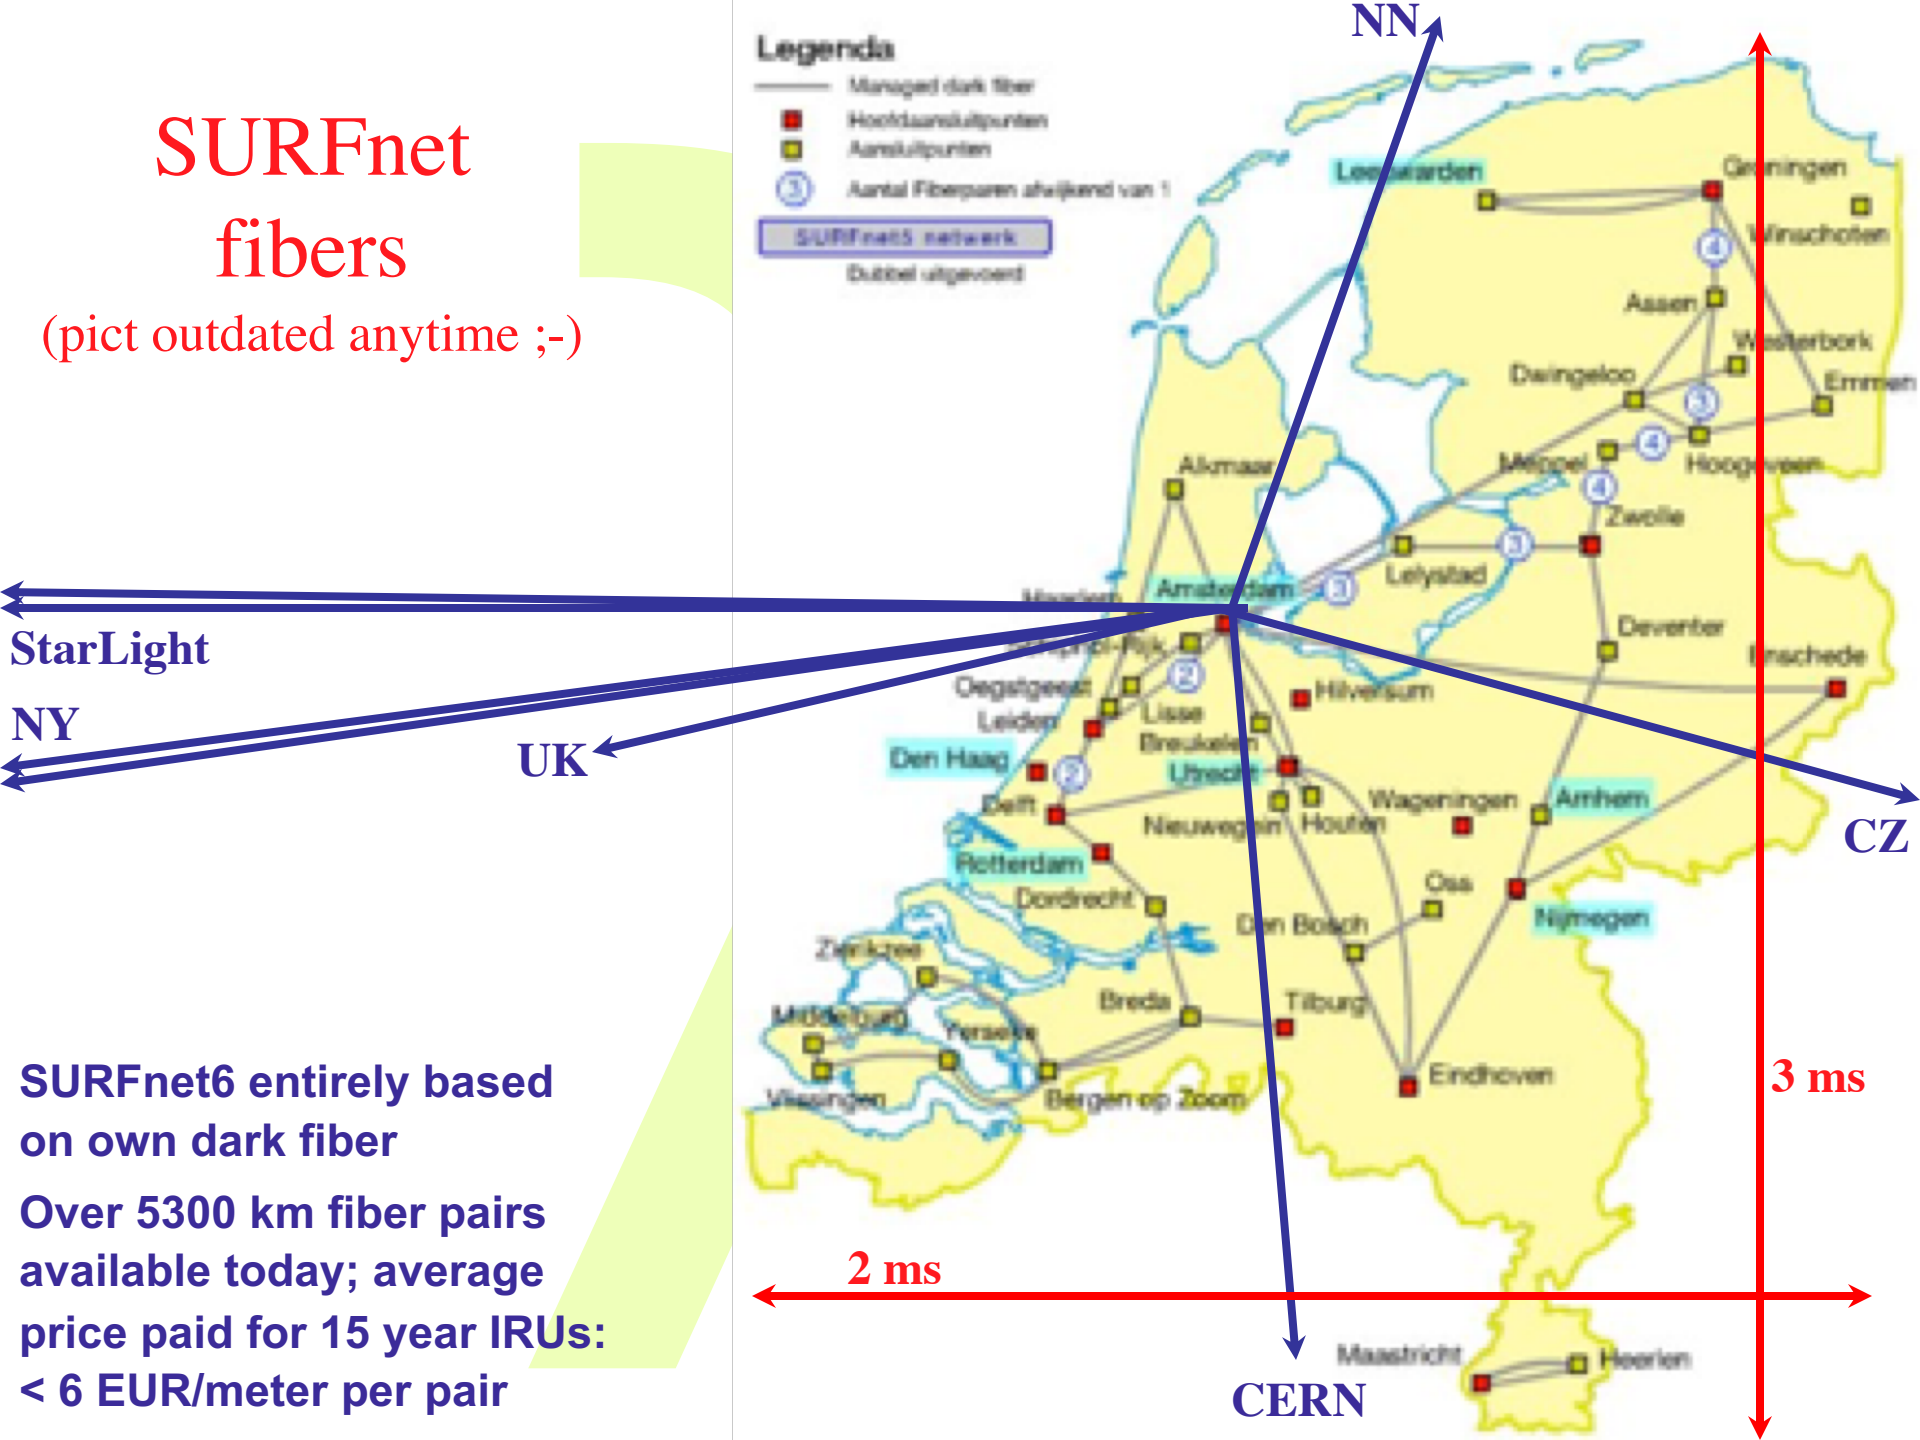

LightHouse

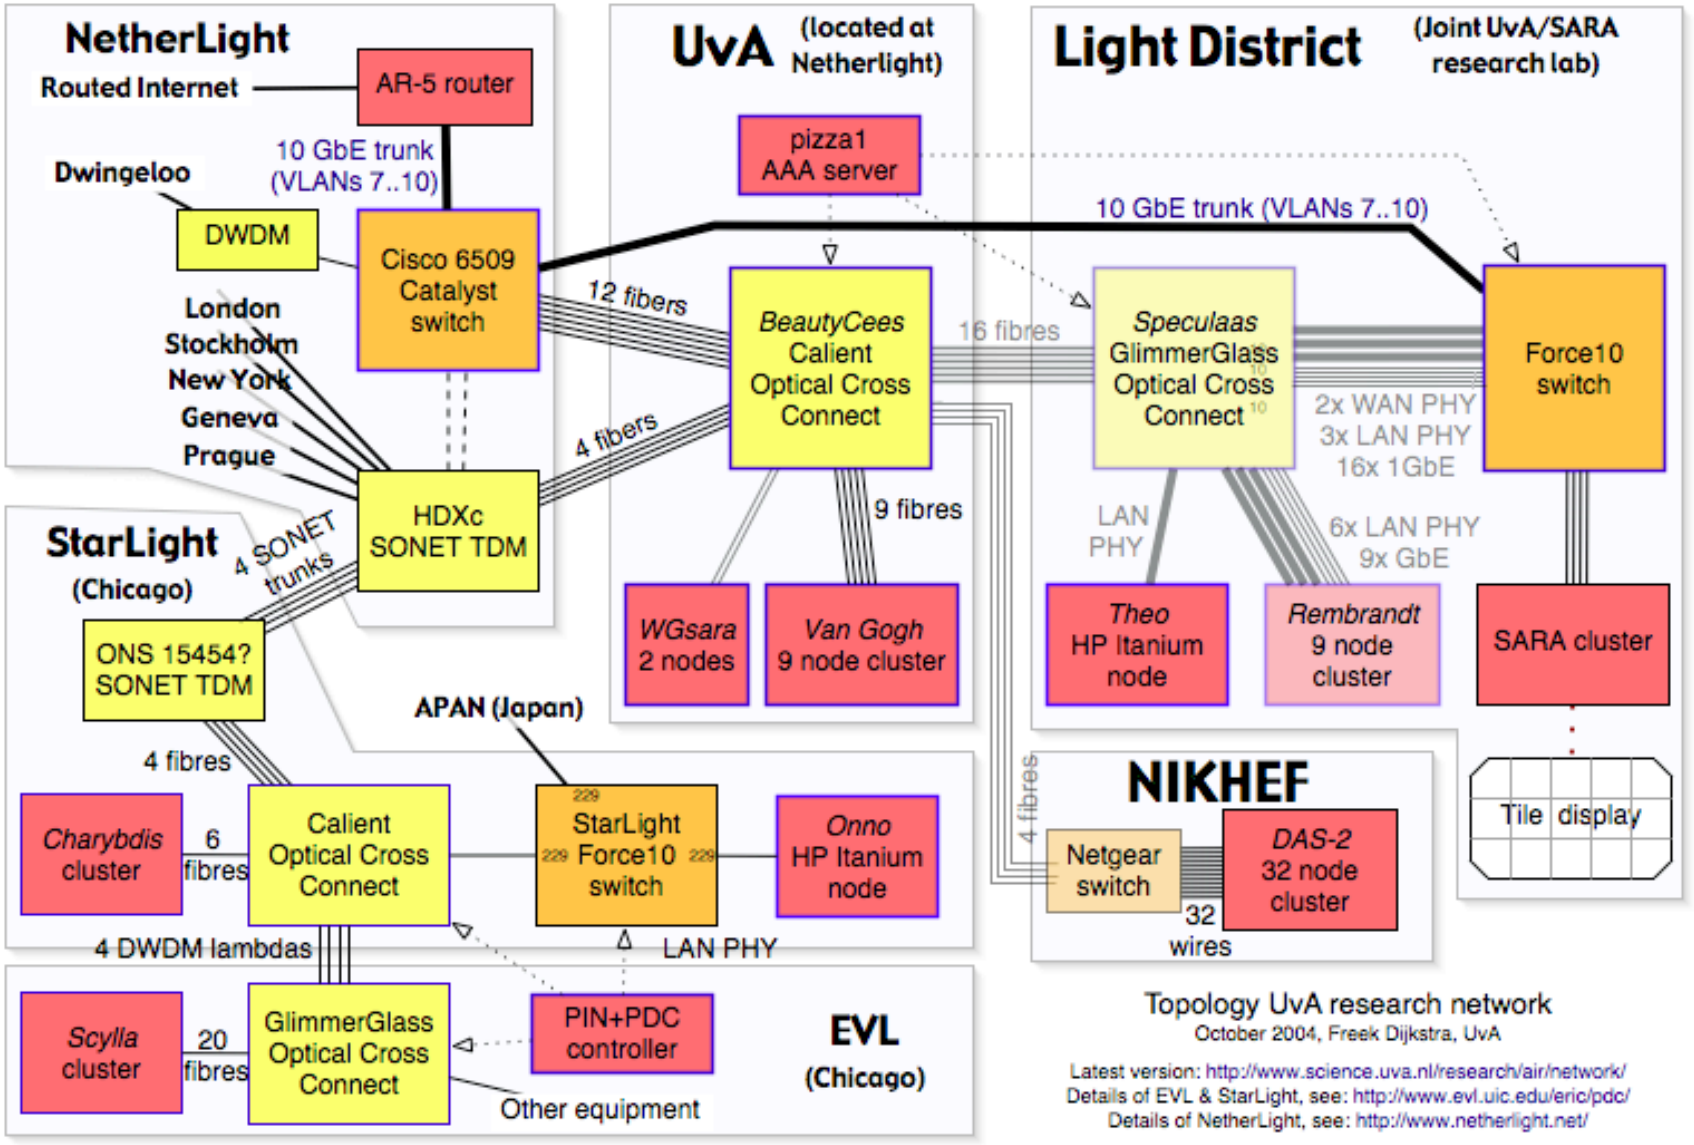
NetherLight
Routed Internet
AR-5 router
10 GbE trunk (VLANs 7..10)
Dwingeloo
DWDM
Cisco 6509 Catalyst switch
London
Stockholm
New York
Geneva
Prague
12 fibers
4 fibers
HDXc SONET TDM
StarLight (Chicago)
4 SONET trunks
ONS 15454? SONET TDM
4 fibres
Charybdis cluster
6 fibres
Calient Optical Cross Connect
4 DWDM lambdas
Scylla cluster
20 fibres
GlimmerGlass Optical Cross Connect
PIN+PDC controller
Other equipment
EVL (Chicago)
APAN (Japan)
StarLight Force10 switch
LAN PHY
Onno HP Itanium node
UvA (located at Netherlight)
pizza1 AAA server
BeautyCees Calient Optical Cross Connect
9 fibres
WGsara 2 nodes
Van Gogh 9 node cluster
10 GbE trunk (VLANs 7..10)
16 fibres
Speculaas GlimmerGlass Optical Cross Connect
2x WAN PHY
3x LAN PHY
16x 1GbE
LAN PHY
6x LAN PHY
9x GbE
Theo HP Itanium node
Rembrandt 9 node cluster
Light District (Joint UvA/SARA research lab)
Force10 switch
SARA cluster
Tile display
4 fibres
NIKHEF
Netgear switch
32 wires
DAS-2 32 node cluster
Topology UvA research network
October 2004, Freek Dijkstra, UvA
Latest version: http://www.science.uva.nl/research/air/network/
Details of EVL & StarLight, see: http://www.evl.uic.edu/eric/pdc/
Details of NetherLight, see: http://www.netherlight.net/

Link to web site

http://pb-cdl.local./~delaat/sc2004/

Transport in the corners

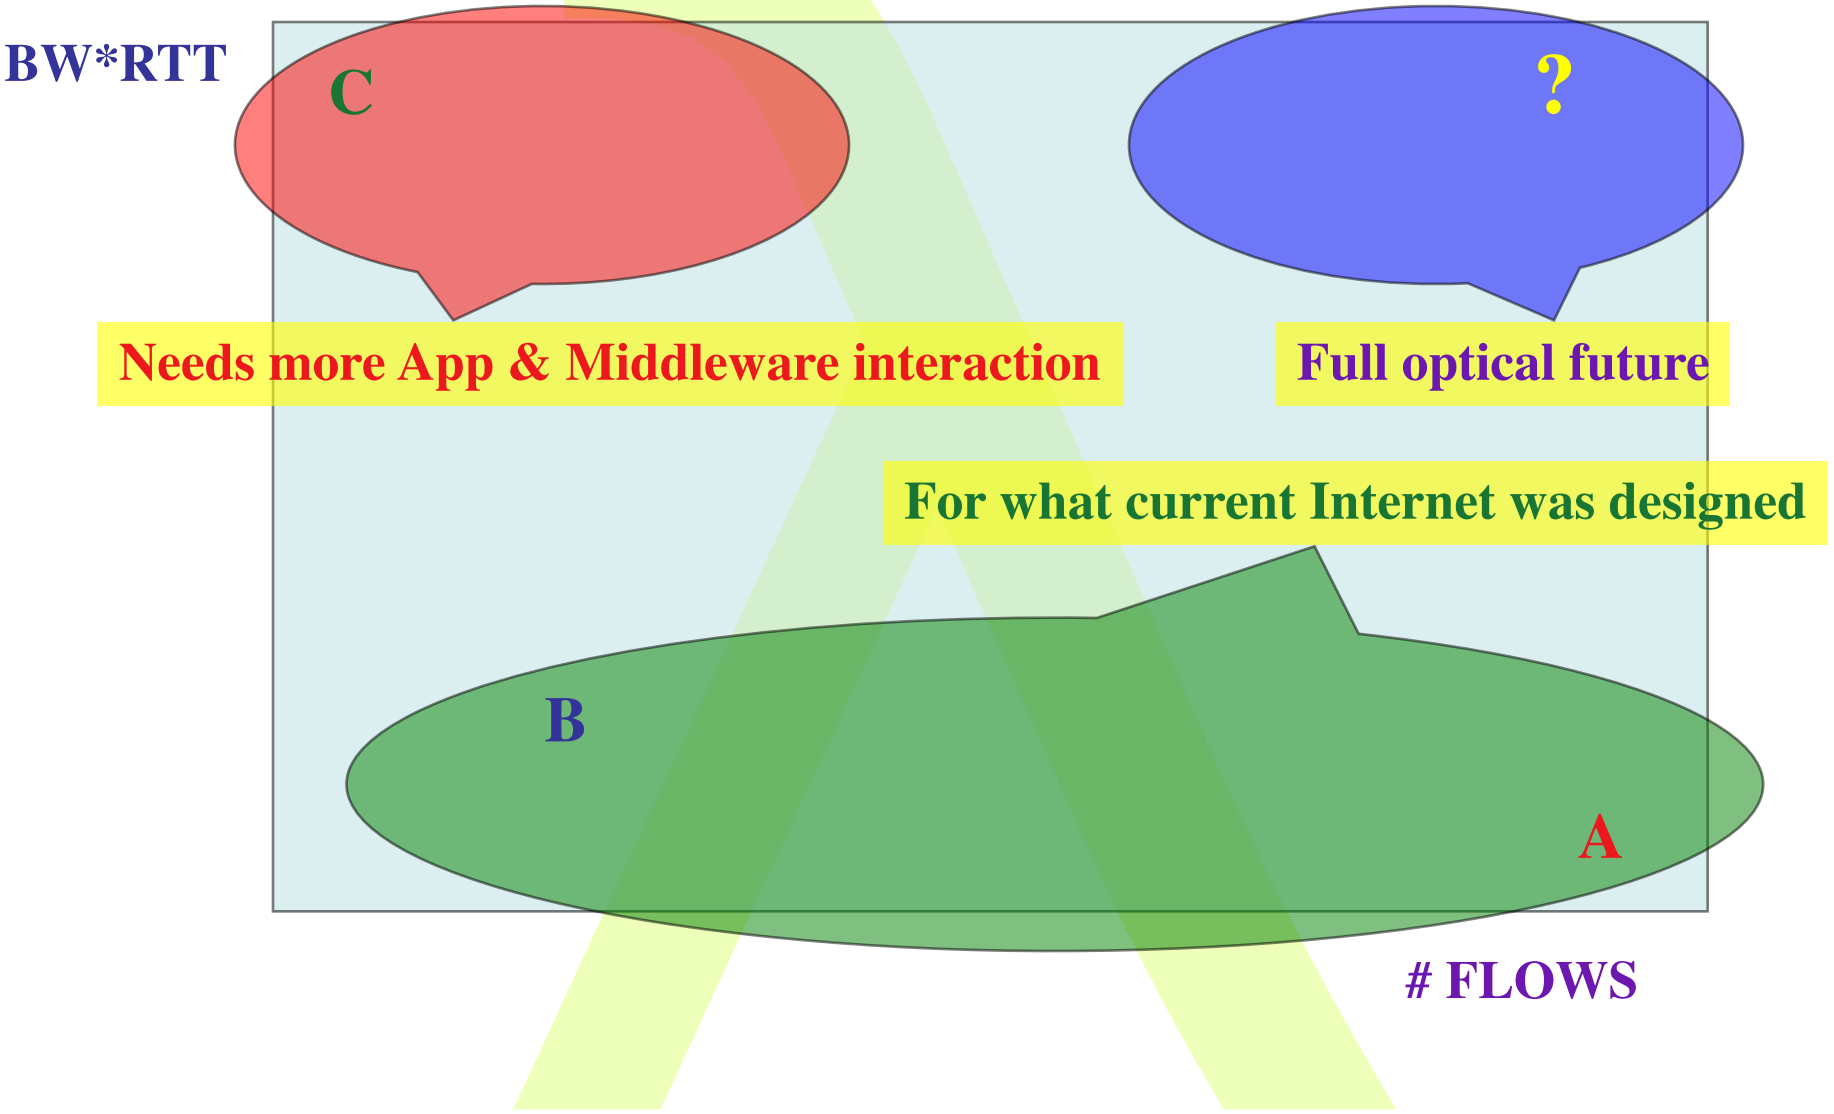

Not quite The END

Thanks to

SURFnet: Kees Neggers,UIC&iCAIR: Tom DeFanti, Joel Mambretti, CANARIE: Bill St. Arnaud

Freek Dijkstra, Hans Blom, Leon Gommans, Bas van oudenaarde, Arie Taal, Pieter de Boer, Bert Andree, Martijn de Munnik, Antony Antony, Rob Meijer, VL-team.

Partially complete list:

Caas

Chase

Cess

Kess

Case

The END

Thanks to
SURFnet: Kees Neggers,UIC&iCAIR: Tom DeFanti, Joel Mambretti, CANARIE: Bill St. Arnaud
Freek Dijkstra, Hans Blom, Leon Gommans, Bas van oudenaarde, Arie Taal, Pieter de Boer, Bert Andree, Martijn de Munnik, Antony Antony, Rob Meijer, VL-team.

Partially complete list:

Caas
Chase
Cess
Kess
Case

[1957-2004]