Dutch LambdaGrid state @ sc2005
(more Lambda than Grid :-)

Cees de Laat

University of Amsterdam

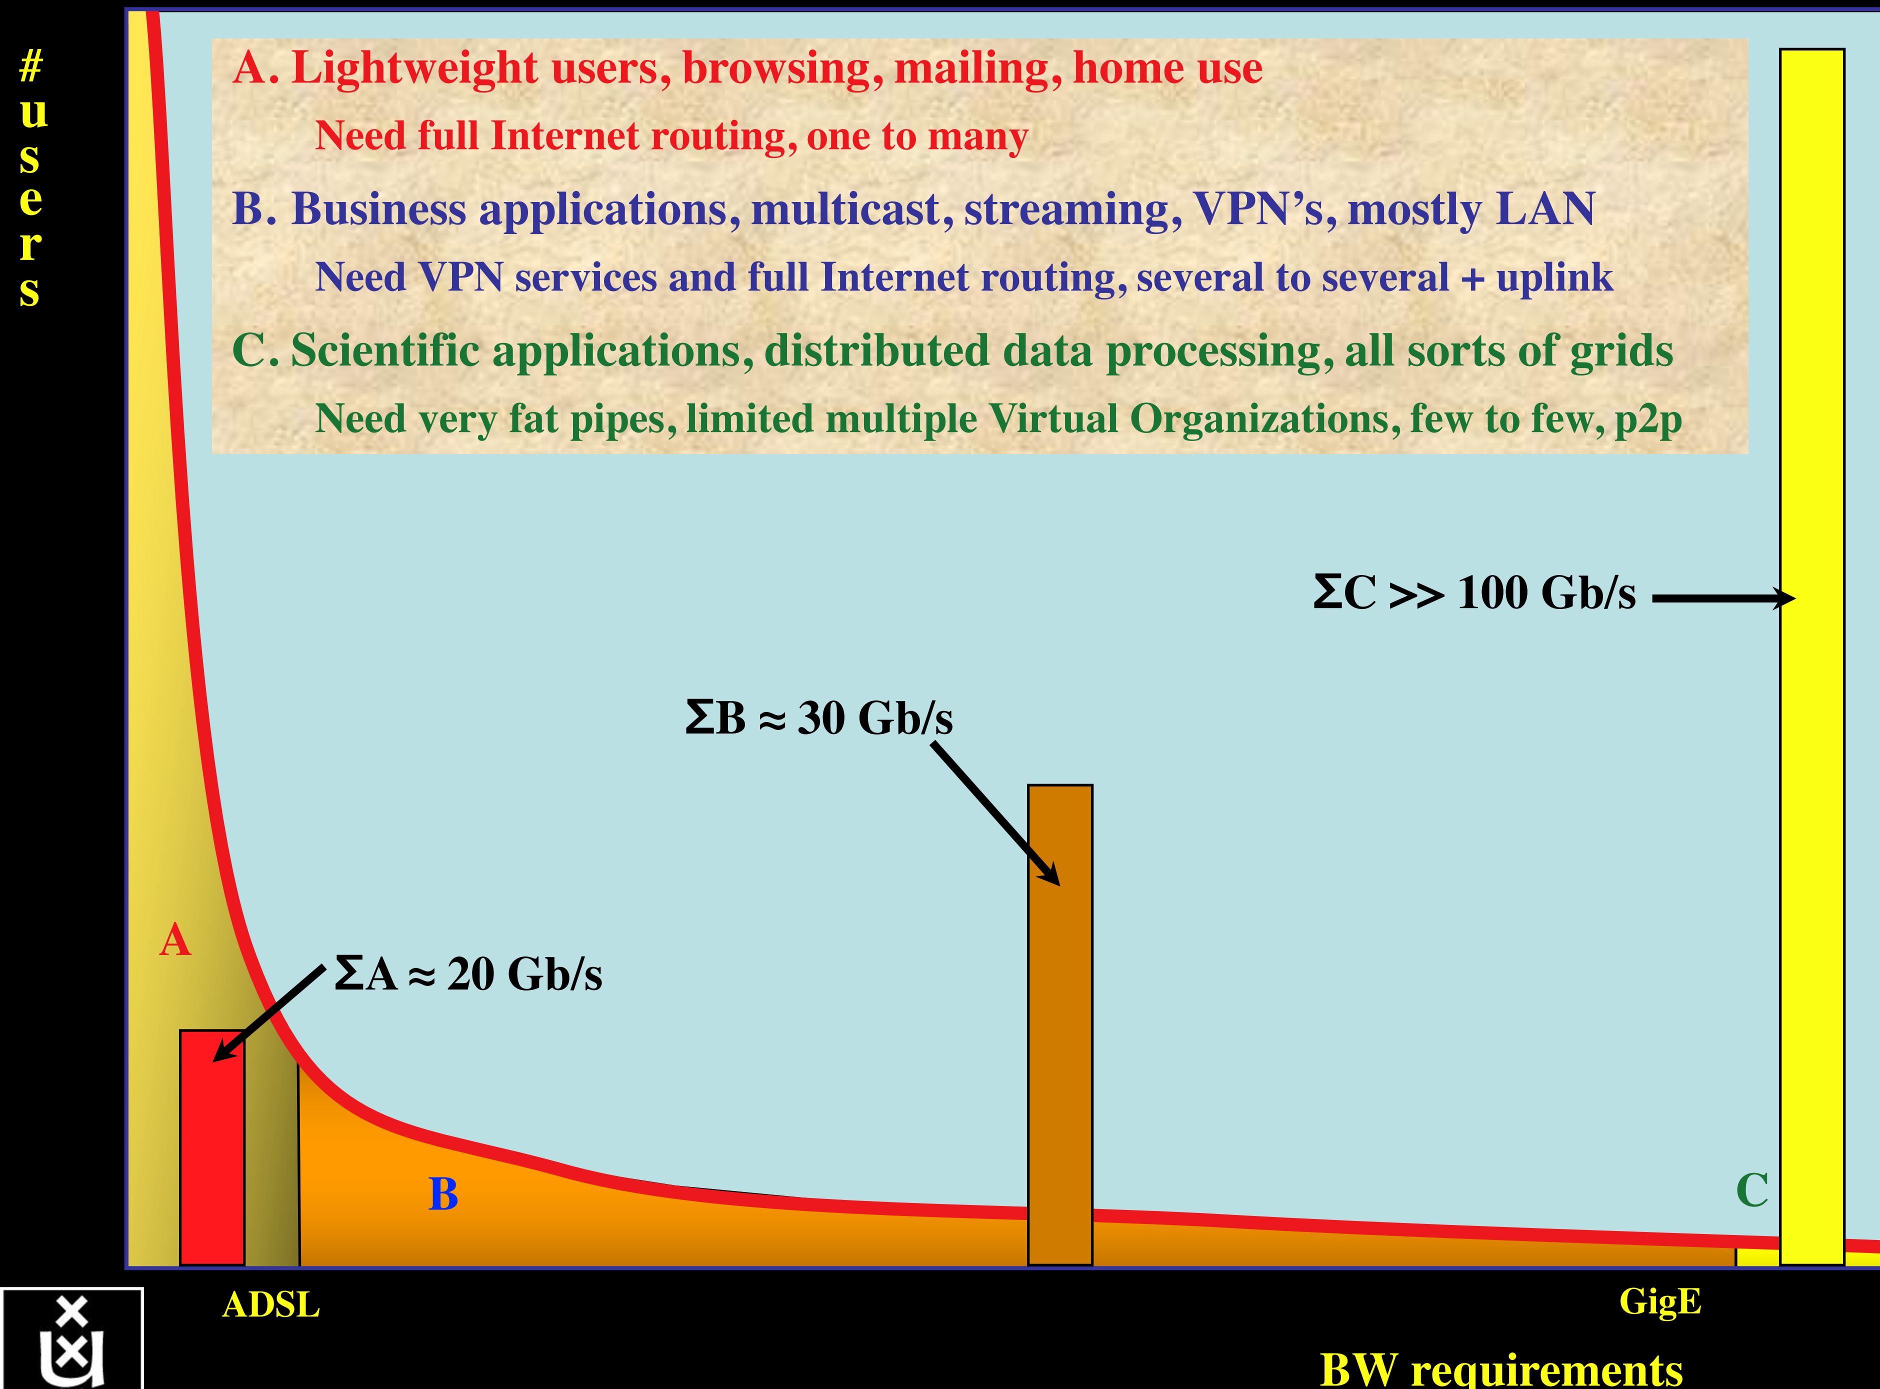
A. Lightweight users, browsing, mailing, home use
Need full Internet routing, one to many
B. Business applications, multicast, streaming, VPN's, mostly LAN
Need VPN services and full Internet routing, several to several + uplink
C. Scientific applications, distributed data processing, all sorts of grids
Need very fat pipes, limited multiple Virtual Organizations, few to few, p2p
users
ΣC >> 100 Gb/s
ΣB ≈ 30 Gb/s
ΣA ≈ 20 Gb/s
A
B
C
ADSL
GigE
BW requirements

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 ≈ 0.5-1.5 k$/port

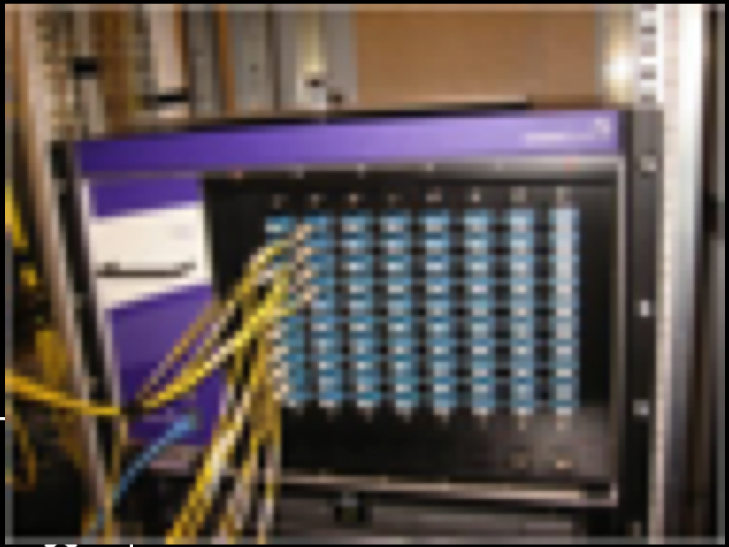

L2 ≈ 5-10 k$/port

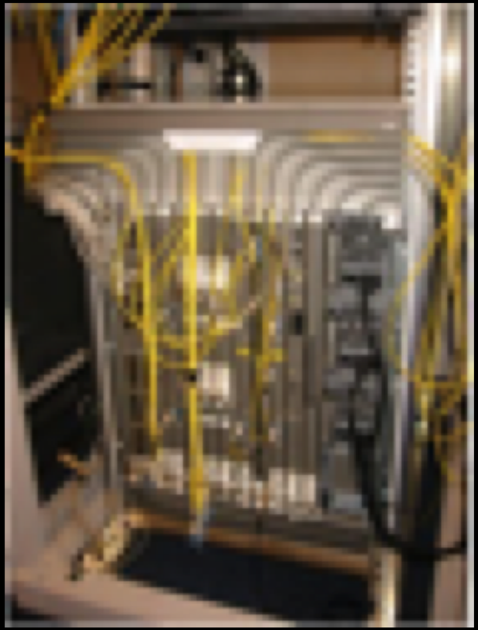

L3 ≈ 75+ k$/port

UVA's
64*64
Optical Switch
@ LightHouse

Costs 1/100th of a similar throughput router

or

1/10th of a similar throughput Ethernet switch

but has only specific services!

GLIMMERGLASS

How low can you go?

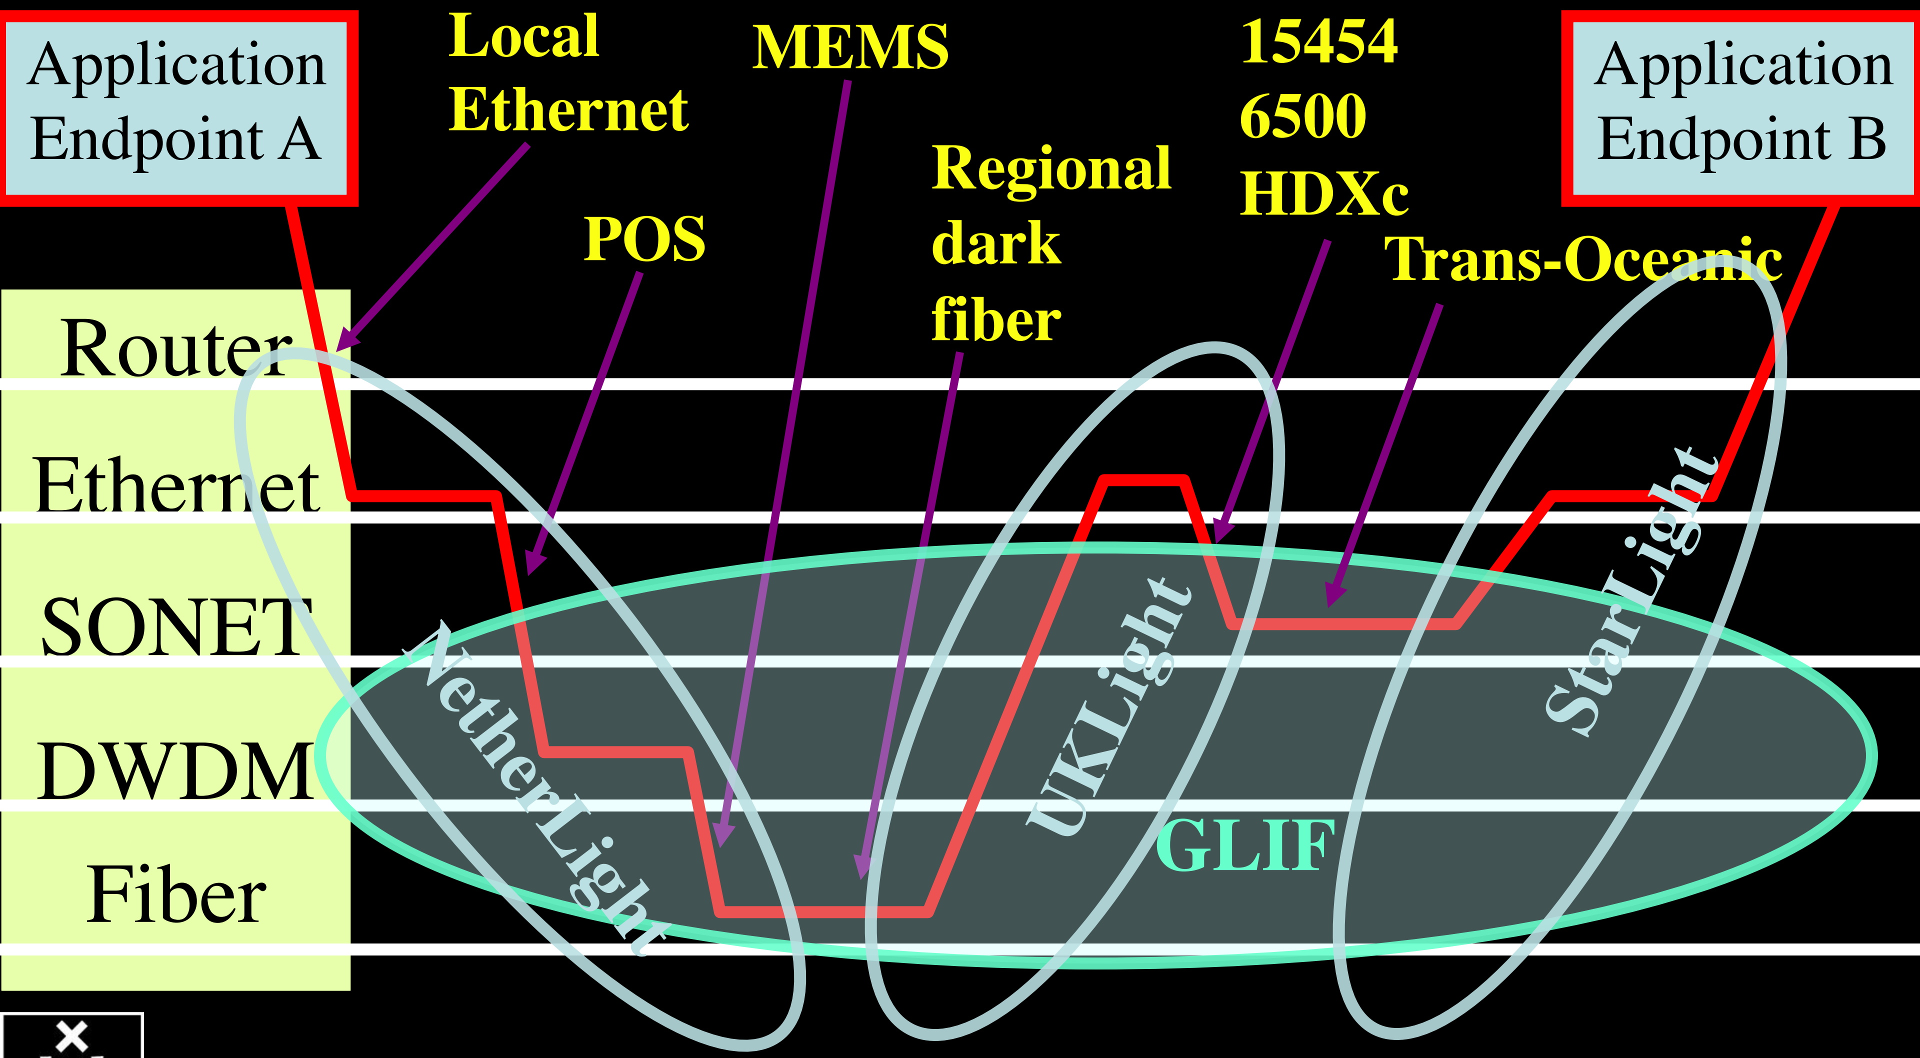

Token Based Networking

Access Control, Resource Management and Path Selection in Optical Networks using Tokens

Tokens performing Resource Management and Access Control in Virtual Machine Turntable Experiment.

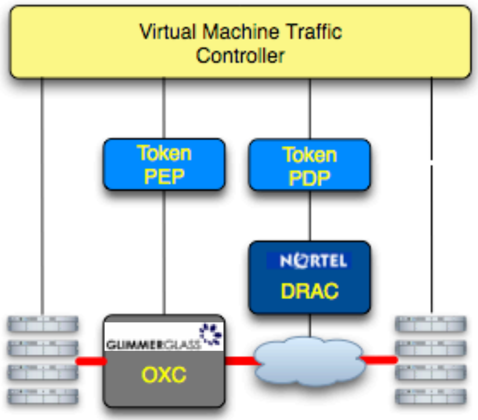

Tokens performing Path Selection and Access Control at Optical Inter-Connection Points

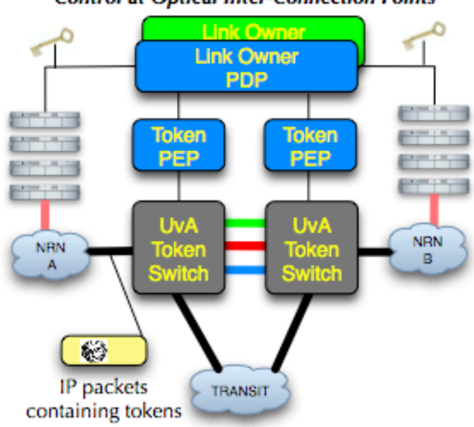

Tokens will allow:

- Separation of (slow) authorization process and real time usage.
- Binding to many different types of attributes: user, time, resource, etc.
- Policy Decision to be abstracted from Policy Enforcement Point.
- Anonymous usage
- Resource Management

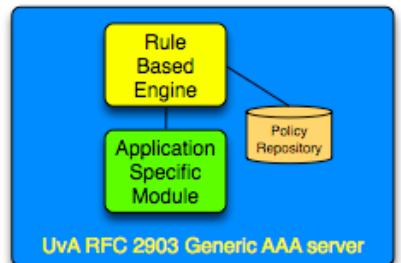

Token marked IP packets will allow:

- Economic Link Owners to assign usage rights without routing changes.
- Recognition at Inter-Connection Points (Optical Exchanges). When authentic and valid, token marked traffic will use the Link Owners path.
- Implementations that support different business models
- Hardware (NPU based) recognition rate expected to be a 10 Gb/s.

UvA Universiteit van Amsterdam
Leon Gommans
Fred Wan, Cees de Laat
Mihai Cristea
GigaPort NextGRID

SURFnet 6 principles

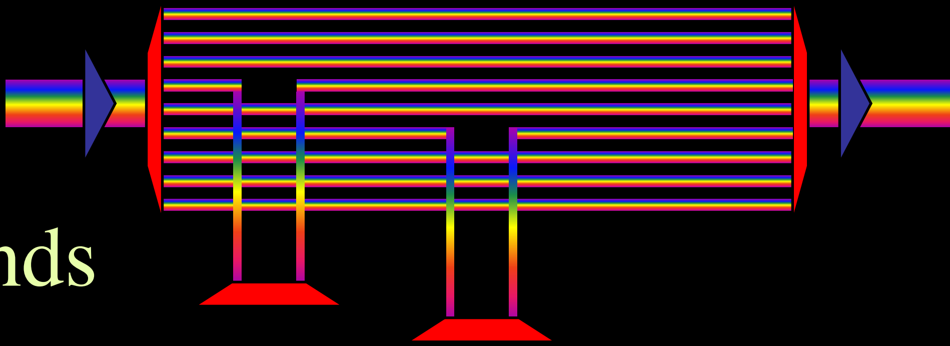

- Based on dark fiber
- 4 DWDM rings of 9 bands
  - each 4, later 8, colors
  - Each capable of 10, later 40 Gb/s
- Universities have POP's on ring, each 1 band
- Connect with 1 or 10 Gb/s Ethernet
- Routing in Amsterdam in 2 core POP's!
- International connectivity in Amsterdam
- Lambda service between ring POP's and to NetherLight

Common Photonic Layer (CPL) in SURFnet6

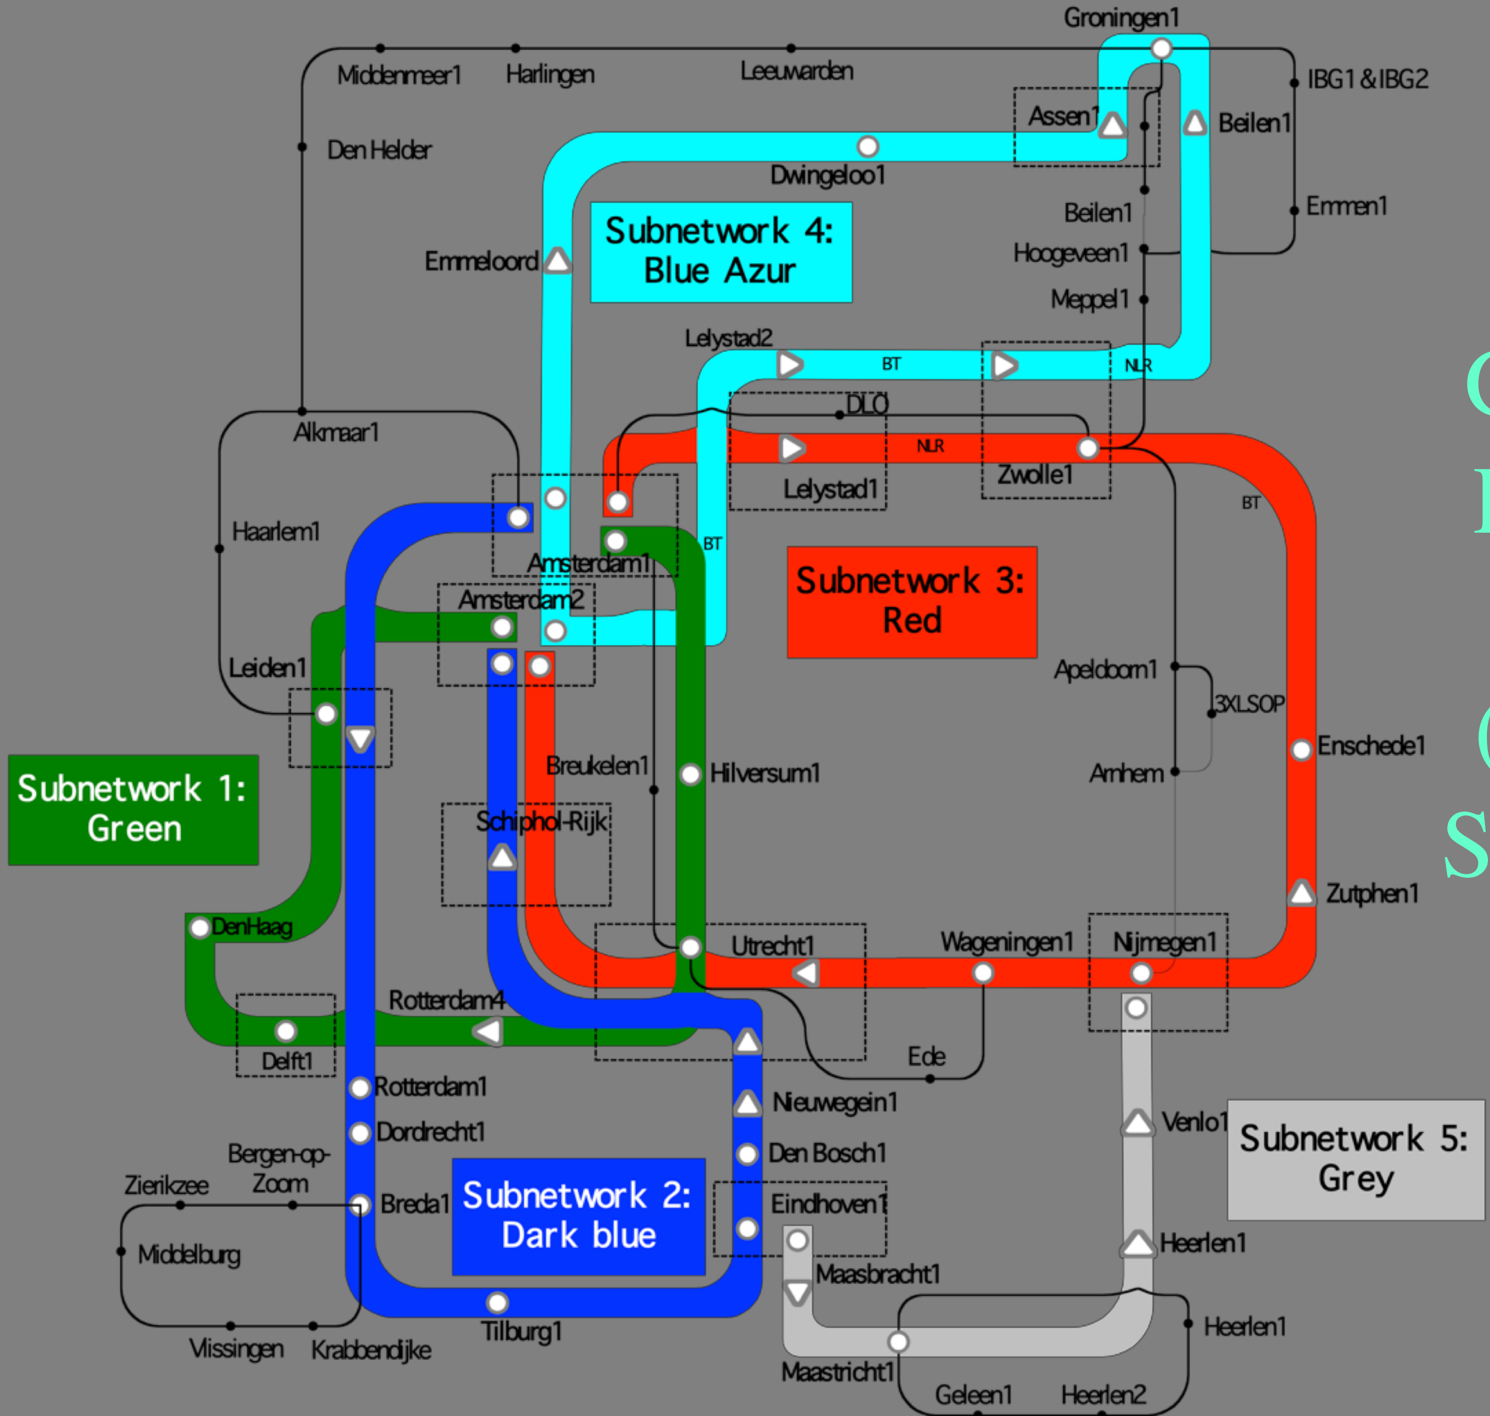

SURFnet

StarPlane DWDM backplane

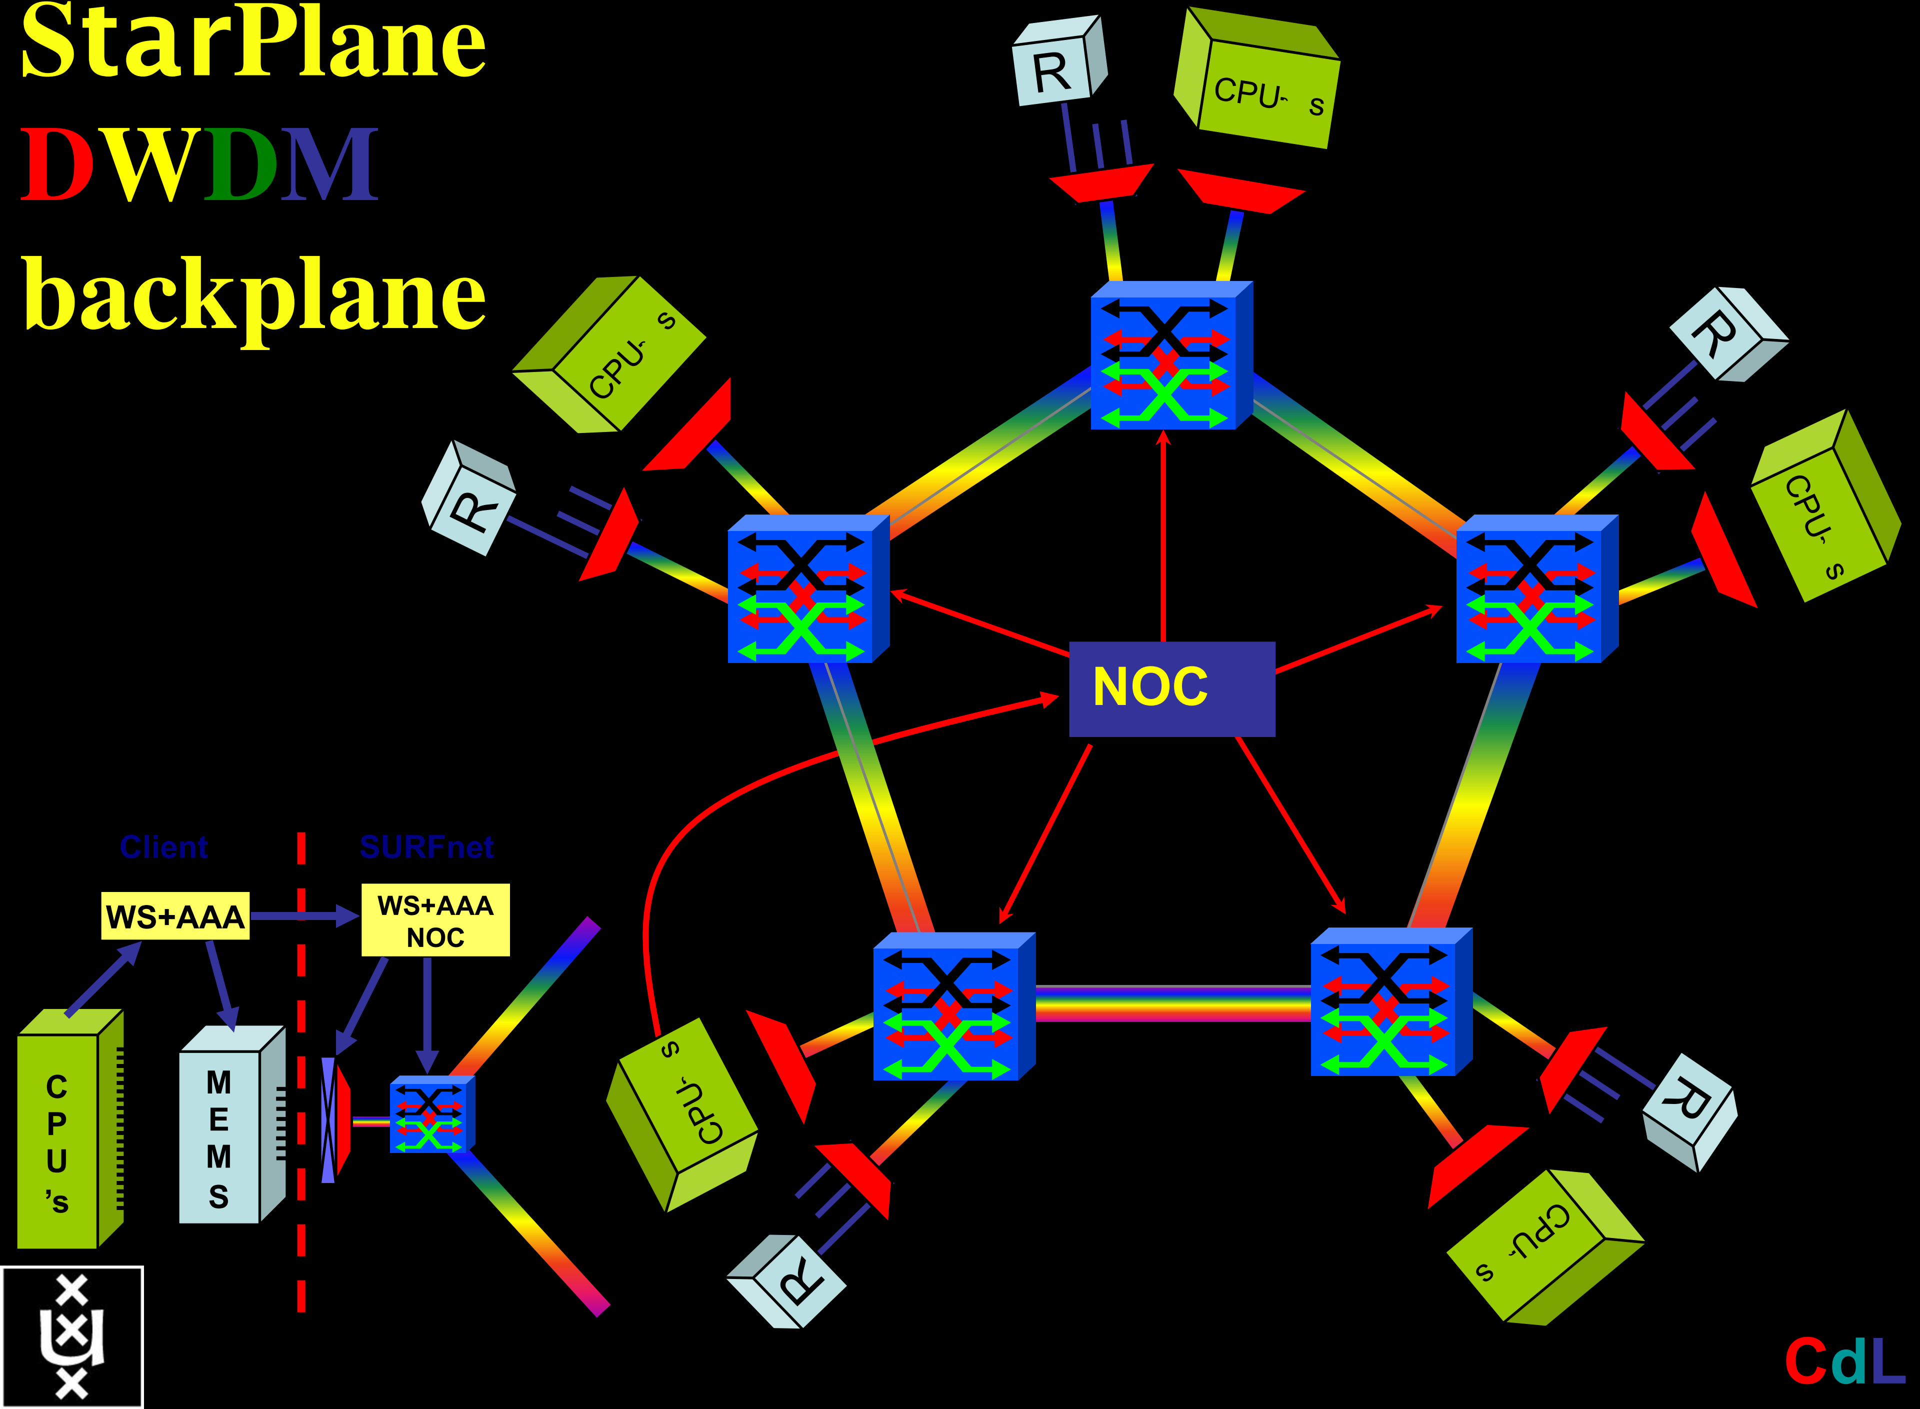

GRID-Colocation problem space

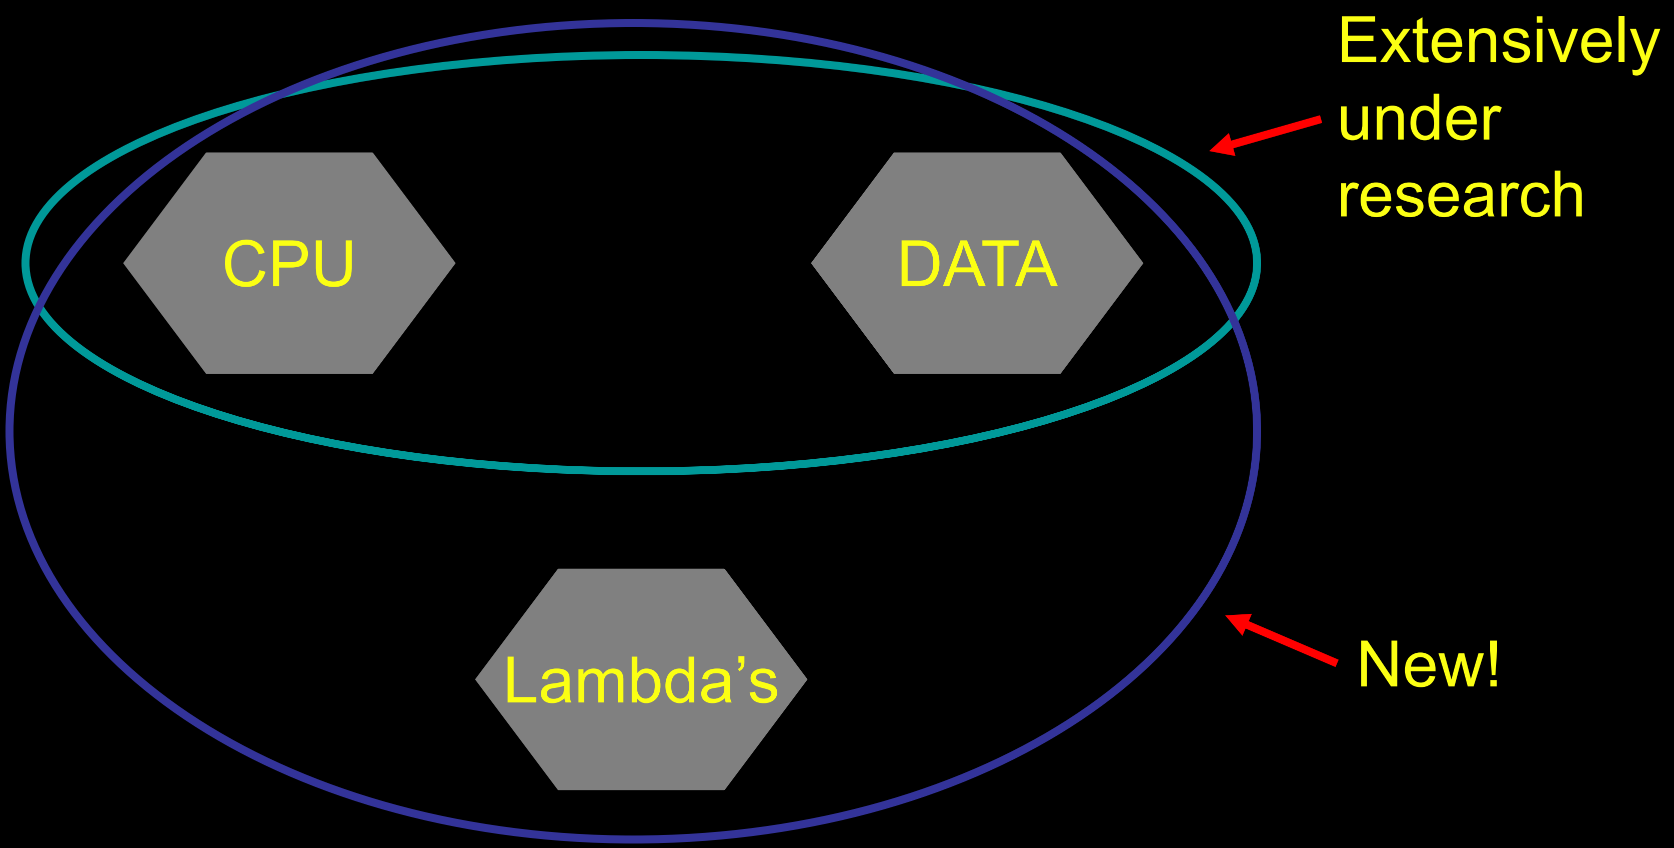

Day 2 set-up: branching out…

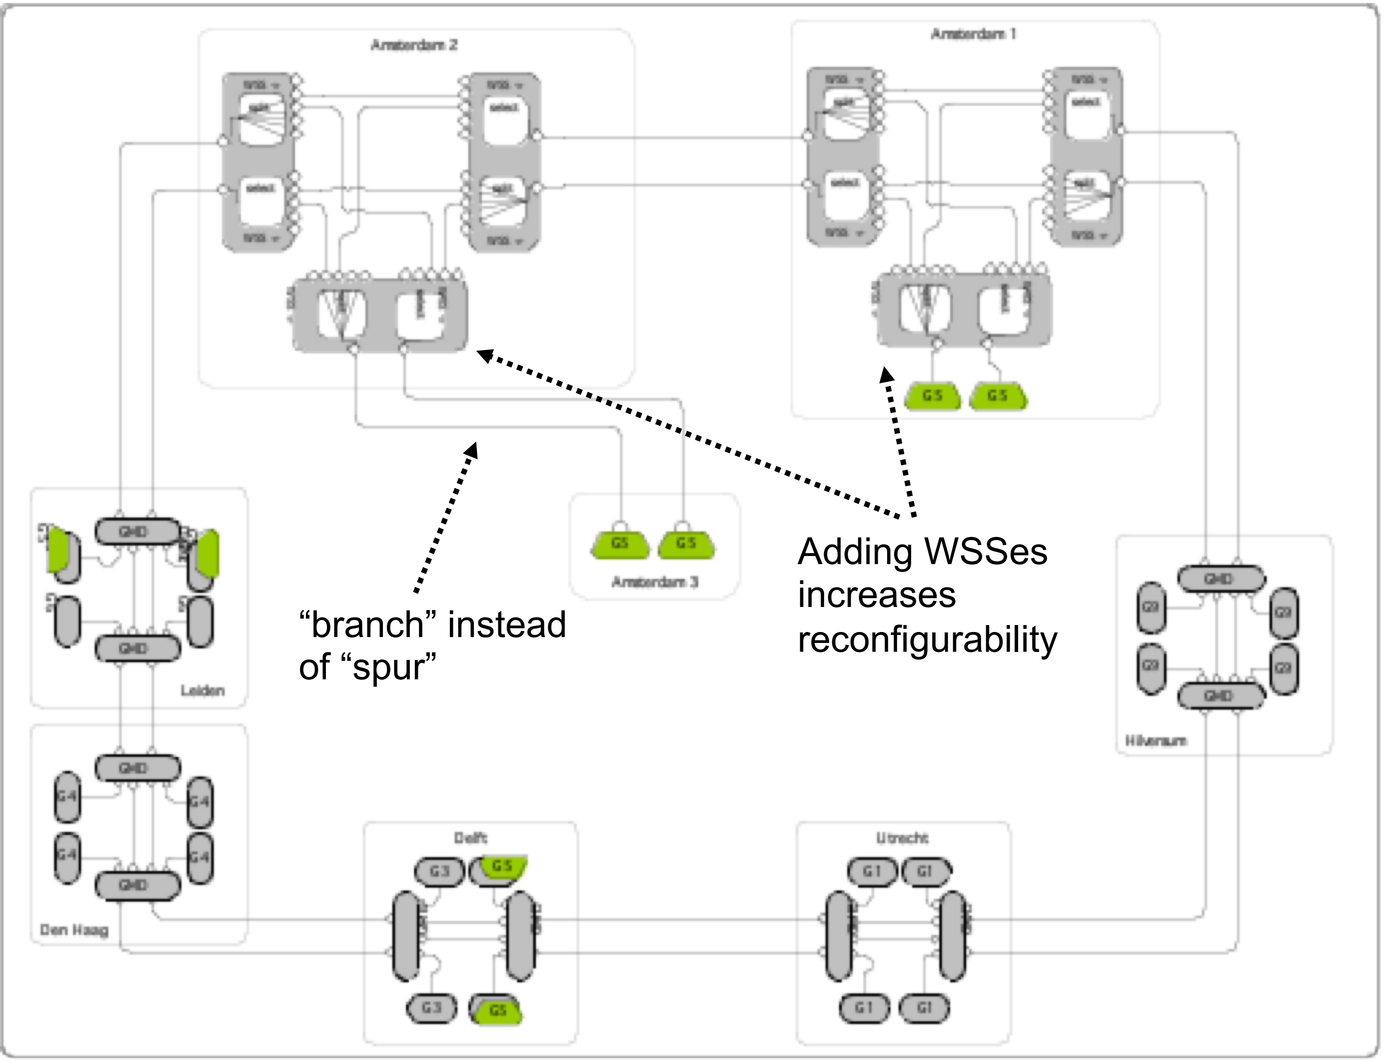

- Add WSSes at Amsterdam sites
- Is NOT supported in March 2006
- Full reconfigurability achieved
- Only limits are
  - Presence of card
  - Wavelength blocking
- No changes to basic ‘static’ mesh

Day 2 detail

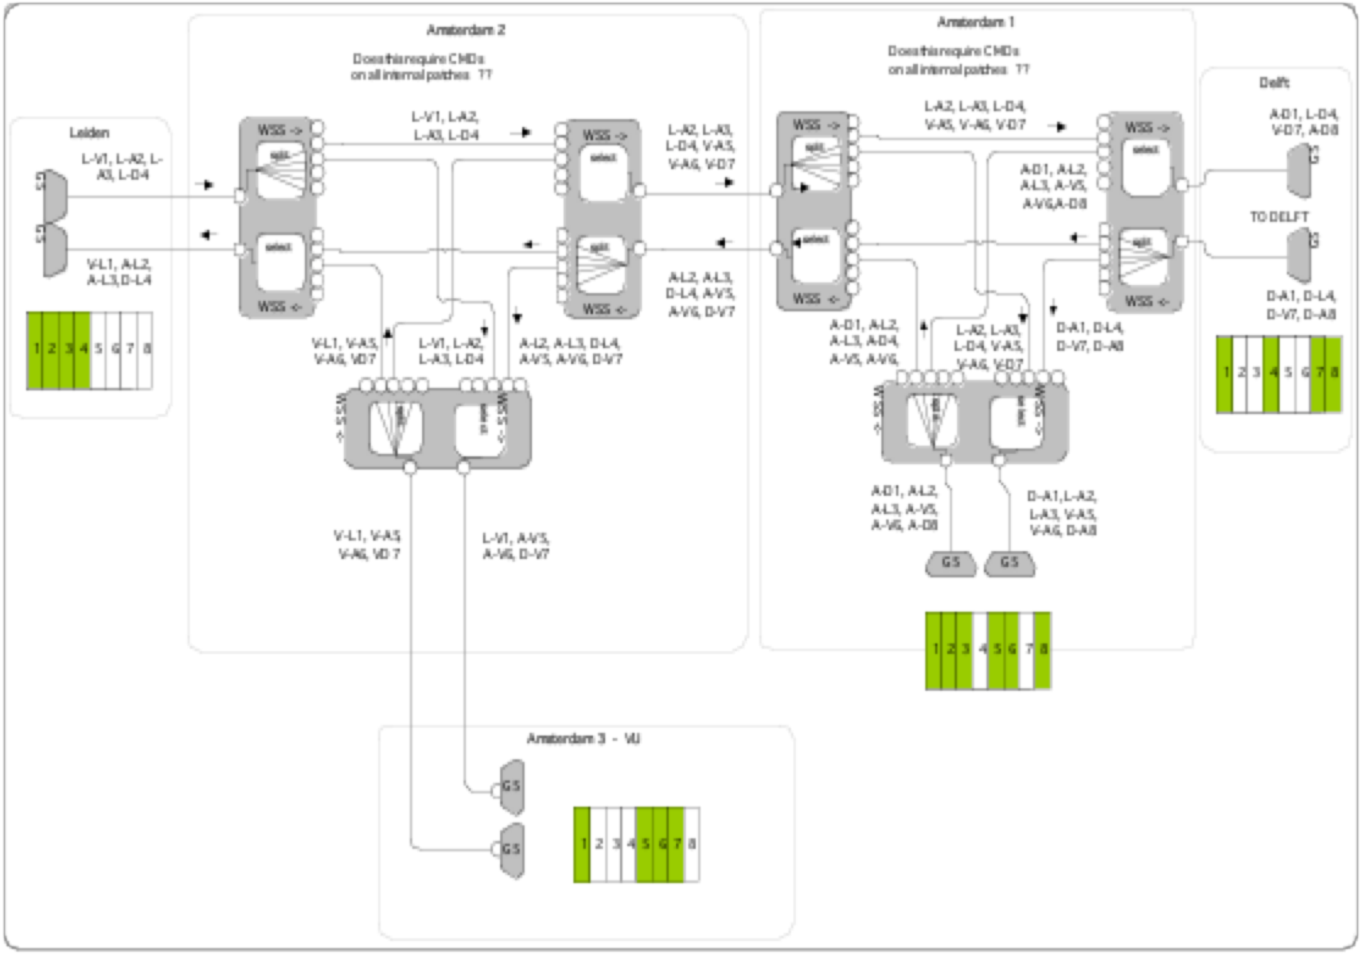

- Wavelength assignment remains – no external changes
- Adding WSSes allows redirecting wavelengths from/to VU and AMS

Day 2 – black box reconfigurability

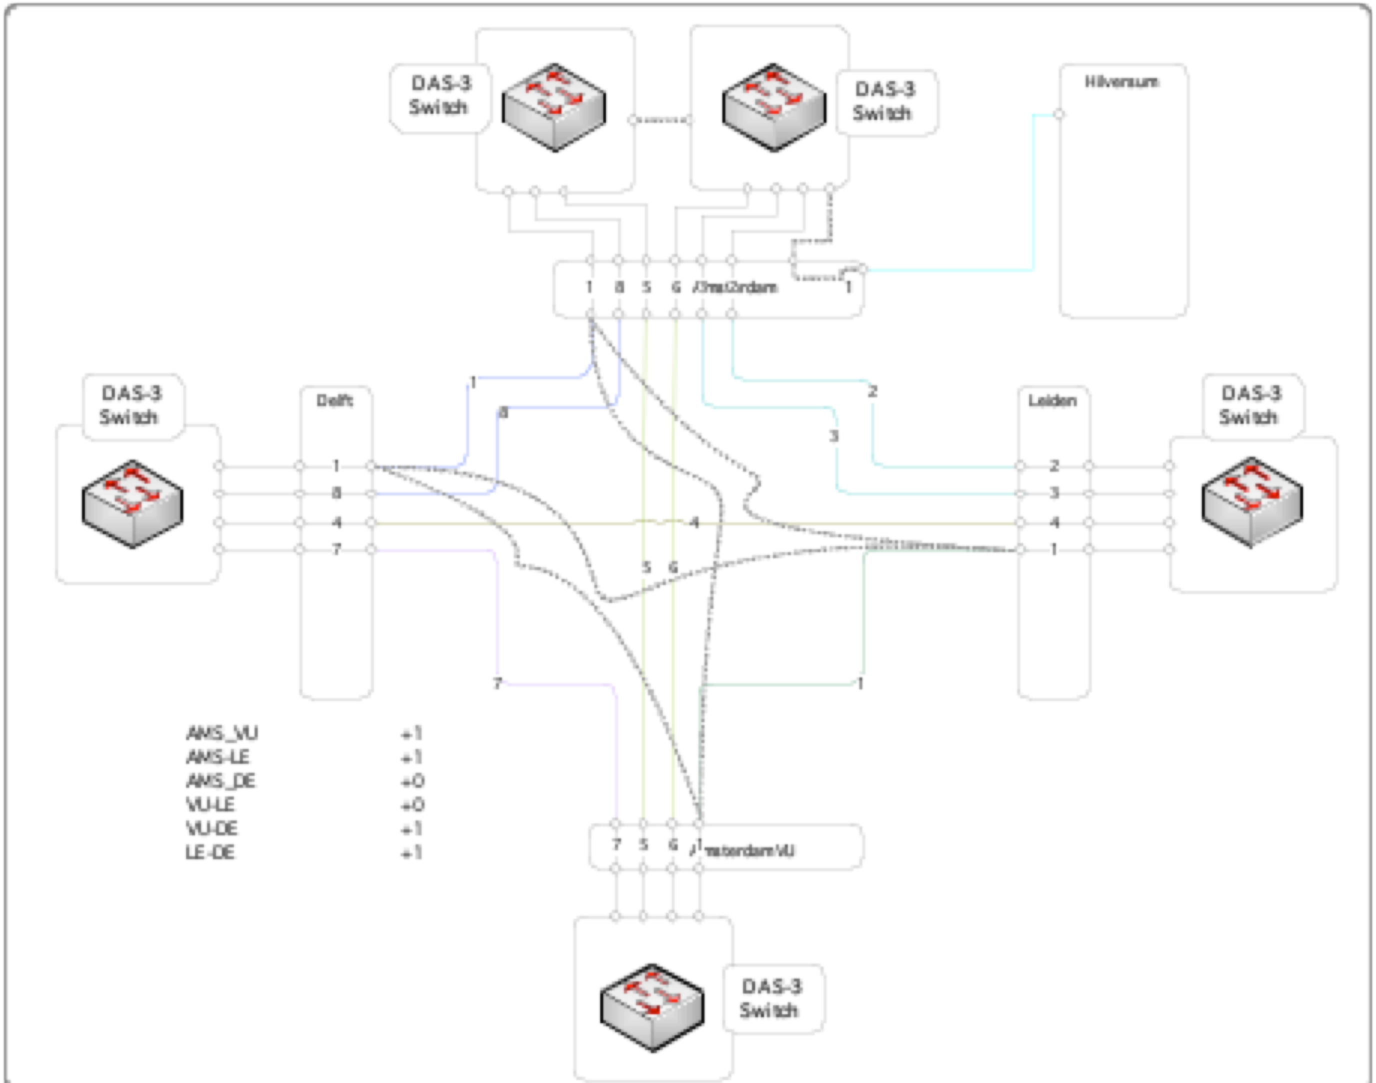

- Compared to day 1 now four instead of one possible redirection
- Redirection only limited by presence of cards and internal wavelength blocking

Day 2 – increased reconfigurability - adding cards

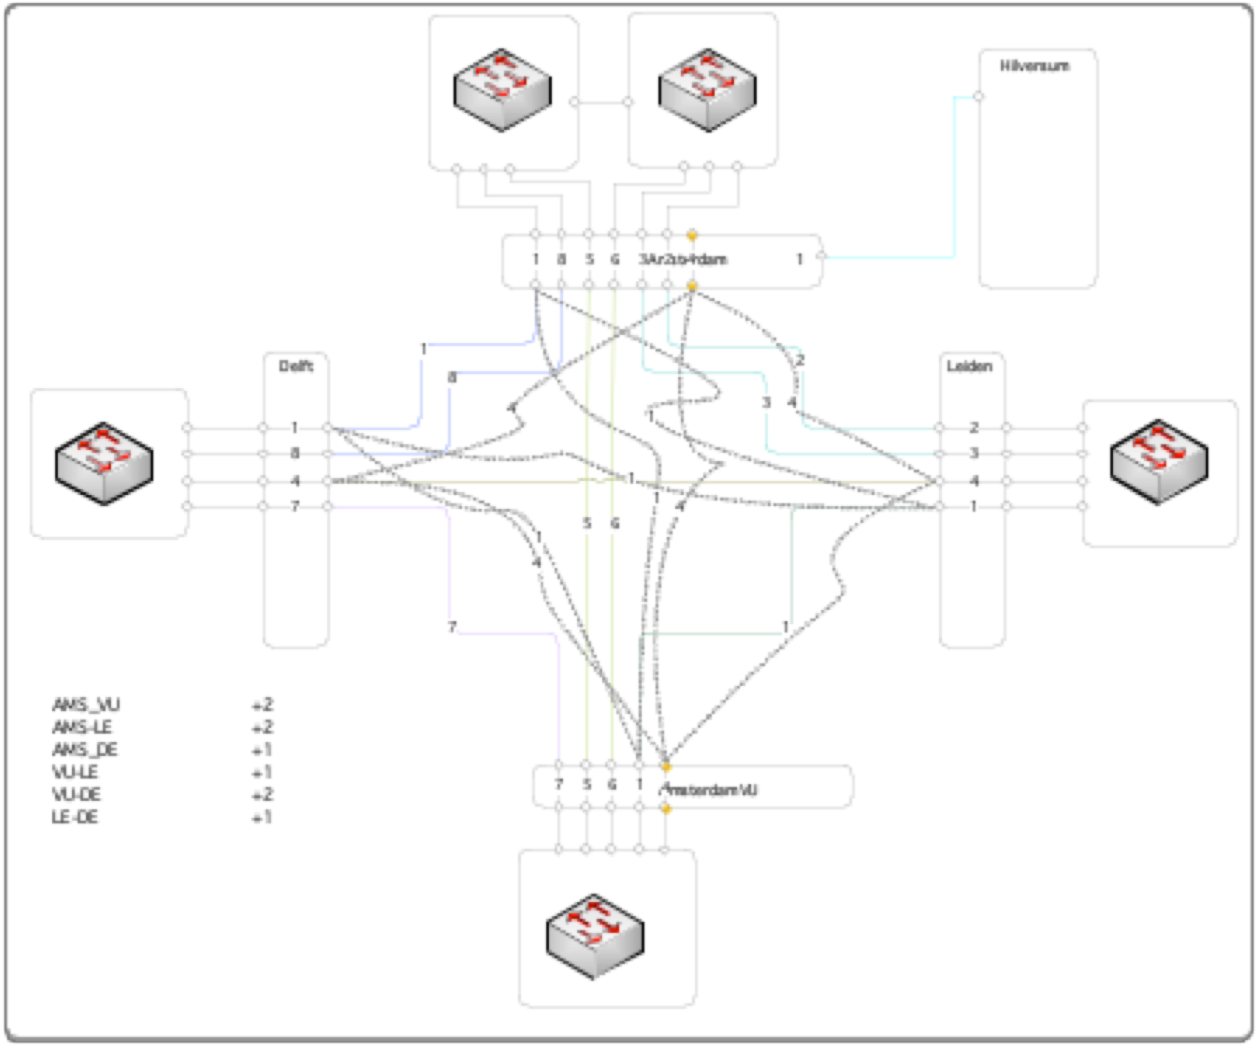

- Adding two cards allows to create more connectivity between ALL sites!
- Some sites can connectivity threefold (from 10 Gb/s to 30 Gb/s)

StarPlane

application-specific management of optical networks

The StarPlane project addresses two concerns in optical networks:

1. The Basic StarPlane Management Infrastructure

StarPlane allows applications to take advantage of the increased bandwidth and potential flexibility in optical networks by letting them create their own network topology in a simple way.

2. The Applications and Their Needs

StarPlane will discover how this new freedom to manipulate the network will benefit the applications.

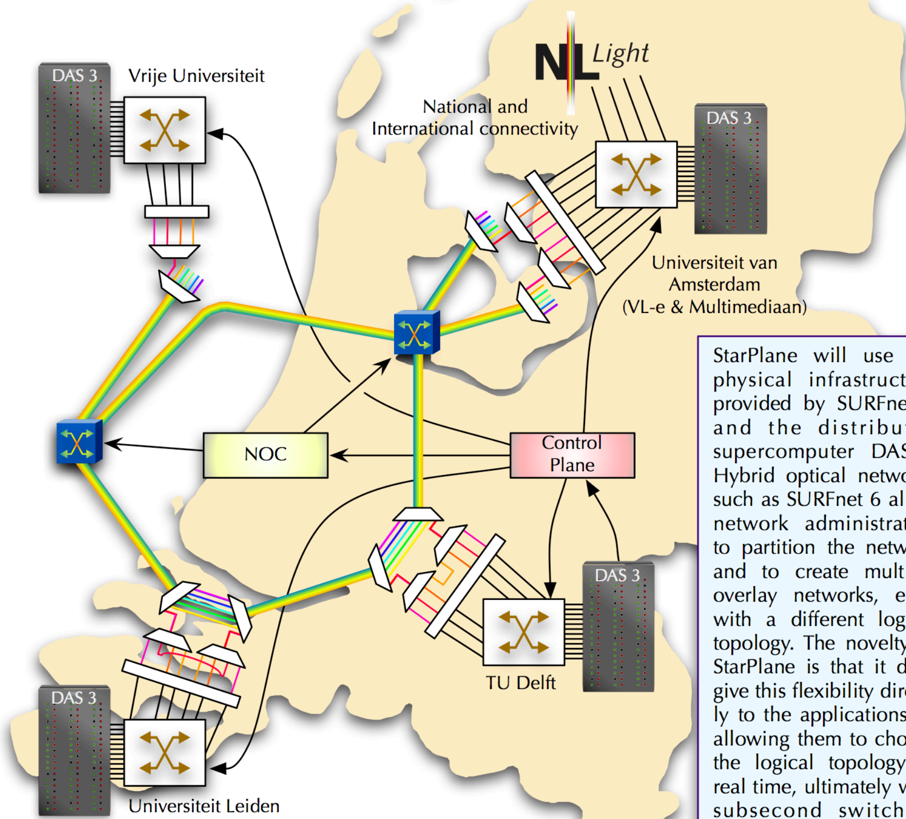

StarPlane will use the physical infrastructure provided by SURFnet 6 and the distributed supercomputer DAS-3. Hybrid optical networks such as SURFnet 6 allow network administrators to partition the network and to create multiple overlay networks, each with a different logical topology. The novelty of StarPlane is that it does give this flexibility directly to the applications by allowing them to choose the logical topology in real time, ultimately with subsecond switching times.

Staff members of the research team:

Prof. dr. ir. H.E. Bal	VU	professor	Parallel programming
Dr. ir. H. Bos	VU	assistant professor	Computer networks
Dr. ir. C.T.A.M. de Laat	UvA	associate professor	Internet and Grids
Prof. dr. P.M.A. Sloot	UvA	professor	Computational science

DWDM network provided by SURFnet

SURFnet

Funded by NWO in the GLANCE program

NWO Netherlands Organisation for Scientific Research

vrije Universiteit amsterdam

UvA Universiteit van Amsterdam

www.starplane.org

GLIF Mission Statement

Global Lambda Integrated Facility

- **GLIF is a world-scale Lambda-based Laboratory for application and middleware development on emerging LambdaGrids, where applications rely on dynamically configured networks based on optical wavelengths**

- **GLIF is an environment (networking infrastructure, network engineering, system integration, middleware, applications) to accomplish real work**

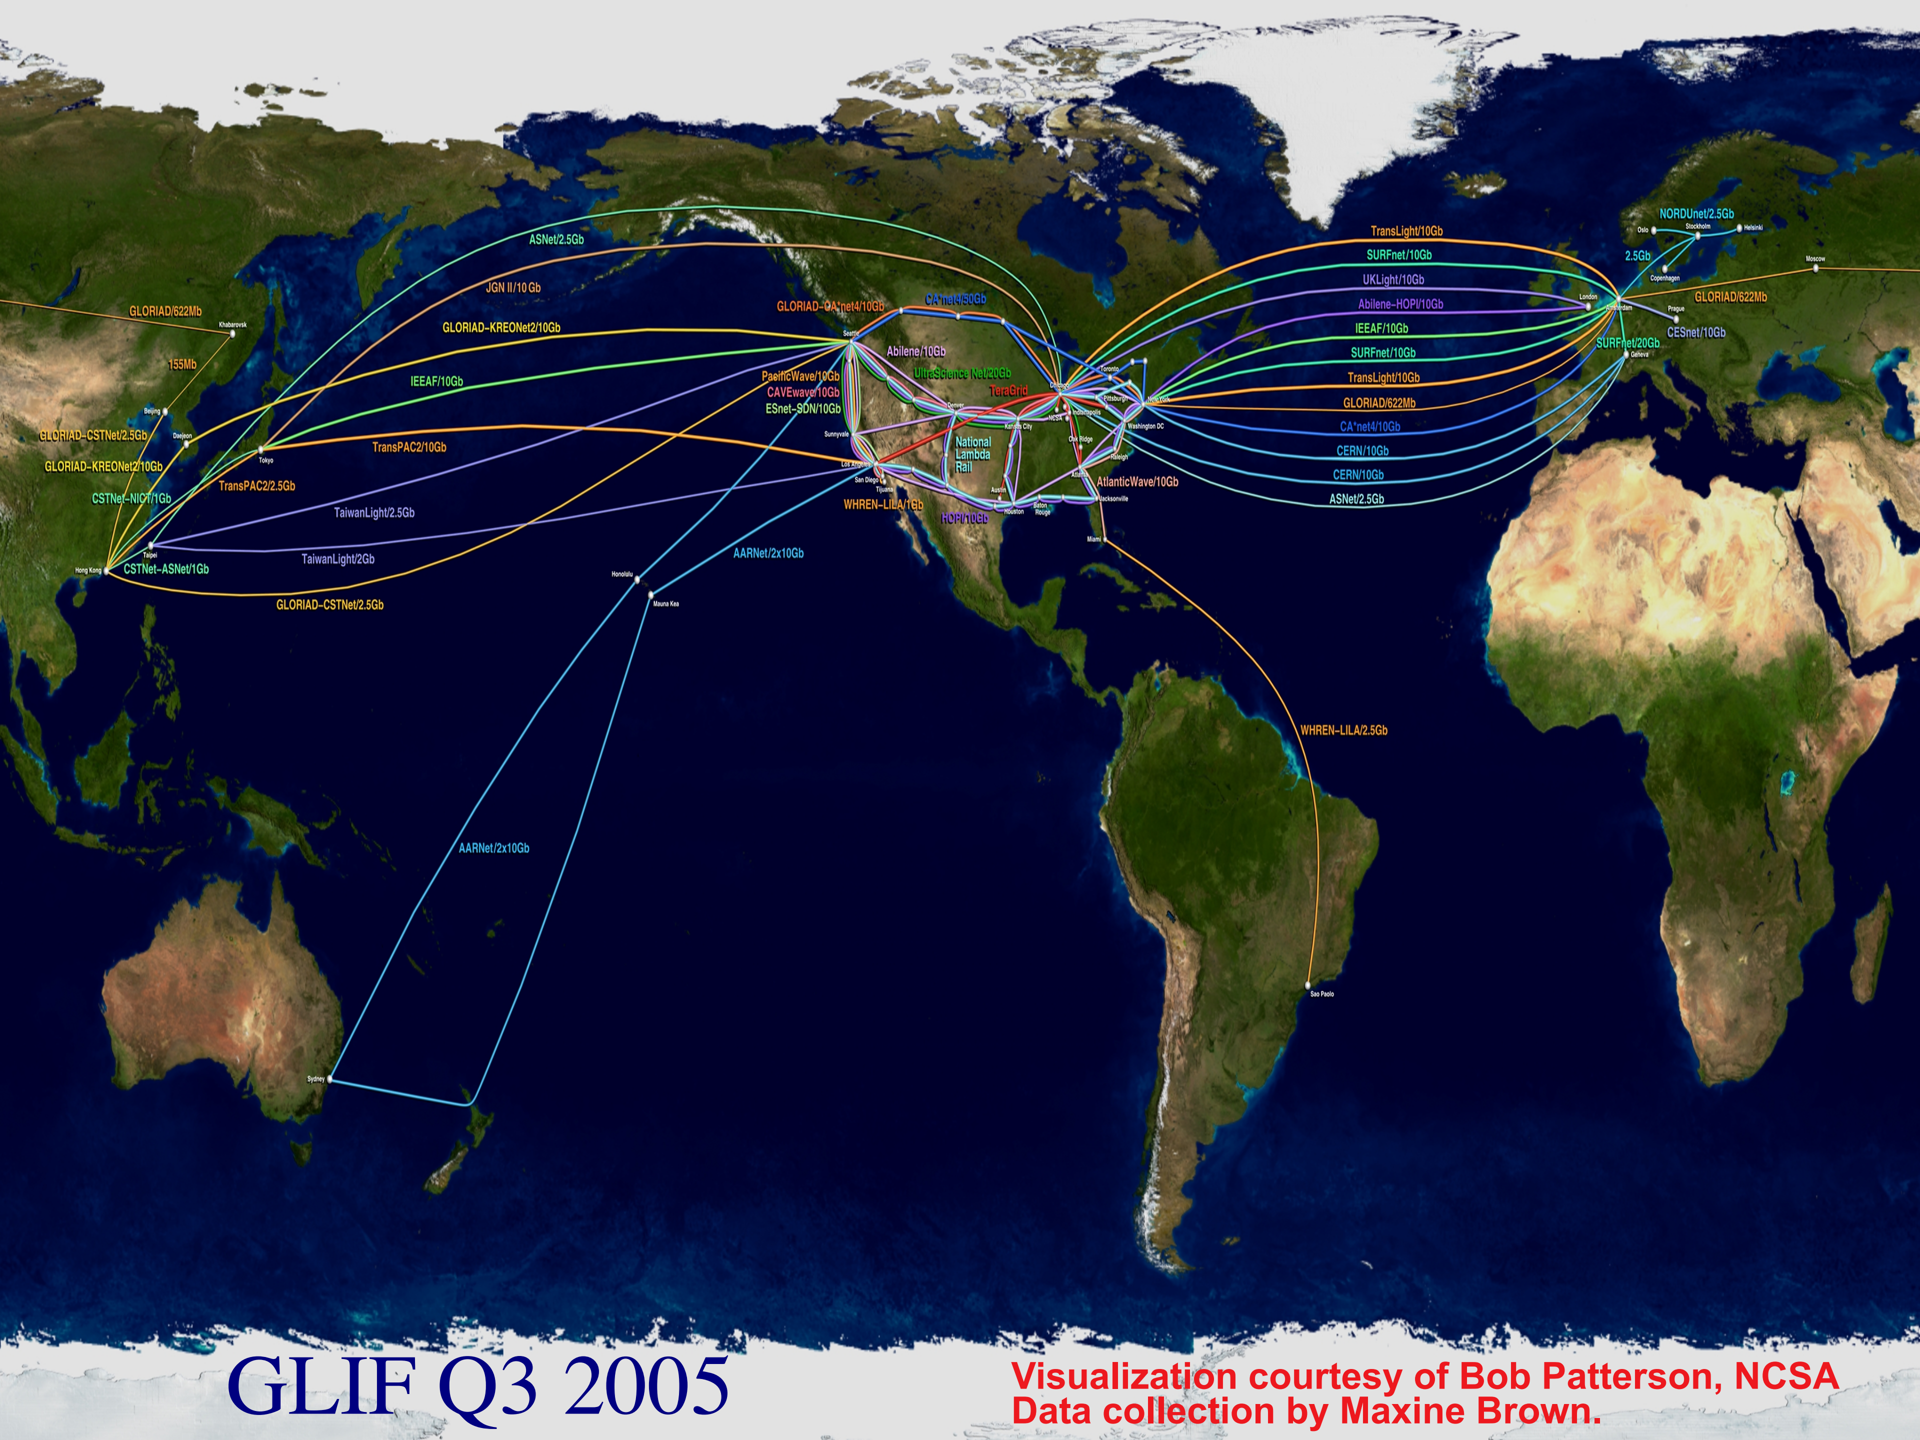

GLIF Q3 2005
Visualization courtesy of Bob Patterson, NCSA
Data collection by Maxine Brown.
ASNet/2.5Gb
JGN II/10 Gb
GLORIAD/622Mb
GLORIAD-KREONet2/10Gb
155Mb
IEEAF/10Gb
GLORIAD-CSTNet/2.5Gb
GLORIAD-KREONet2/10Gb
CSTNet-NICT/1Gb
TransPAC2/10Gb
TransPAC2/2.5Gb
TaiwanLight/2.5Gb
TaiwanLight/2Gb
CSTNet-ASNet/1Gb
GLORIAD-CSTNet/2.5Gb
GLORIAD-CA*net4/10Gb
CA*net4/50Gb
Abilene/10Gb
UltraScience Net/20Gb
TeraGrid
PacificWave/10Gb
CAVEwave/10Gb
ESnet-SDN/10Gb
National
Lambda
Rail
WHREN-LILA/1Gb
HOPI/10Gb
AtlanticWave/10Gb
AARNet/2x10Gb
AARNet/2x10Gb
WHREN-LILA/2.5Gb
TransLight/10Gb
SURFnet/10Gb
UKLight/10Gb
Abilene-HOPI/10Gb
IEEAF/10Gb
SURFnet/10Gb
TransLight/10Gb
GLORIAD/622Mb
CA*net4/10Gb
CERN/10Gb
CERN/10Gb
ASNet/2.5Gb
NORDUnet/2.5Gb
2.5Gb
GLORIAD/622Mb
CESnet/10Gb
SURFnet/20Gb

GLIF Structure

aka de tulpenbollen-inpakkers

GLIF Governance and policy

Our small-scale Lambda Workshop is now turning into a global activity. TransLight and similar projects contribute to the infrastructure part of GLIF. A good and well understood governance structure is key to the manageability and success of GLIF. Our prime goal is to decide upon and agree to the GLIF governance and infrastructure usage policy.

GLIF Lambda infrastructure and Lambda exchange implementations

A major function for previous Lambda Workshops was to get the network engineers together to discuss and agree on the topology, connectivity and interfaces of the Lambda facility. Technology developments need to be folded into the architecture and the expected outcome of this meeting is an agreed view on the interfaces and services of Lambda exchanges and a connectivity map of Lambdas for the next year, with a focus on iGrid 2005 and the emerging applications.

Persistent Applications

Key to the success of the GLIF effort is to connect the major applications to the Facility. We, therefore, need a list of prime applications to focus on and a roadmap to work with those applications to get them up to speed. The demonstrations at SC2004 and iGrid 2005 can be determined in this meeting.

Control Plane and Grid Integration

The GLIF can only function if we agree on the interfaces and protocols that talk to each other in the control plane on the contributed Lambda resources. The main players in this field are already meeting, almost on a bi-monthly schedule. Although not essential, this GLIF meeting could also host a breakout session on control plane middleware.

Resource Brokering:
Your Ticket Into NetherLight

Application architecture:

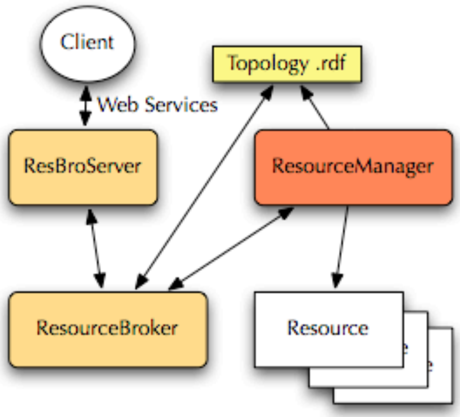

Lambda networking allows the creation of application specific light paths.

Lambda networking facilities empower users to request services and provision **end-to-end light paths** if and when they need it.

NetherLight, located in Amsterdam, The Netherlands, is one of such facilities.

The Amsterdam **LightHouse** is a joint research laboratory of the UvA and SARA.

Resources in the LightHouse can be used by collaborators to prove the concepts of hybrid networks.

Lightpath setup components:

Topology information → We make use of **semantic web** techniques.
The description of the network is contained in RDF files.

Reservation system → We provide **web services** interfaces to the client for:
- resources and path inquiries;
- reservation handling.

Management system → We provision the paths on the LightHouse equipment.

Semantic web
The Network Description Language, an RDF Schema, describes networks in a standard, interoperable way.

Web Services
A WSDL file describes the interfaces to the service available to clients. Clients can interact with the service directly or via a portal.

Our SC|05 demonstration
We show the setup of dynamic connections between two computing nodes through the LightHouse/ NetherLight Optical Exchange.

UvA Universiteit van Amsterdam

Paola Grosso
Jeroen van der Ham

GigaPort

Transport of flows

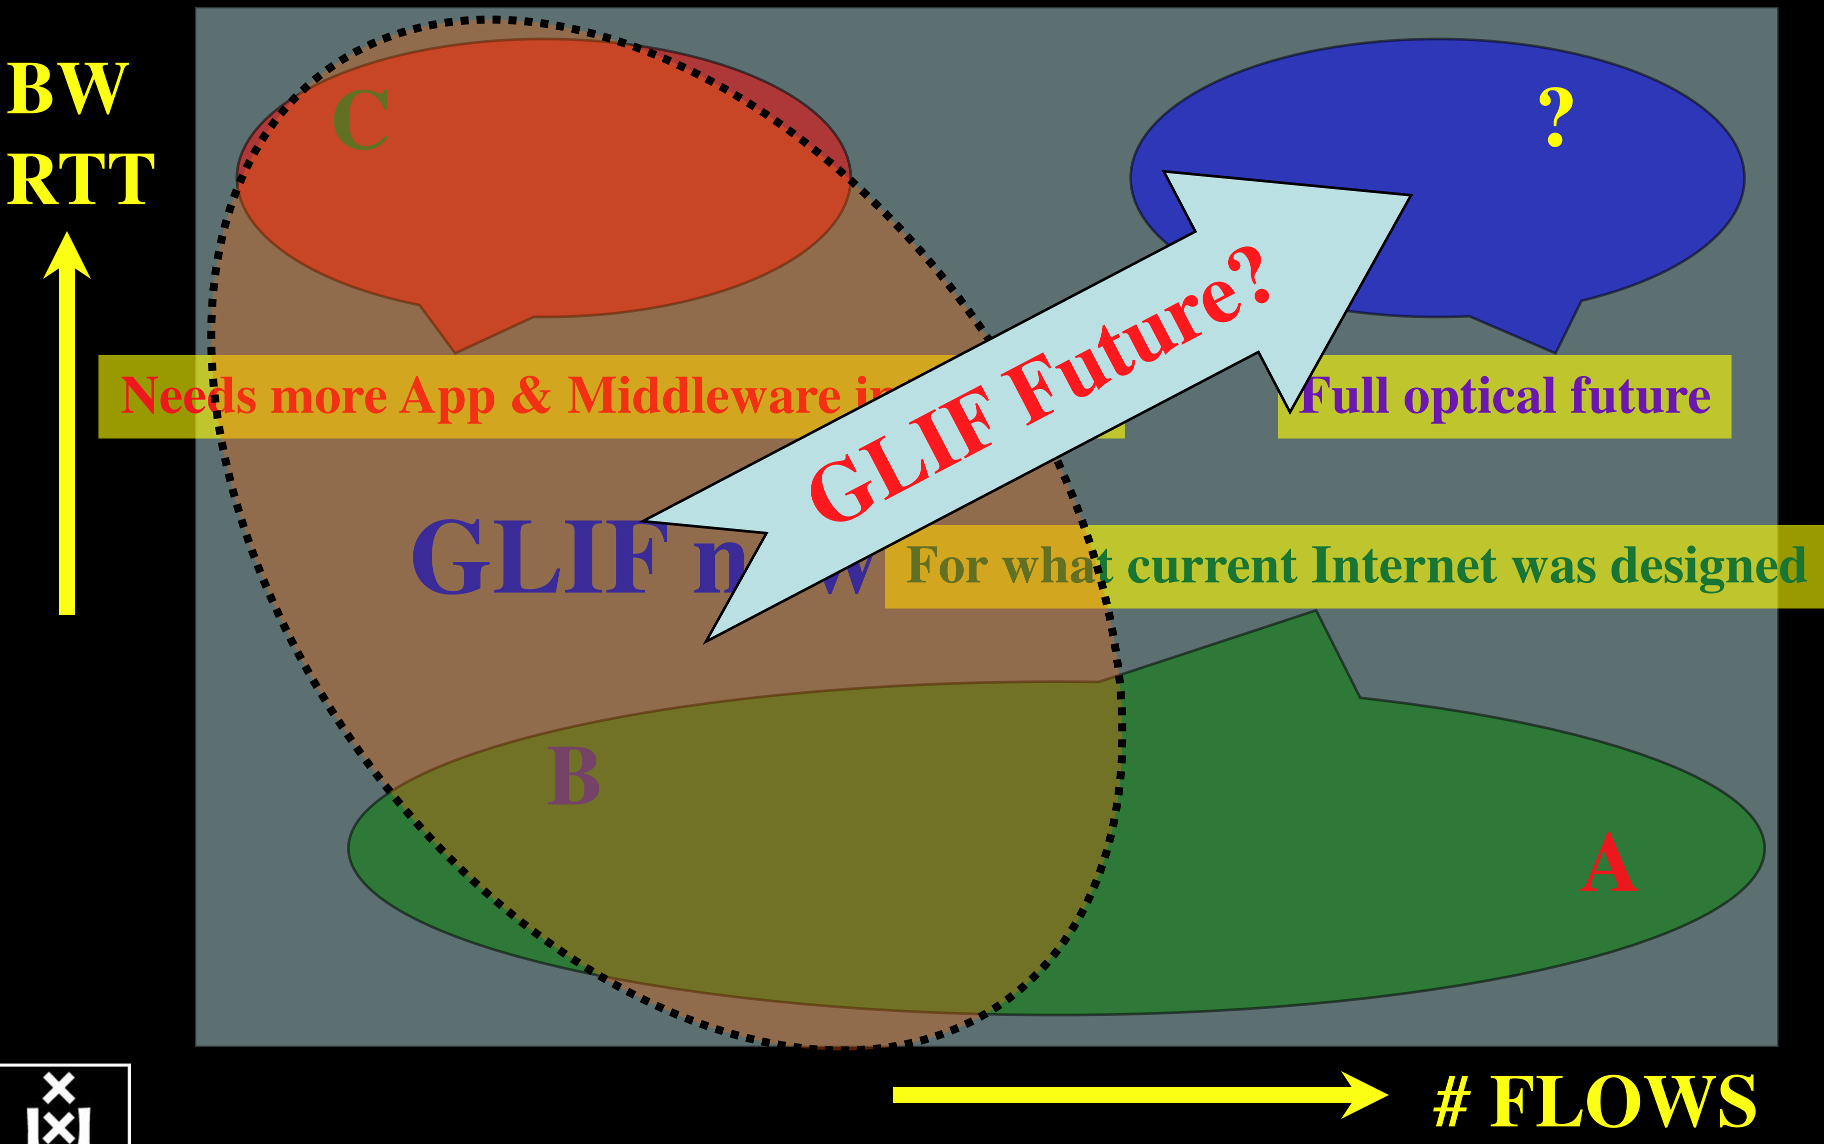

ASK YURI DEMCHENKO

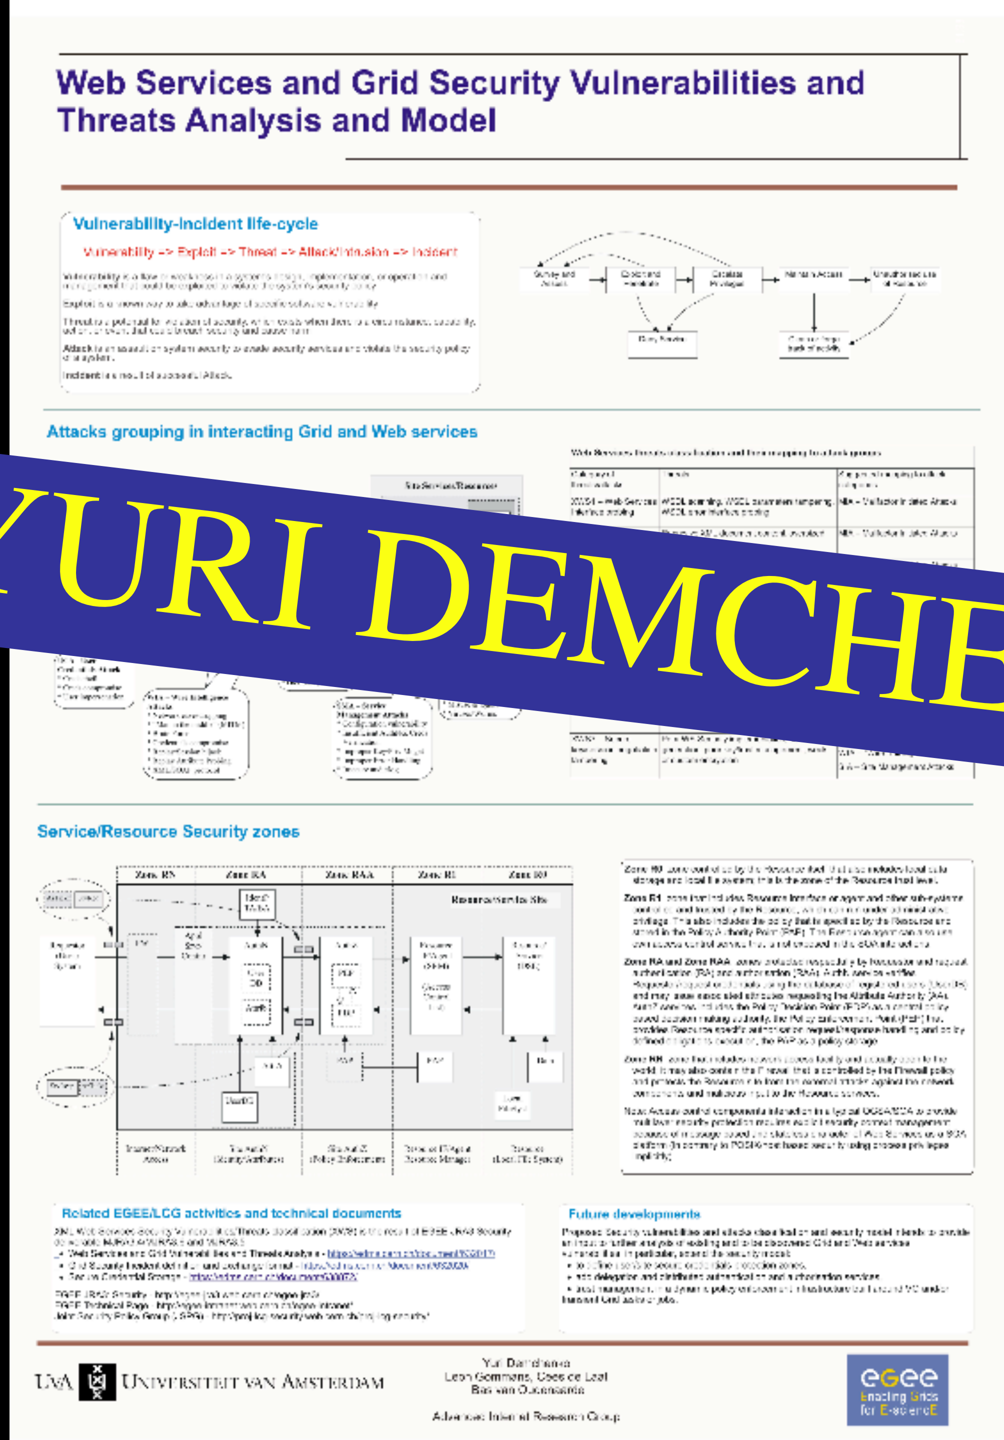

DEMO US117b – GLVF TOPS

- Very large 2D and 3D datasets located at SARA in Amsterdam
- Render cluster (29 nodes) at SARA
- 20 Gbps connectivity between cluster and NetherLight, via University of Amsterdam switch
- 2 * 10 Gbps lambda between Amsterdam and San Diego (SURFnet, CaveWave, WANPHY Amsterdam-SD)
- UCSD LambdaVision display at iGrid: 100 Megapixel Tiled Display
- 1 Frame = 100 Mpixel*24 bits = 2.4 Gbits
- Uncompressed and compressed mode
- Streaming pixels from Amsterdam to San Diego

OptIPuter Partner

OptIPuter

Final results

- Third demoslot, 2 times 10 Gbps available
  - Sustained bandwidth used between rendering and display 18 Gbps
  - Peak bandwidth of 19.5 Gbps!
- New world record of transatlantic bandwidth usage by one single application visualizing scientific content

Ask Paul Wielinga

OptIPuter Partner

The "Dead Cat" demo

Produced by:
- Michael Scarpa
- Robert Belleman
- Peter Sloot
- Cees de Laat

Many thanks to:
- AMC
- SARA
- GigaPort
- UvA/AIR
- Silicon Graphics
- Zoölogisch Museum

I am not that bad!

Photo 1977

Questions ?

More info:

http://www.glif.is/

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

delaat@uva.nl