TransLight

A Global LambdaGrid for e-Science

Maxine Brown
Electronic Visualization Laboratory
University of Illinois at Chicago

The Keukenhof, Leiden, The Netherlands

STARLIGHT℠

STAR TAP International Connectivity

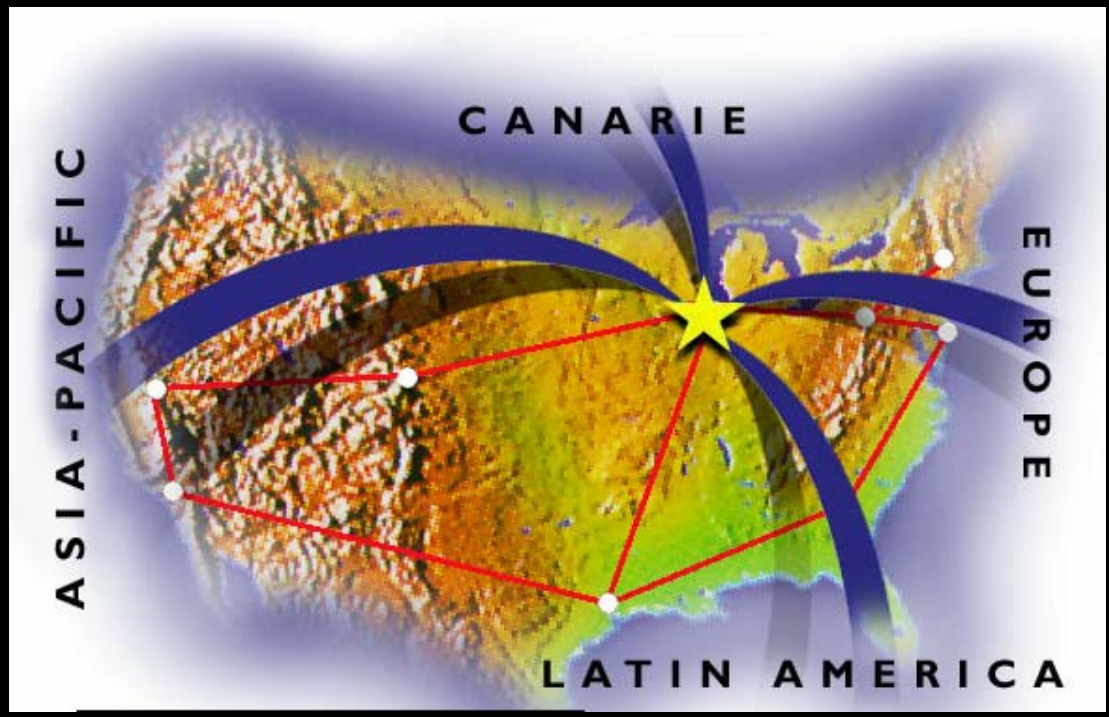

STAR TAP was created in 1997 to provide a persistent infrastructure for the long-term interconnection and interoperability of advanced international networking in support of applications, performance measuring, and technology evaluations. By 2000, STAR TAP successfully became a model for Next-Generation Internet eXchanges (NGIXs).

StarLight: Global Optical Internet Exchange

StarLight is an experimental optical infrastructure and proving ground for network services optimized for high-performance applications.

Operational since summer 2001, **StarLight is perhaps the world's largest 1GigE and 10GigE switch/router exchange** for high-performance access to participating networks and is becoming a true optical switching facility for wavelengths.

Abbott Hall, Northwestern University's Chicago downtown campus

US NRN International Peering Locations

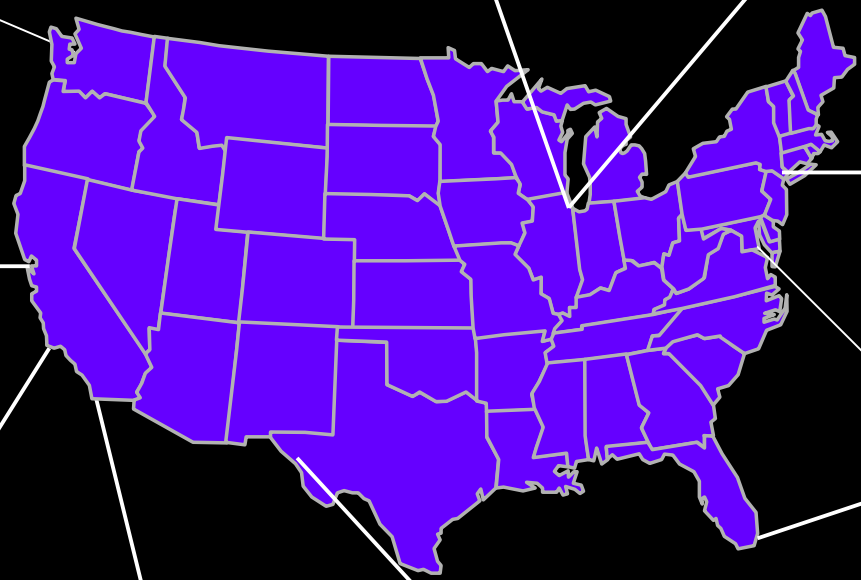

STAR TAP

Abilene
DREN
NISN
vBNS+
6TAP IPv6

CERNET
GEMnet
HARNET
HEAnet
KOREN/KREONET2
Renater2
TANet2

STARLIGHT

Abilene
ESnet
NREN
6TAP IPv6

AMPATH
APAN
ASnet
CA*net4
CERN
EU DataTAG
NORDUnet

RBnet
SURFnet

Other:
I-WIRE
MREN
OMNInet
TeraGrid DTFnet

PACIFIC WAVE/ SEATTLE

Abilene
AARnet
APAN
CA*net4
GEMnet

SingAREN

DREN
ESnet
TAnet2

SUNNYVALE

Abilene
ESnet
6TAP IPv6

WIDE

NEW YORK CITY

Abilene
ESnet
CA*net4
GÉANT

HEAnet
IEEAF/Tyco/
SURFnet
NORDUnet
SINET

WASHINGTON DC

Abilene
GÉANT

AMPATH/ MIAMI

Abilene
ANSP (Brazil)
REACCIUN-2 (Venezuela)
RETINA (Argentina)
REUNA2 (Chile)
RNP (Brazil)

LOS ANGELES

Abilene
UNINET

SAN DIEGO

Abilene
CUDI

EL PASO

Abilene
CUDI

STARLIGHT

NSF Network Types

- NSF describes three classes of Research & Education networks beyond the commodity Internet:
  - Production Networks
  - Research Networks
  - Experimental Networks

Production Networks

- High-performance networks, which are always available and dependable.
- These networks reach all researchers.

Research Networks

- Smaller-scale networks for basic scientific and engineering network research.
- Research Networks are not expected to be reliable, nor are they expected to support applications.

Experimental Networks

- High-performance trials of new technologies that support ***application-dictated*** development of software toolkits, middleware, computing and networking.
- Provide ***known and knowable characteristics*** with deterministic and repeatable behavior on a persistent basis, while encouraging experimentation with innovative concepts.
- Experimental Networks are seen as the ***missing link*** between Research and Production Networks.

http://www.evl.uic.edu/activity/NSF/index.html
http://www.calit2.net/events/2002/nsf/index.html

The Next S-Curves of Exponential Technology Growth

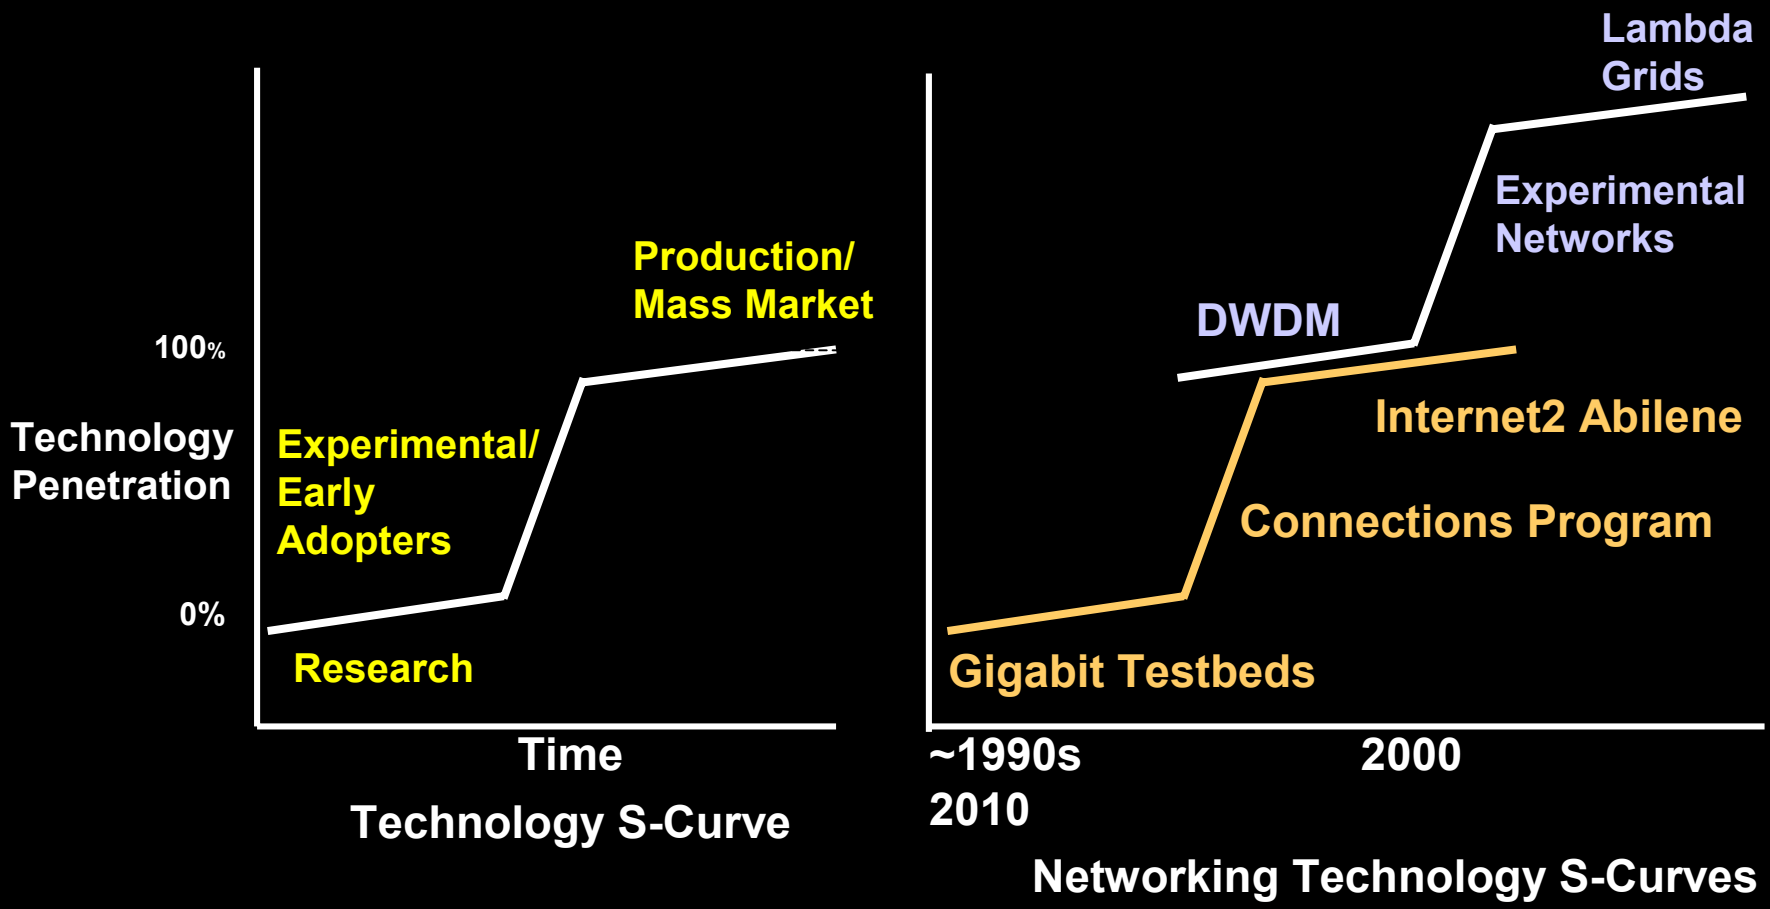

StarLight is Home for all Types of Networks

- A **Production Network** 1GigE and 10GigE exchange
- An **Experimental Network** optical (lambda) exchange and *proving ground for network services* optimized for high-performance computing applications
- A **Research Network** 1GigE and 10GigE MEMS-switched exchange
- A co-location space with 66 racks for networking *and* computing and data-management equipment
- An OIX with fiber and/or circuits from SBC/Ameritech, Qwest, AT&T, Global Crossing, Looking Glass Networks, Level 3, RCN, T-Systems, *I-WIRE*
- A **lambda-switching** facility, with links coming from USA, NetherLight, DataTAG, CA*net4, and proposed from UK-Light and APAN forming ***TransLight***

Who Needs Experimental Networks? iGrid 2002, September 2002, Amsterdam

- **28 demonstrations from 16 countries:** Australia, Canada, CERN, France, Finland, Germany, Greece, Italy, Japan, The Netherlands, Singapore, Spain, Sweden, Taiwan, United Kingdom, United States
- Applications demonstrated: art, bioinformatics, chemistry, cosmology, cultural heritage, education, high-definition media streaming, manufacturing, medicine, neuroscience, physics, tele-science



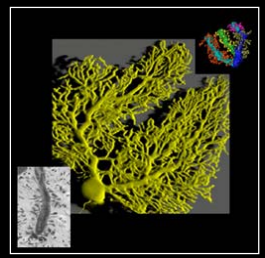


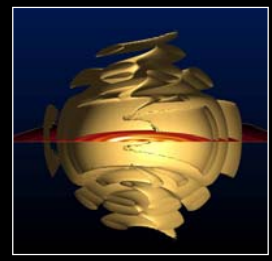

- Grid technologies demonstrated: Major emphasis on grid middleware, data management grids, data replication grids, visualization grids, data/visualization grids, computational grids, access grids, grid portals
- **25Gb transatlantic bandwidth (100Mb/attendee, 250x iGrid2000!)**

www.startap.net/igrid2002

STARLIGHT℠

Who Needs Experimental Networks?
iGrid 2002, September 2002, Amsterdam

A worldwide community of e-scientists, application programmers, networking engineers, electrical/computer engineers, artists and computer scientists working together









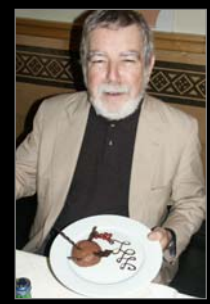



STARLIGHT SM

Who Needs Experimental Networks?

NSF Major Research Equipment (MRE) and Information Technology Research (ITR) projects

ALMA: Atacama Large Millimeter Array
www.alma.nrao.edu

GriPhyN
Data Intensive Science

GriPhyN: Grid Physics Network
www.griphyn.org

TeraGrid
www.teragrid.org

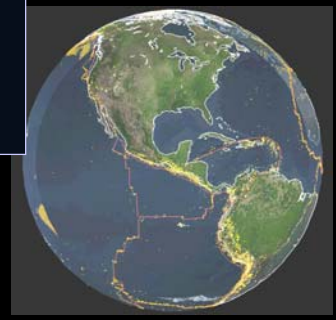

Network for Earthquake Engineering Simulation
www.neesgrid.org

iVDgL

International Virtual Data Grid Laboratory
www.ivdgl.org

OptIPuter

The OptIPuter
www.calit2.net/news/2002/9-25-optiputer.html

PPDG

Particle Physics Data Grid
www.ppdg.net

GEON
CYBERINFRASTRUCTURE FOR THE GEOSCIENCES

GEON: Geosciences Network
www.geongrid.org

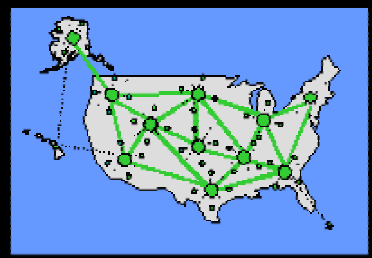

NEON: National Ecological Observatory Network
www.sdsc.edu/NEON

EarthScope
www.earthscope.org

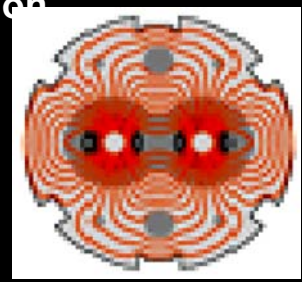

Large Hadron Collider (LHC)
http://lhc-new-omepage.web.cern.ch

STARLIGHT℠

What is a Lambda?

- A *lambda*, in networking, is a fully dedicated wavelength of light in an optical network, typically used today for 1-10Gbps.
- Lambdas are circuit-based technology, but can carry packet-based information.
- We are now mostly working with 1Gb dedicated layer2 circuits that act like lambdas
- For this discussion, *lambda* means large and desirable units of networking, which is how the applications see them, conceptually offering the promise of end-to-end custom connectivity or allowing sufficiently massive over-provisioning of bandwidth so that the connection is effectively uncongested.

What is a LambdaGrid?

- A ***grid*** **is a set of networked, middleware-enabled computing resources.**
- A ***LambdaGrid*** **is a grid in which the lambda networks themselves are resources that can be scheduled, like all other computing resources. The ability to schedule and provision lambdas provides *deterministic* end-to-end network performance for real-time or time-critical applications, which cannot be achieved on today's grids.**

TransLight is a Global-Scale Experimental Network (Coming Summer 2003)

- Supports prototypes of the most aggressive e-science applications coming this decade
- Consists of many GigE lambdas among North America, Europe and Asia via StarLight *available for scheduling*
- Enables researchers to experiment with deterministic provisioning of dedicated circuits, and then compare results with standard production networks
- Experiments include moving large amounts of data, real-time collaboration and visualization, and distributed computing

TransLight Activities

- A TransLight Governance Board to create policy for scheduling circuits; initial members are SURFnet, CANARIE, DataTAG/CERN, StarLight/Euro-Link, with others to be invited soon
- Provide global lambdas for scheduling
- Build more hubs with switches, co-location space and fiber access like StarLight and NetherLight

TransLight Fabric

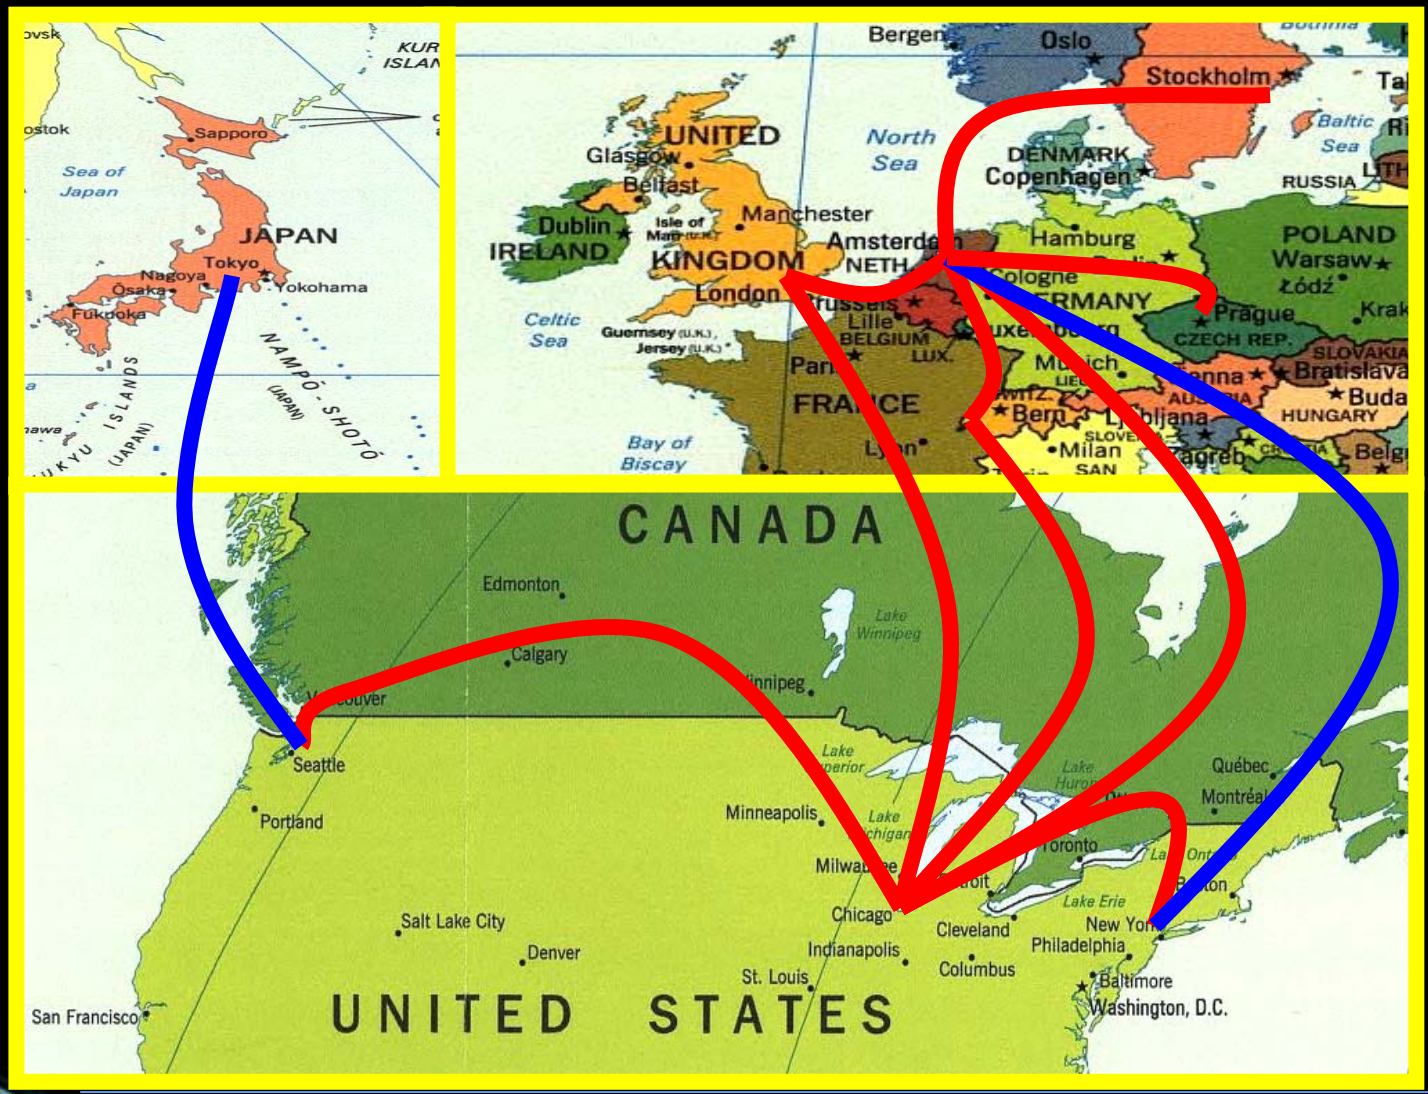

Red Lines:
TransLight Lambda 10Gb Links, Funded

Blue Lines:
TransLight Lambda 10Gb Links, Donated

STARLIGHT℠

The TransLight Project
~54 Int'l GigE Lambdas in 2003/2004

- 16 Canadian Lambdas from StarLight
  - 8 GigEs Chicago to Eastern Canada and NYC
  - 8 GigEs Chicago to Western Canada and Seattle
- 22 European Lambdas to StarLight
  - 10 GigEs Amsterdam to Chicago
  - 4 GigEs CERN to Chicago
  - 8 GigEs London to Chicago
- 12 European Lambdas to NetherLight
  - 4 GigEs CERN to Amsterdam
  - 2 GigEs Prague to Amsterdam
  - 2 GigEs Stockholm to Amsterdam
  - 8 GigEs London to Amsterdam

And many Metro/Regional/ National Lambdas

TransLight Persistent Experiments

- New network protocols – TCP is not usable for long fat applications
- Optical point-to-point connects for instruments
- Circuits are not scalable, but neither are router$
- Application-centric intelligent signaling and management plane middleware – putting intelligence on the edges (not the core)
- Tuning compute, data, visualization, networking using clusters to get orders of magnitude improvement
- Security at 10Gb line speed

Extremely High-Resolution Digital Montage Visualization

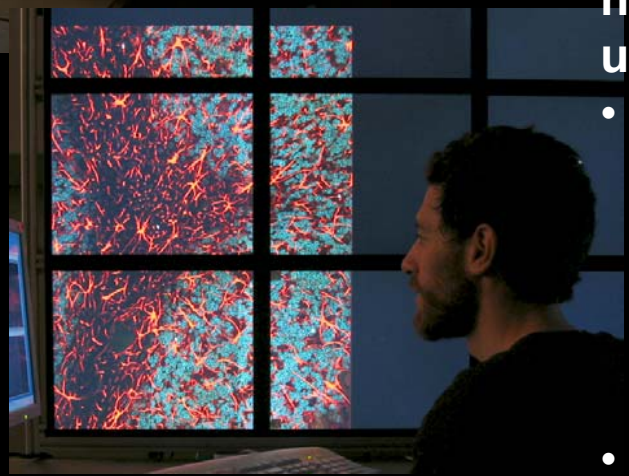

- Large digital montage viewer for tiled LCD displays view high-resolution montages from NCMIR, Scripps, USGS.
- Example: In a Homeland Security / Emergency Response Application, USGS has aerial photos of 133 urban areas:
  - 5643 tiles each 5000x5000 pixel resolution ~ 375,600x375,600 pixels for each urban area (394GB per area)
  - Total data ~ 51 TB

www.evl.uic.edu/cavern

STARLIGHT SM

StarLight as a Research Network Using Optical Switching

- No need to look at every packet when transferring a terabyte of information
  - 1% the cost of routing
  - 10% the cost of switching
  - 64x64 10Gb:
    - $100,000 O-O-O switched
    - $1,000,000 O-E-O switched
    - $10,000,000 O-E-O Routed
- Spend the savings on computing and collaboration systems instead!
- Replaces patch panels; allows rapid reconfiguration of 1 and 10Gb experiments

StarLight as a Research Network
Optical Micro-Mirror Switching

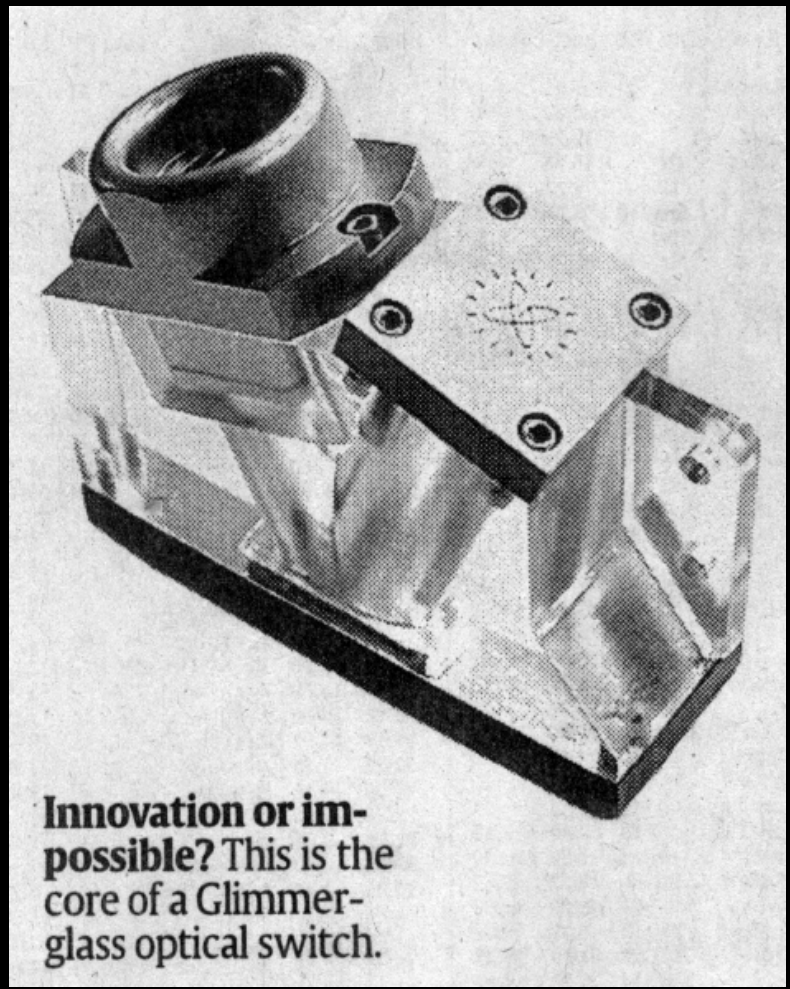

Innovation or impossible? This is the core of a Glimmerglass optical switch.

StarLight as a Research Network
Optical Switches at StarLight and NetherLight

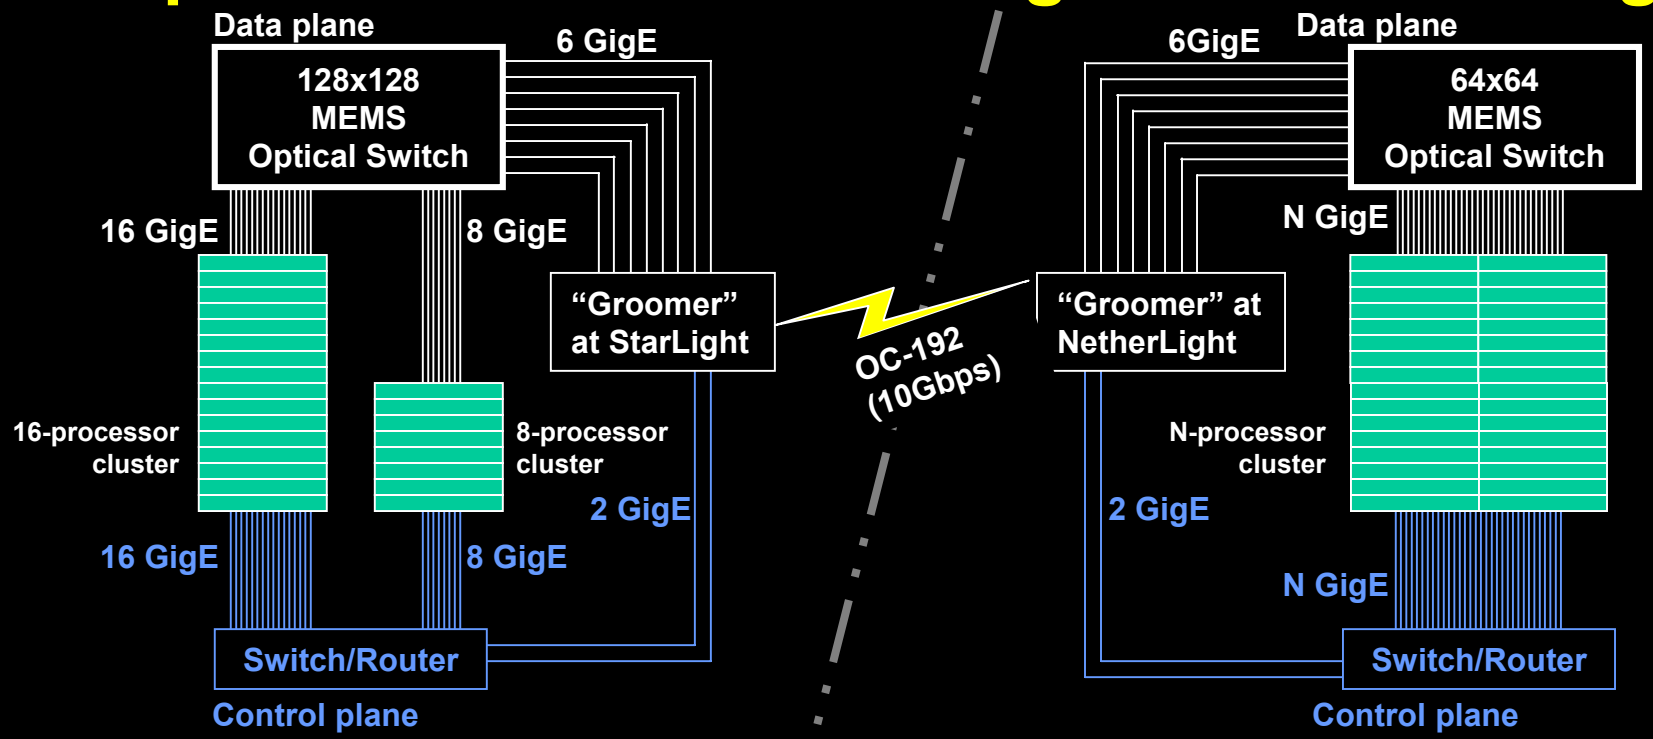

STARLIGHT℠

NETHERLIGHT

A "groomer" is a box that accepts multiple circuits of varying types (e.g., 1GigE, 10GigE) and aggregates and/or disseminates over the 10Gbps transoceanic link. *As the amount of transoceanic connectivity increases, we aim to "bandwidth match" the amount of data being sent and/or received by clusters across continents.*

Thanks to…

- **StarLight planning, research, collaborations, and outreach efforts are made possible, in major part, by funding from:**
  - National Science Foundation (NSF) awards ANI-9980480, ANI-9730202, EIA-9802090, EIA-9871058, ANI-0225642, and EIA-0115809
  - NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement ACI-9619019 to NCSA
  - State of Illinois I-WIRE Program, and major UIC cost sharing
  - Northwestern University for providing space, $$, engineering and management
- **NSF/CISE/ANIR and DoE/Argonne National Laboratory for StarLight and I-WIRE network engineering and planning leadership**
- **NSF/CISE/ACIR and NCSA/SDSC for DTF/TeraGrid/ETF opportunities**
- **UCAID/Abilene for Internet2 and ITN transit; IU for the GlobalNOC**
- **Bill St. Arnaud of CANARIE, Kees Neggers of SURFnet, Olivier Martin of CERN Michael McRobbie of IU, and Harvey Newman of CalTech for networking leadership**

For More Information

Maxine Brown
maxine@uic.edu

www.startap.net/starlight
www.startap.net/translight