Advances Toward Economic and Efficient Terabit LANs and WANs

Cees de Laat

Advanced Internet Research Group (AIRG)

University of Amsterdam

News Flash:

You are sitting on a quarter terabit Lan!

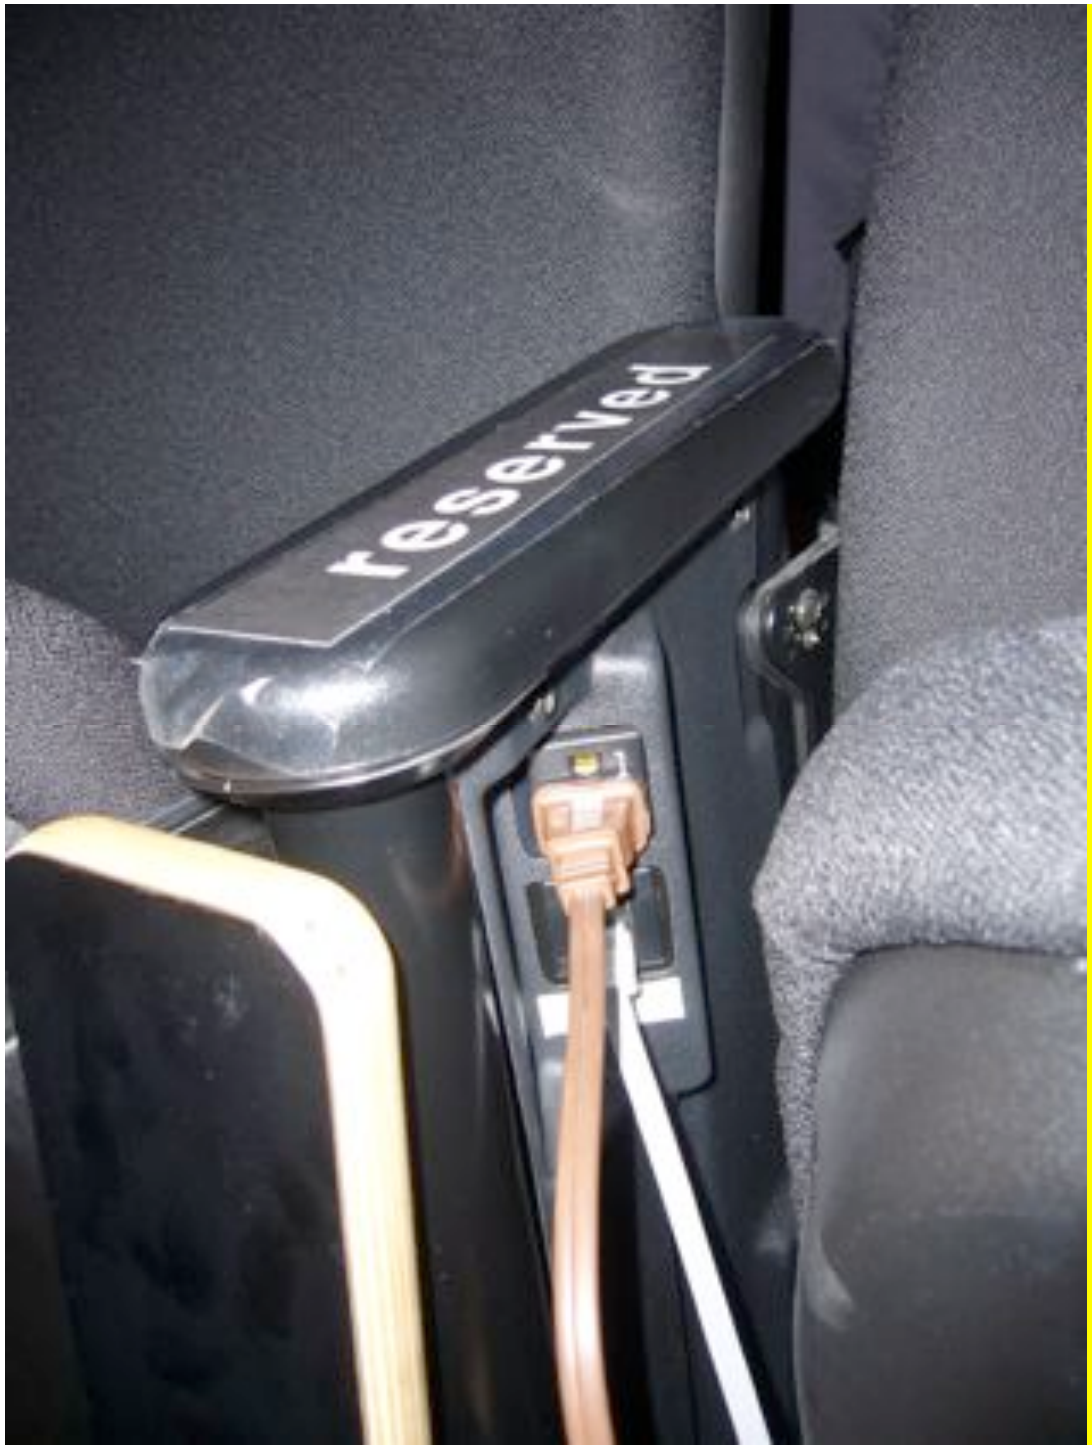

Check your seats!

You have a network connection of 1 Gigabit/s in your chair!

250 seats * 1 Gb/s -> 0.25 Tb/s

How to define a Terabit/s LAN?

- Depends where you put the boundary!
  - Backplane of a switch?
- What is the functionality
  - Many times point to point (dwdm/wss/mems)
  - Broadcast/multicast capabillity
  - Packet service versus streaming
- Application requirements
  - Packet services or streaming
- LAN-WAN interface
  - Addressing (mac -> ip)
  - Routing to the edge of the system
  - Integrate LAN and WAN

Terabits to the Desktop by 2010

- Simplified User View
- Terabit Fiber Connection To The Desktop
- Integrated Photonics And Electronics
- Single Fiber Dense-WDM
- Packets And Flows
- Encryption
- …

" Ethernet "	
1990	10 Mb
1995	100 Mb
1998	1 Gb
2002	10 Gb
2006	100 Gb
2008	1 Tb
2010	10 Tb

OptIPuter

Source: Steven Squires,
Chief Scientist HP

it2

OptIPuter is Prototyping The PC of 2010

- Terabits to the Desktop..
- 100 Megapixels Display
  - 55-Panel
- 1/3 Terabit/sec I/O
  - 30 x 10GE interfaces
  - Linked to OptIPuter
- 1/4 TeraFLOP
  - Driven by 30 Node Cluster of 64 bit Dual Opterons
- 1/8 TB RAM
- 60 TB Disk

Source: Jason Leigh, Tom DeFanti, EVL@UIC
OptIPuter Co-PIs

So, is a Terabit LAN just a bigger ethernet switch?

If yes:
just wait and buy!

intermezzo

Low power GigaBit/s @ home

- 5 years ago Gigabit/s networking was a master student project
- Don’t forget power consumption www.delaat.net @ 28 Watt

So, is a Terabit LAN just a bigger ethernet switch?

If yes:
just wait and buy!

Why no:
Grid -> collective services!
Anti-statistical behavior!

Towards Hybrid Networking!

- Costs of optical equipment 10% of switching 10 % of full routing equipment for same throughput
  - 10G routerblade -> 75-300 k$, 10G switch port -> 5-10 k$, MEMS port -> 0.5-1.5 k$
  - DWDM lasers for long reach expensive, 10-50 k$
- Bottom line: look for a hybrid architecture which serves all classes in a cost effective way ==> map A -> L3 , B -> L2 , C -> L1
- Give each packet in the network the service it needs, but no more !

L1 ≈ 1 k$/port

L2 ≈ 5-10 k$/port

L3 ≈ 75+ k$/port

Cost of 1 Tb/s LAN

very rough estimate sept 2005

- L2
  - 10 gig nic ~ 1 k
  - Per L2 switch port ~5 k
  - Total 6 M
- L1
  - Per L1 port ~ 1 k
  - Total 2 M
- OOPS, I forgot the cables and the power bill

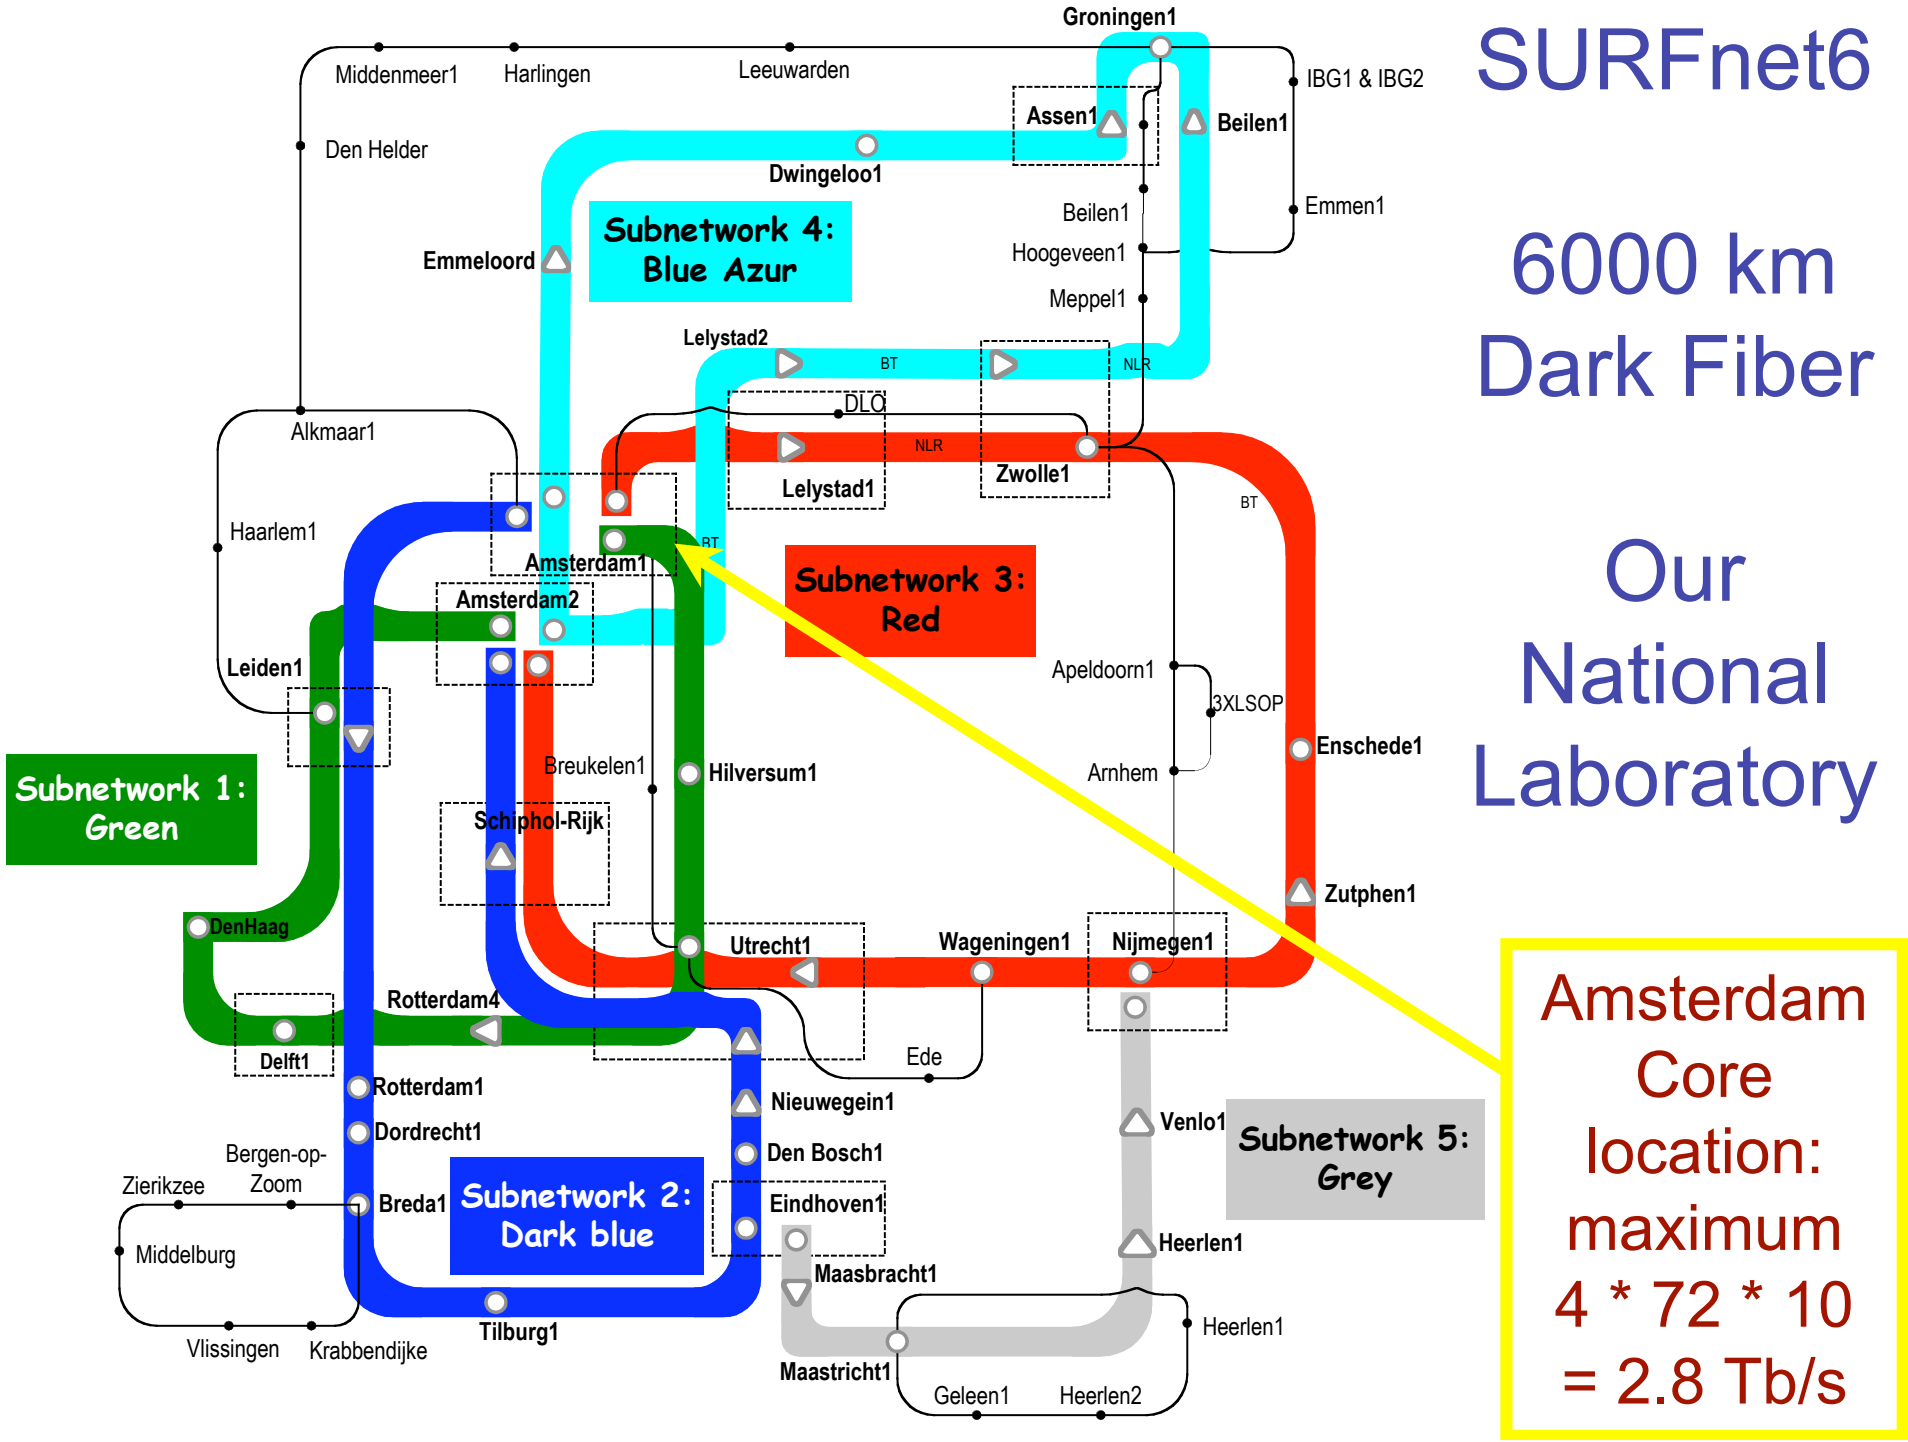

SURFnet6

6000 km Dark Fiber

Our National Laboratory

Amsterdam Core location: maximum 4 * 72 * 10 = 2.8 Tb/s

StarPlane: DWDM backplane for DAS-3

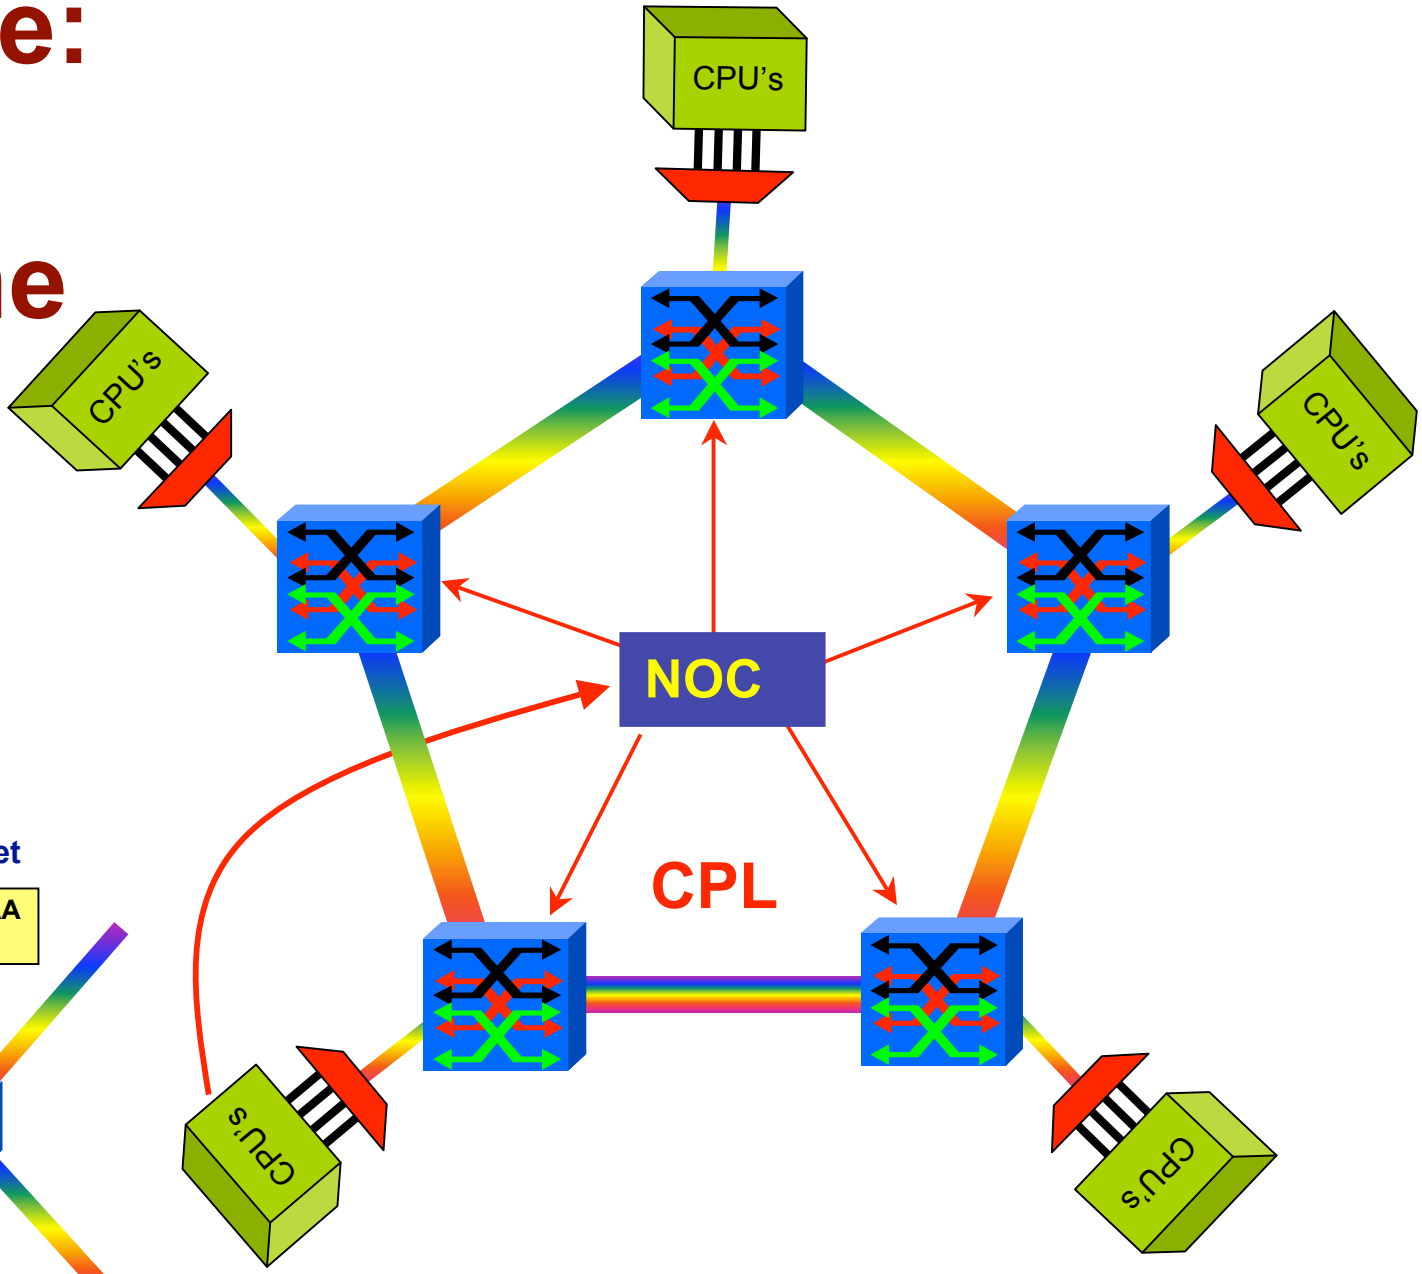

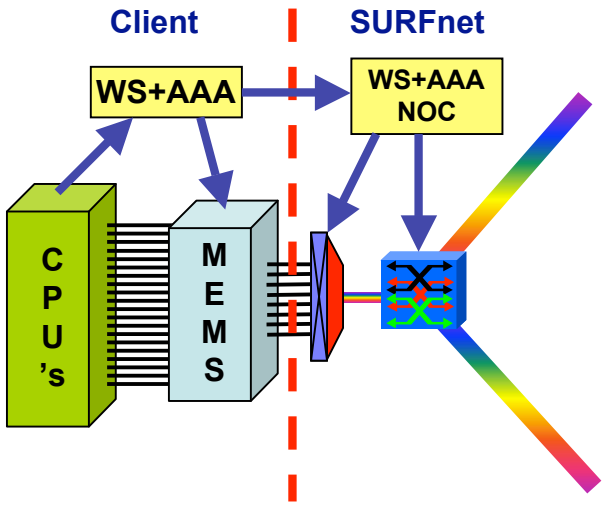

cdλ

Transport of flows

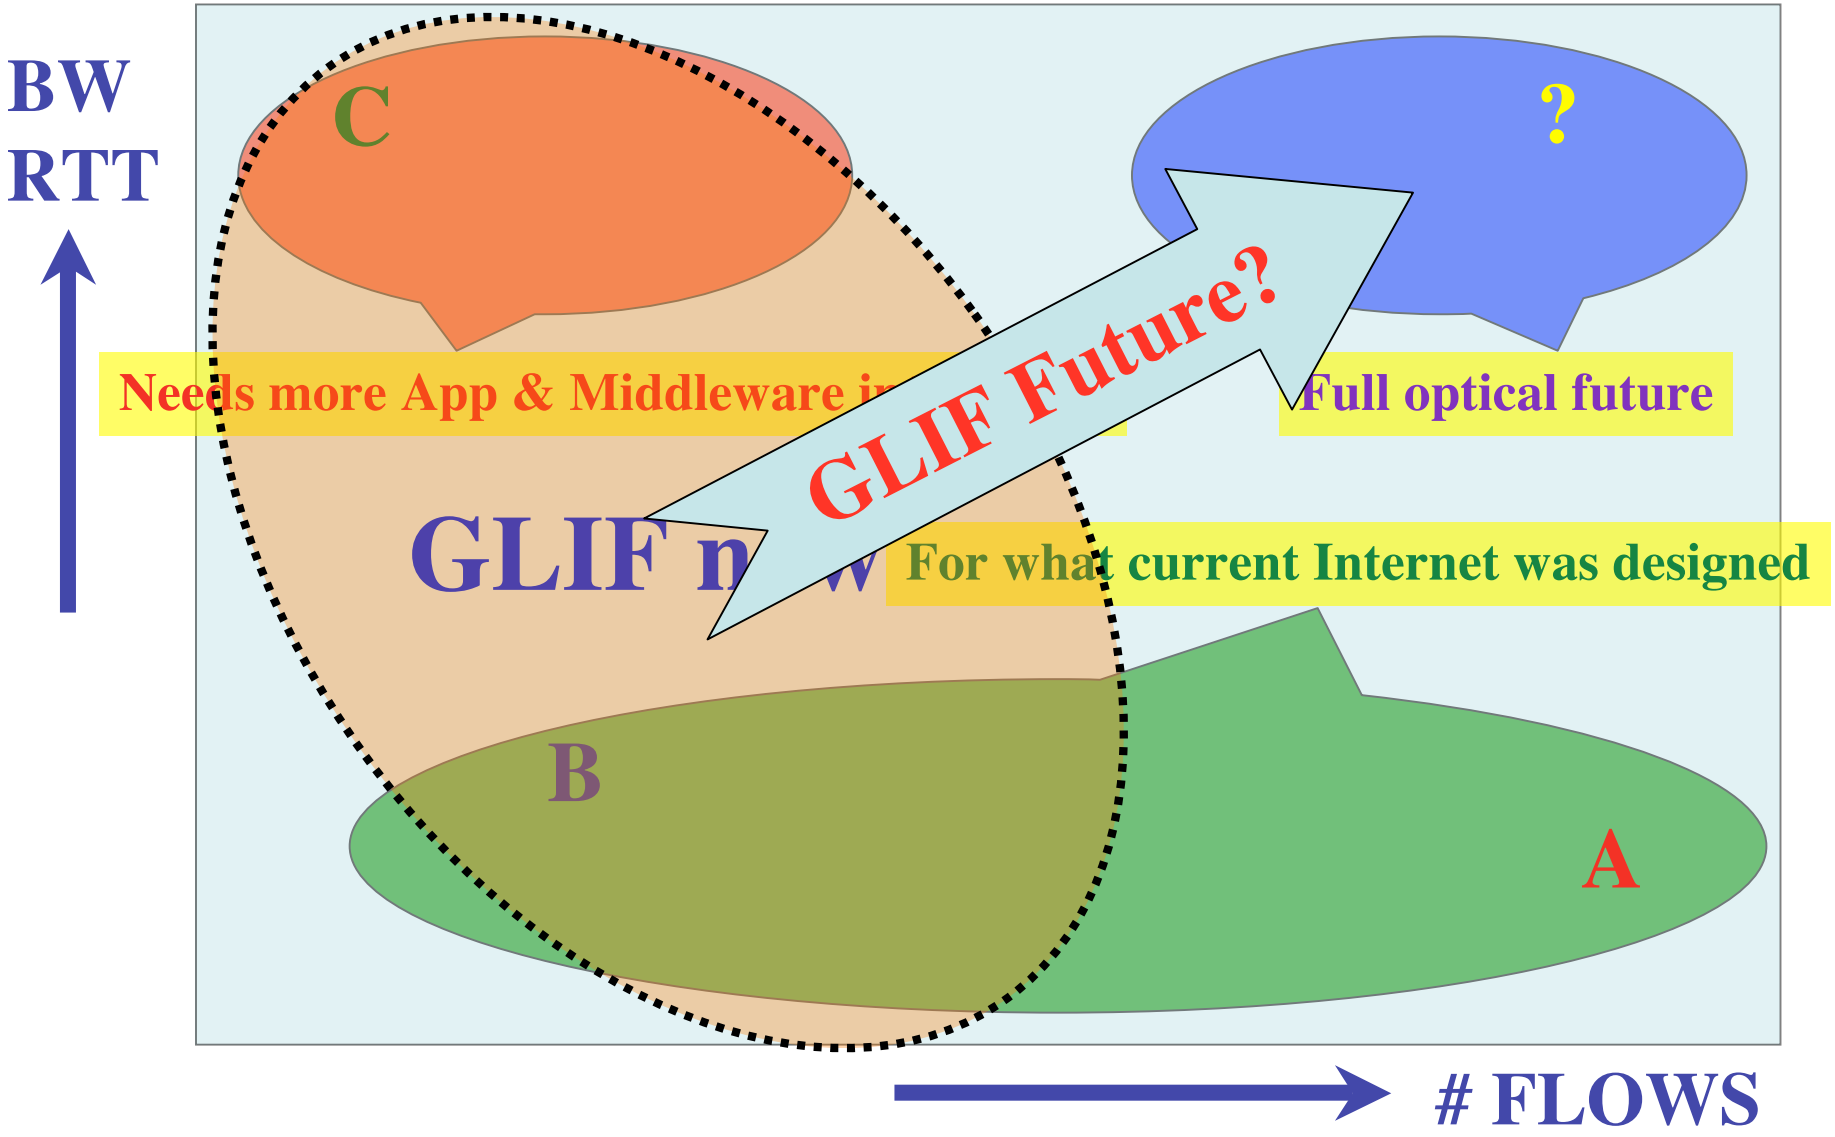

Questions ?

More info:

http://www.science.uva.nl/~delaat

delaat@uva.nl

iGRID2005 publication opportunity

"Future Generation Computer Systems (FGCS): The International Journal of Grid Computing: Theory, Methods and Applications" will publish a SPECIAL iGRID ISSUE in Spring/Summer 2006.

Guest editors: Larry Smarr, Tom DeFanti, Maxine Brown, Cees de Laat

We can accept around 20-25 papers, Papers will be reviewed

* Maximum paper length is limited to 8 pages
* Limit of 1 paper per demonstration.
* Describe your iGrid experiences, results and performance measurements.
* **DEADLINE for submission is ONE MONTH AFTER iGRID -> Oct 31.**

Submission must be via the FGCS website. For author guides and submission information, see <http://ees.elsevier.com/fgcs/>.

Contact: Cees de Laat delaat@science.uva.nl (need reviewers :-)