Distributed Virtual Network Operations Center (DVNOC) - Towards Federated & Customer-focused Cyberinfrastructure

Harika Tandra, Software Engineer
GLORIAD

(presentation based on slides prepared by Greg Cole, Principal Investigator, GLORIAD project)

What is GLORIAD ?

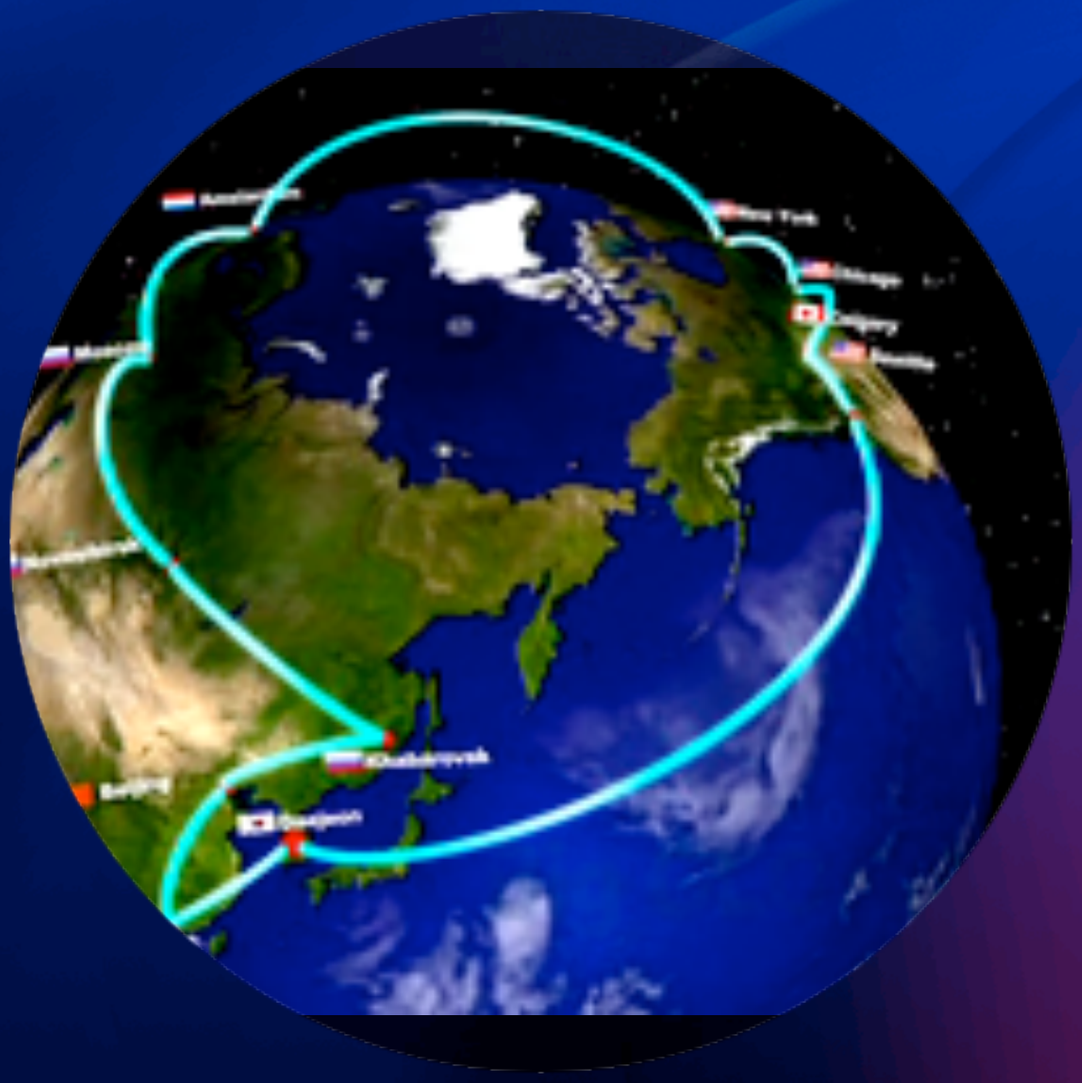

- A cooperative R&E network ringing the northern hemisphere linking scientists, educators and students in Russia, USA, China, Korea, Netherlands, Canada, the Nordic countries – and soon India, Egypt, Singapore – and others with specialized network services; co-funded, co-managed by all international partners
- Variously sized circuits/services arounnd northen hemisphere
- Hybrid circuit-(L1/L2) and packet-switched services(L3)
- Collaborative International Program to Develop/Deploy advanced Cyberinfrastructure and appliations between partnering countries (and others) as effort to expand science, education and cultural cooperation and exchange

GLORIAD MAP

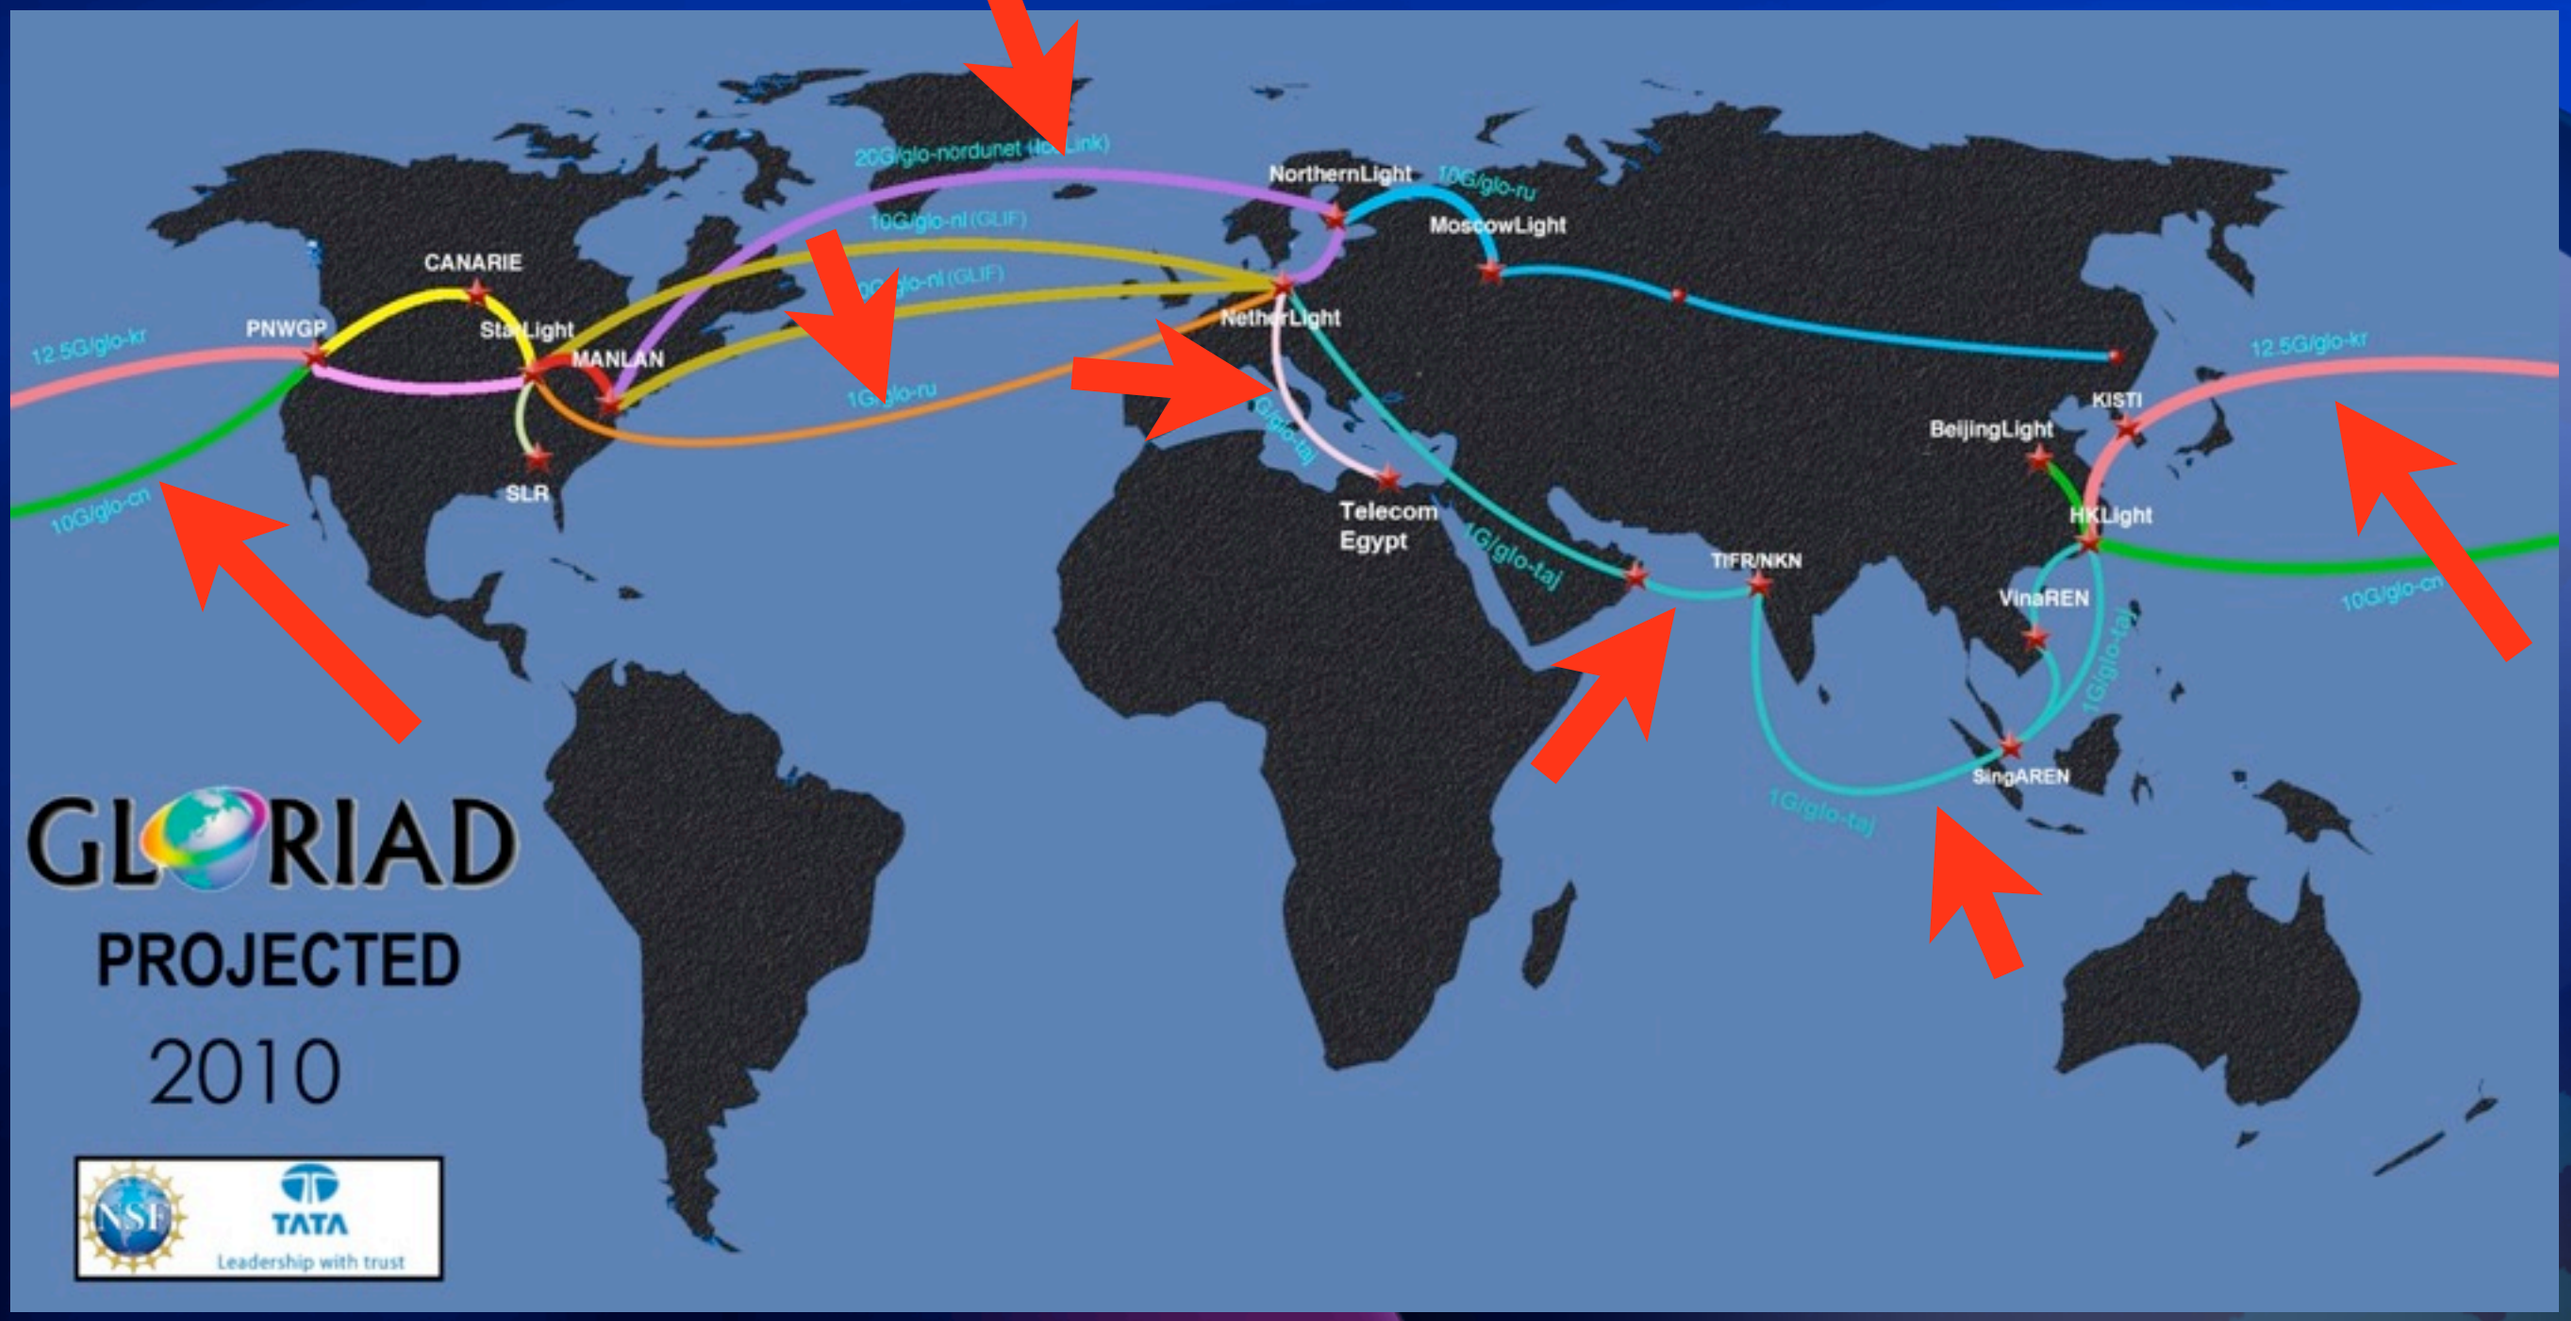
20G/glo-nordunet (Hot Link)
10G/glo-nl (GLIF)
10G/glo-nl (GLIF)
1G/glo-ru
12.5G/glo-kr
10G/glo-cn
CANARIE
PNWGP
StarLight
MANLAN
SLR
NorthernLight
10G/glo-ru
MoscowLight
NetherLight
1G/glo-taj
Telecom Egypt
1G/glo-taj
TIFR/NKN
VinaREN
1G/glo-taj
SingAREN
1G/glo-taj
KISTI
12.5G/glo-kr
BeijingLight
HKLight
10G/glo-cn
GLORIAD
PROJECTED
2010
NSF
TATA
Leadership with trust

GLORIAD mission

- Connecting the unconnected
- Better informing science and education community (and general public) about global opportunities for collaboration
- Promoting decentralized, distributed, transparent and open approach to global R&E networking

DVNOC Tool

Wednesday, February 9, 2011

DVNOC

- Addresses need for all levels of cyberinfrastructure operators (and users) to collaborate on decentralized, distributed and reliable operations of links and services
- Focus on customer-based performance
- Large development effort on part of Chinese, Dutch, Korean, Nordic and US (and we hope, soon, other national) GLORIAD teams

DVNOC Contd..

- Web based application
- Developed using Flash/Flex platform
- Current version: http://viz.gloriad.org/dvnoc/dvnoc.html

DVNOC

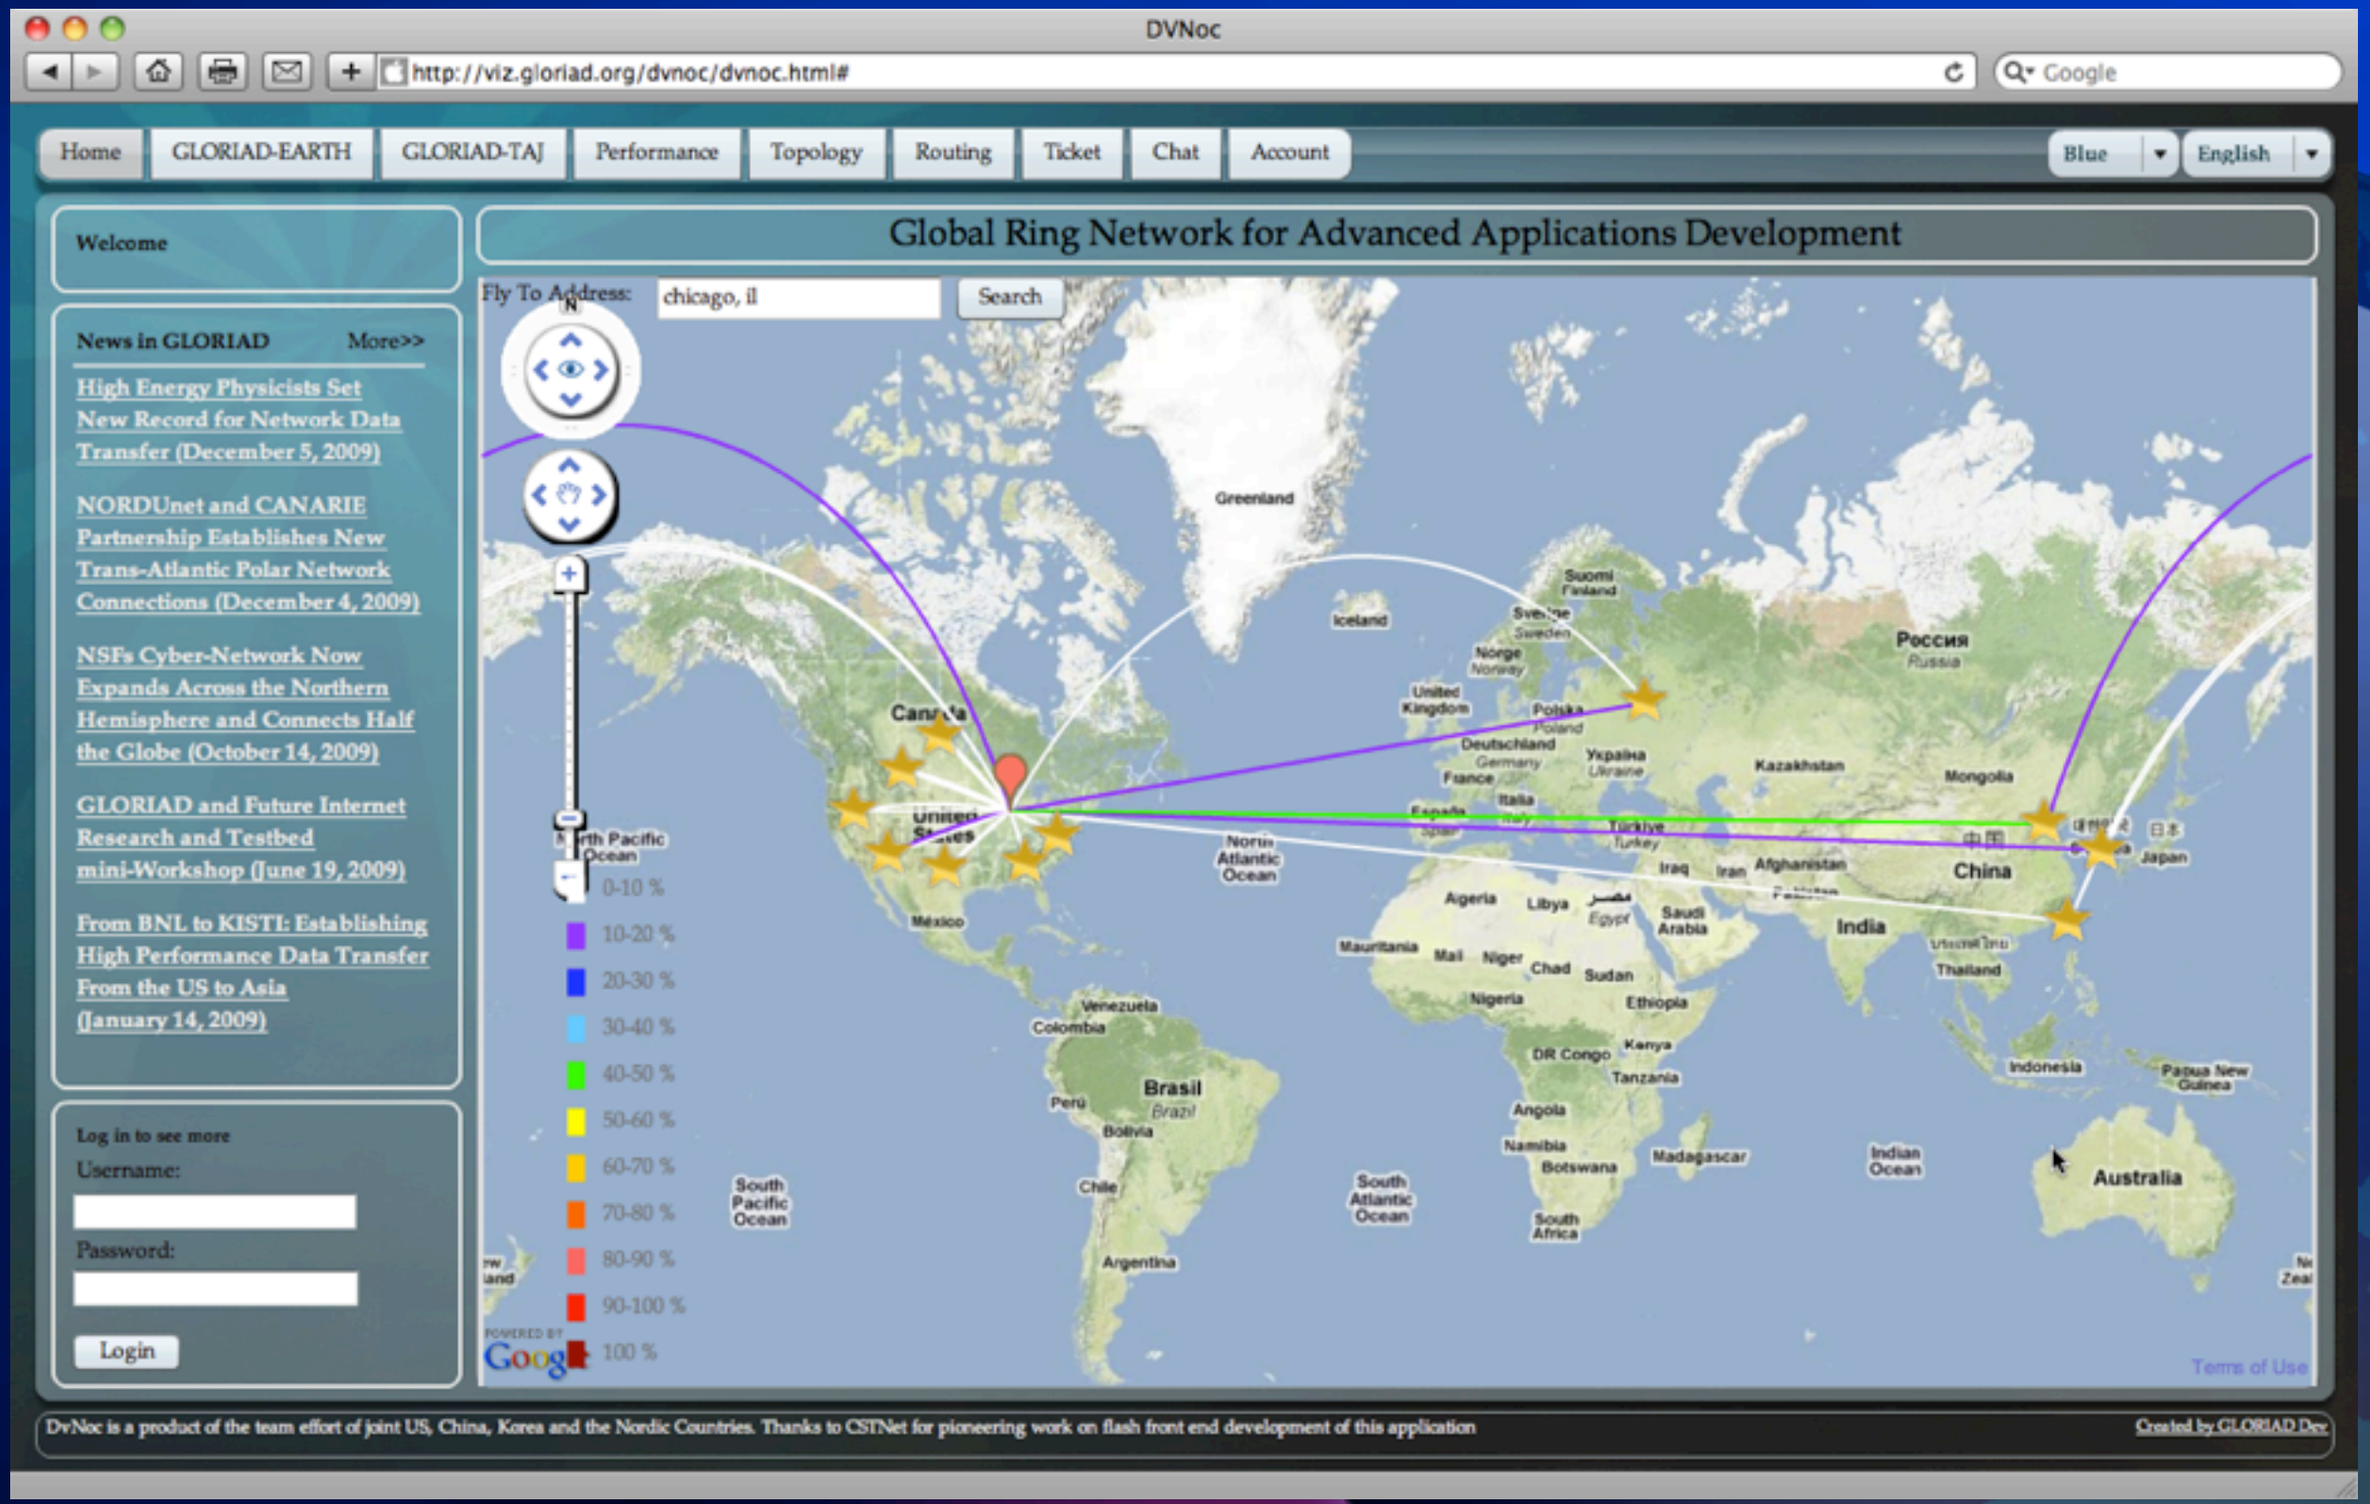
DVNoc
http://viz.gloriad.org/dvnoc/dvnoc.html#
Google
Home
GLORIAD-EARTH
GLORIAD-TAJ
Performance
Topology
Routing
Ticket
Chat
Account
Blue
English
Welcome
Global Ring Network for Advanced Applications Development
Fly To Address:
chicago, il
Search
News in GLORIAD
More>>
High Energy Physicists Set New Record for Network Data Transfer (December 5, 2009)
NORDUnet and CANARIE Partnership Establishes New Trans-Atlantic Polar Network Connections (December 4, 2009)
NSFs Cyber-Network Now Expands Across the Northern Hemisphere and Connects Half the Globe (October 14, 2009)
GLORIAD and Future Internet Research and Testbed mini-Workshop (June 19, 2009)
From BNL to KISTI: Establishing High Performance Data Transfer From the US to Asia (January 14, 2009)
Log in to see more
Username:
Password:
Login
0-10 %
10-20 %
20-30 %
30-40 %
40-50 %
50-60 %
60-70 %
70-80 %
80-90 %
90-100 %
100 %
Greenland
Iceland
Suomi Finland
Sverige Sweden
Norge Norway
United Kingdom
Polska Poland
Deutschland Germany
France
Italia
Україна Ukraine
Россия Russia
Kazakhstan
Mongolia
China
Japan
India
Thailand
Indonesia
Papua New Guinea
Australia
Afghanistan
Iraq
Iran
Türkiye Turkey
Saudi Arabia
Egypt
Libya
Algeria
Mauritania
Mali
Niger
Chad
Sudan
Nigeria
Ethiopia
Kenya
DR Congo
Tanzania
Angola
Namibia
Botswana
South Africa
Madagascar
Indian Ocean
South Atlantic Ocean
South Pacific Ocean
North Pacific Ocean
North Atlantic Ocean
México
Venezuela
Colombia
Perú
Brasil Brazil
Bolivia
Chile
Argentina
Terms of Use
DvNoc is a product of the team effort of joint US, China, Korea and the Nordic Countries. Thanks to CSTNet for pioneering work on flash front end development of this application
Created by GLORIAD Dev

DVNOC - GLORIAD Earth Tab

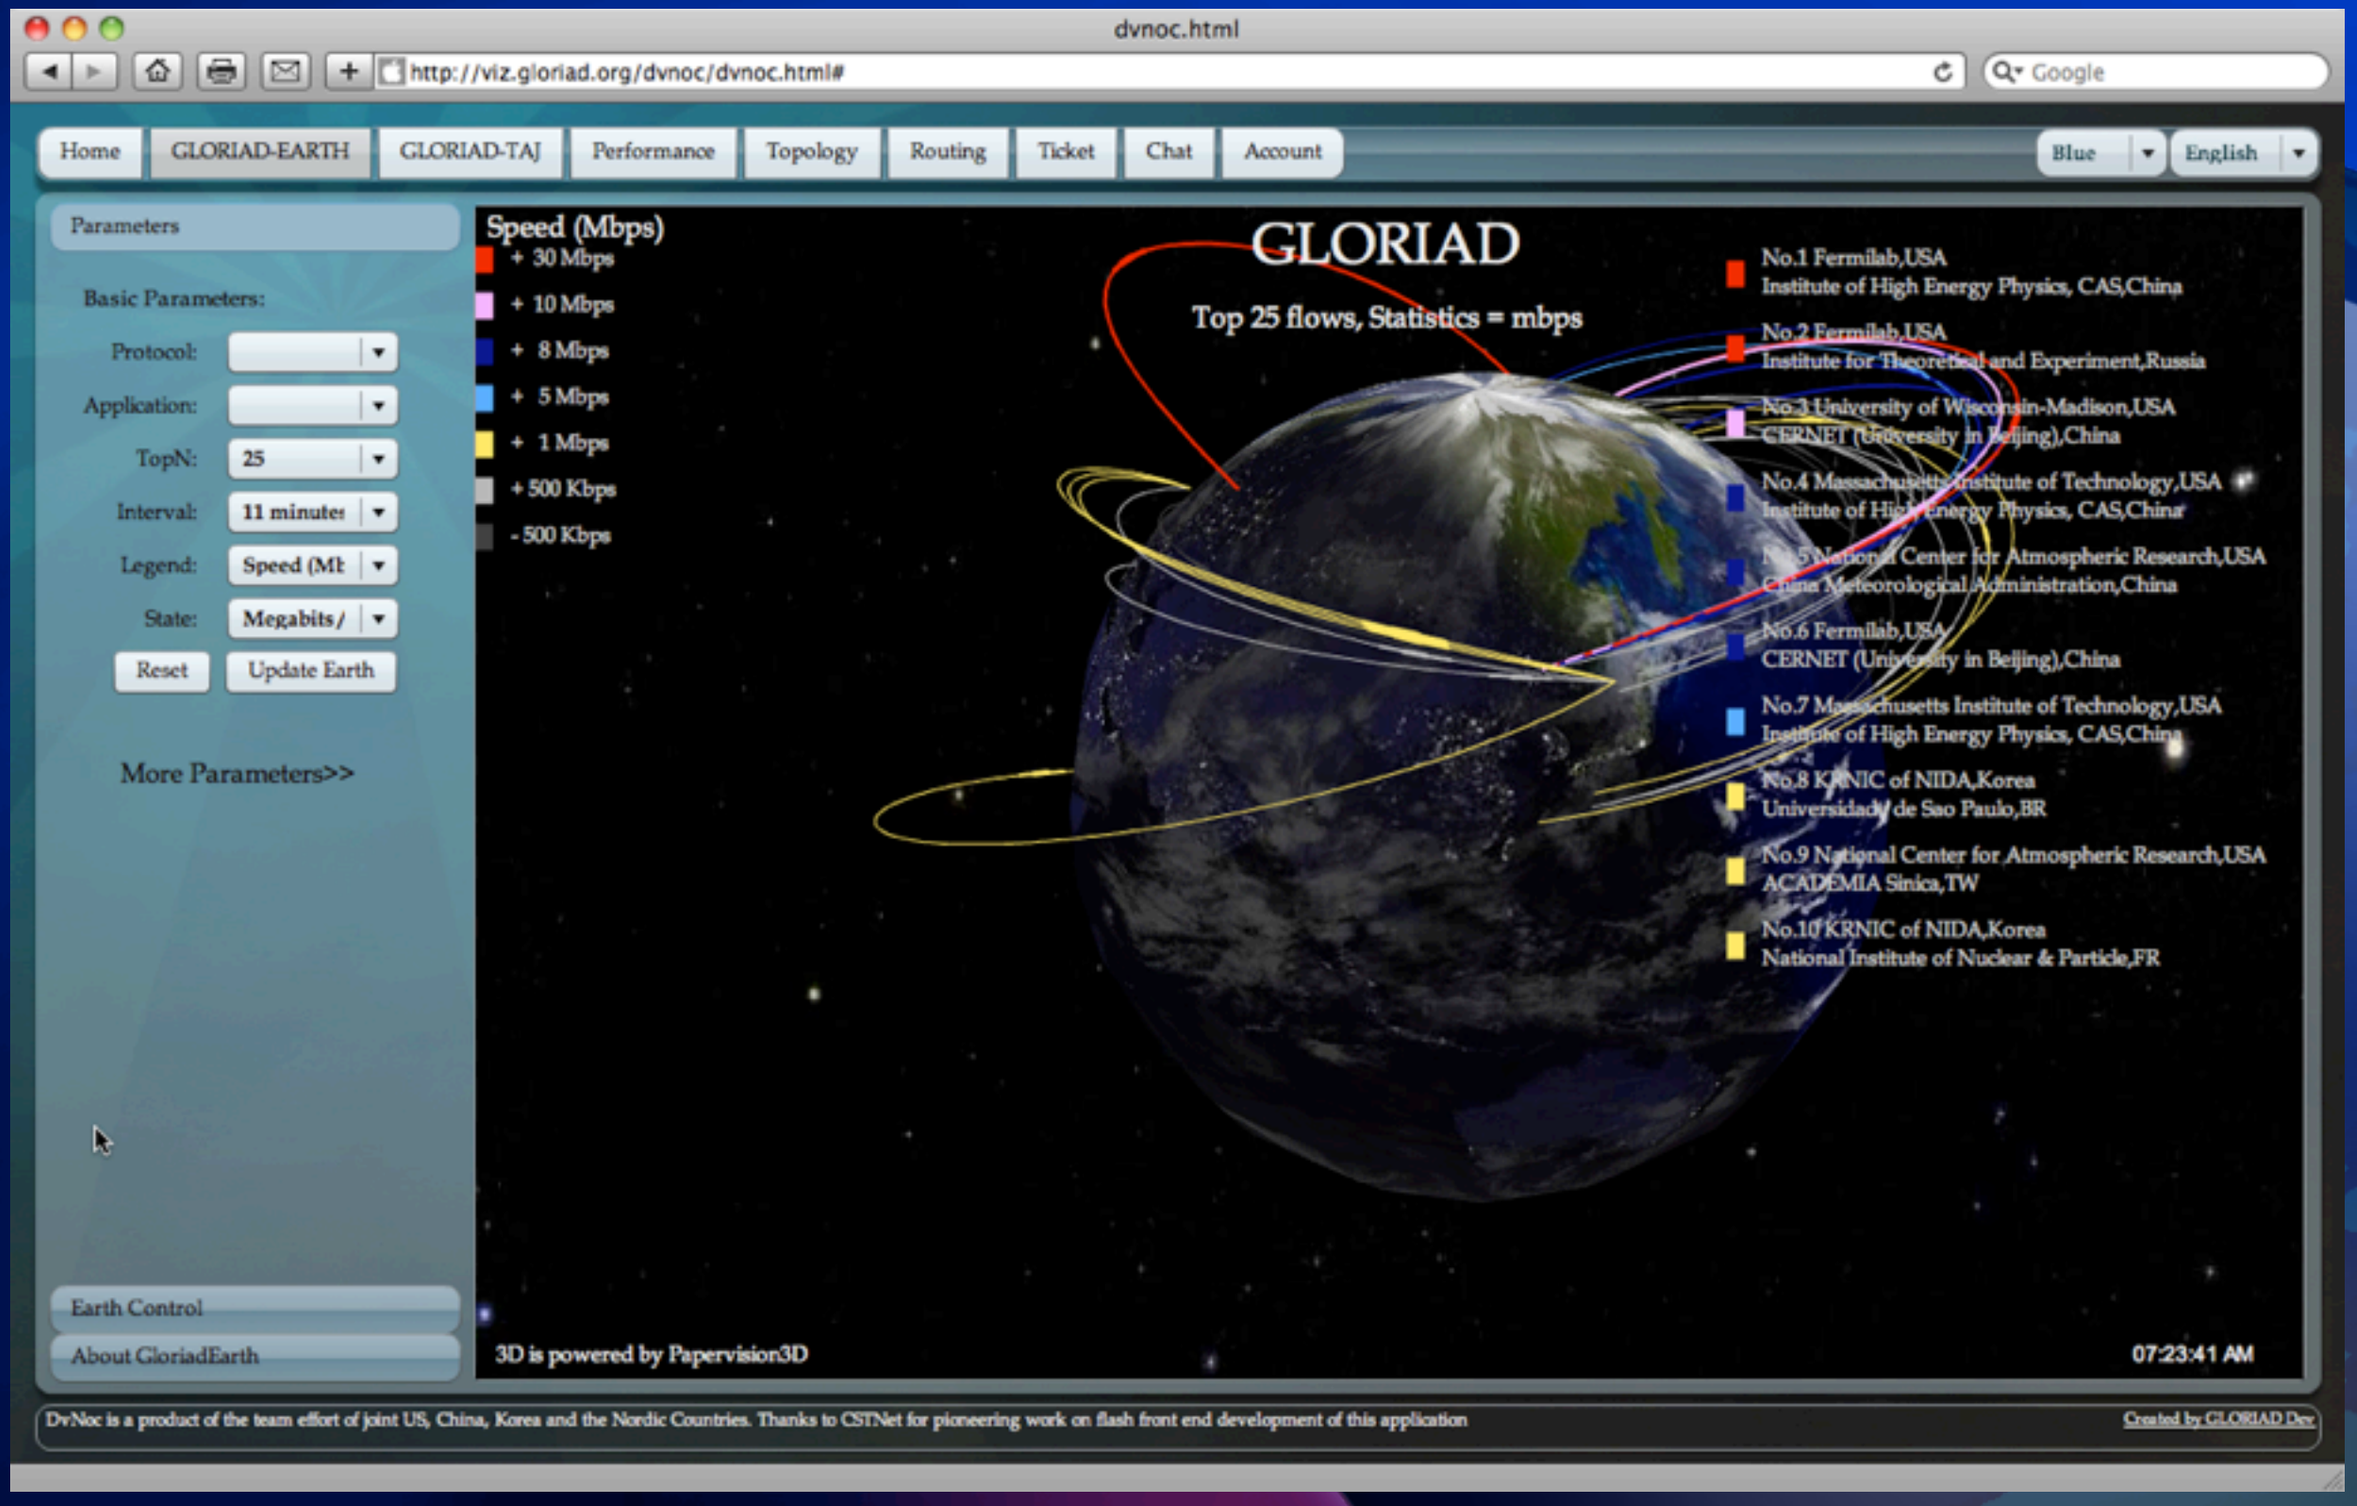

DVNOC - GLORIAD Earth Tab

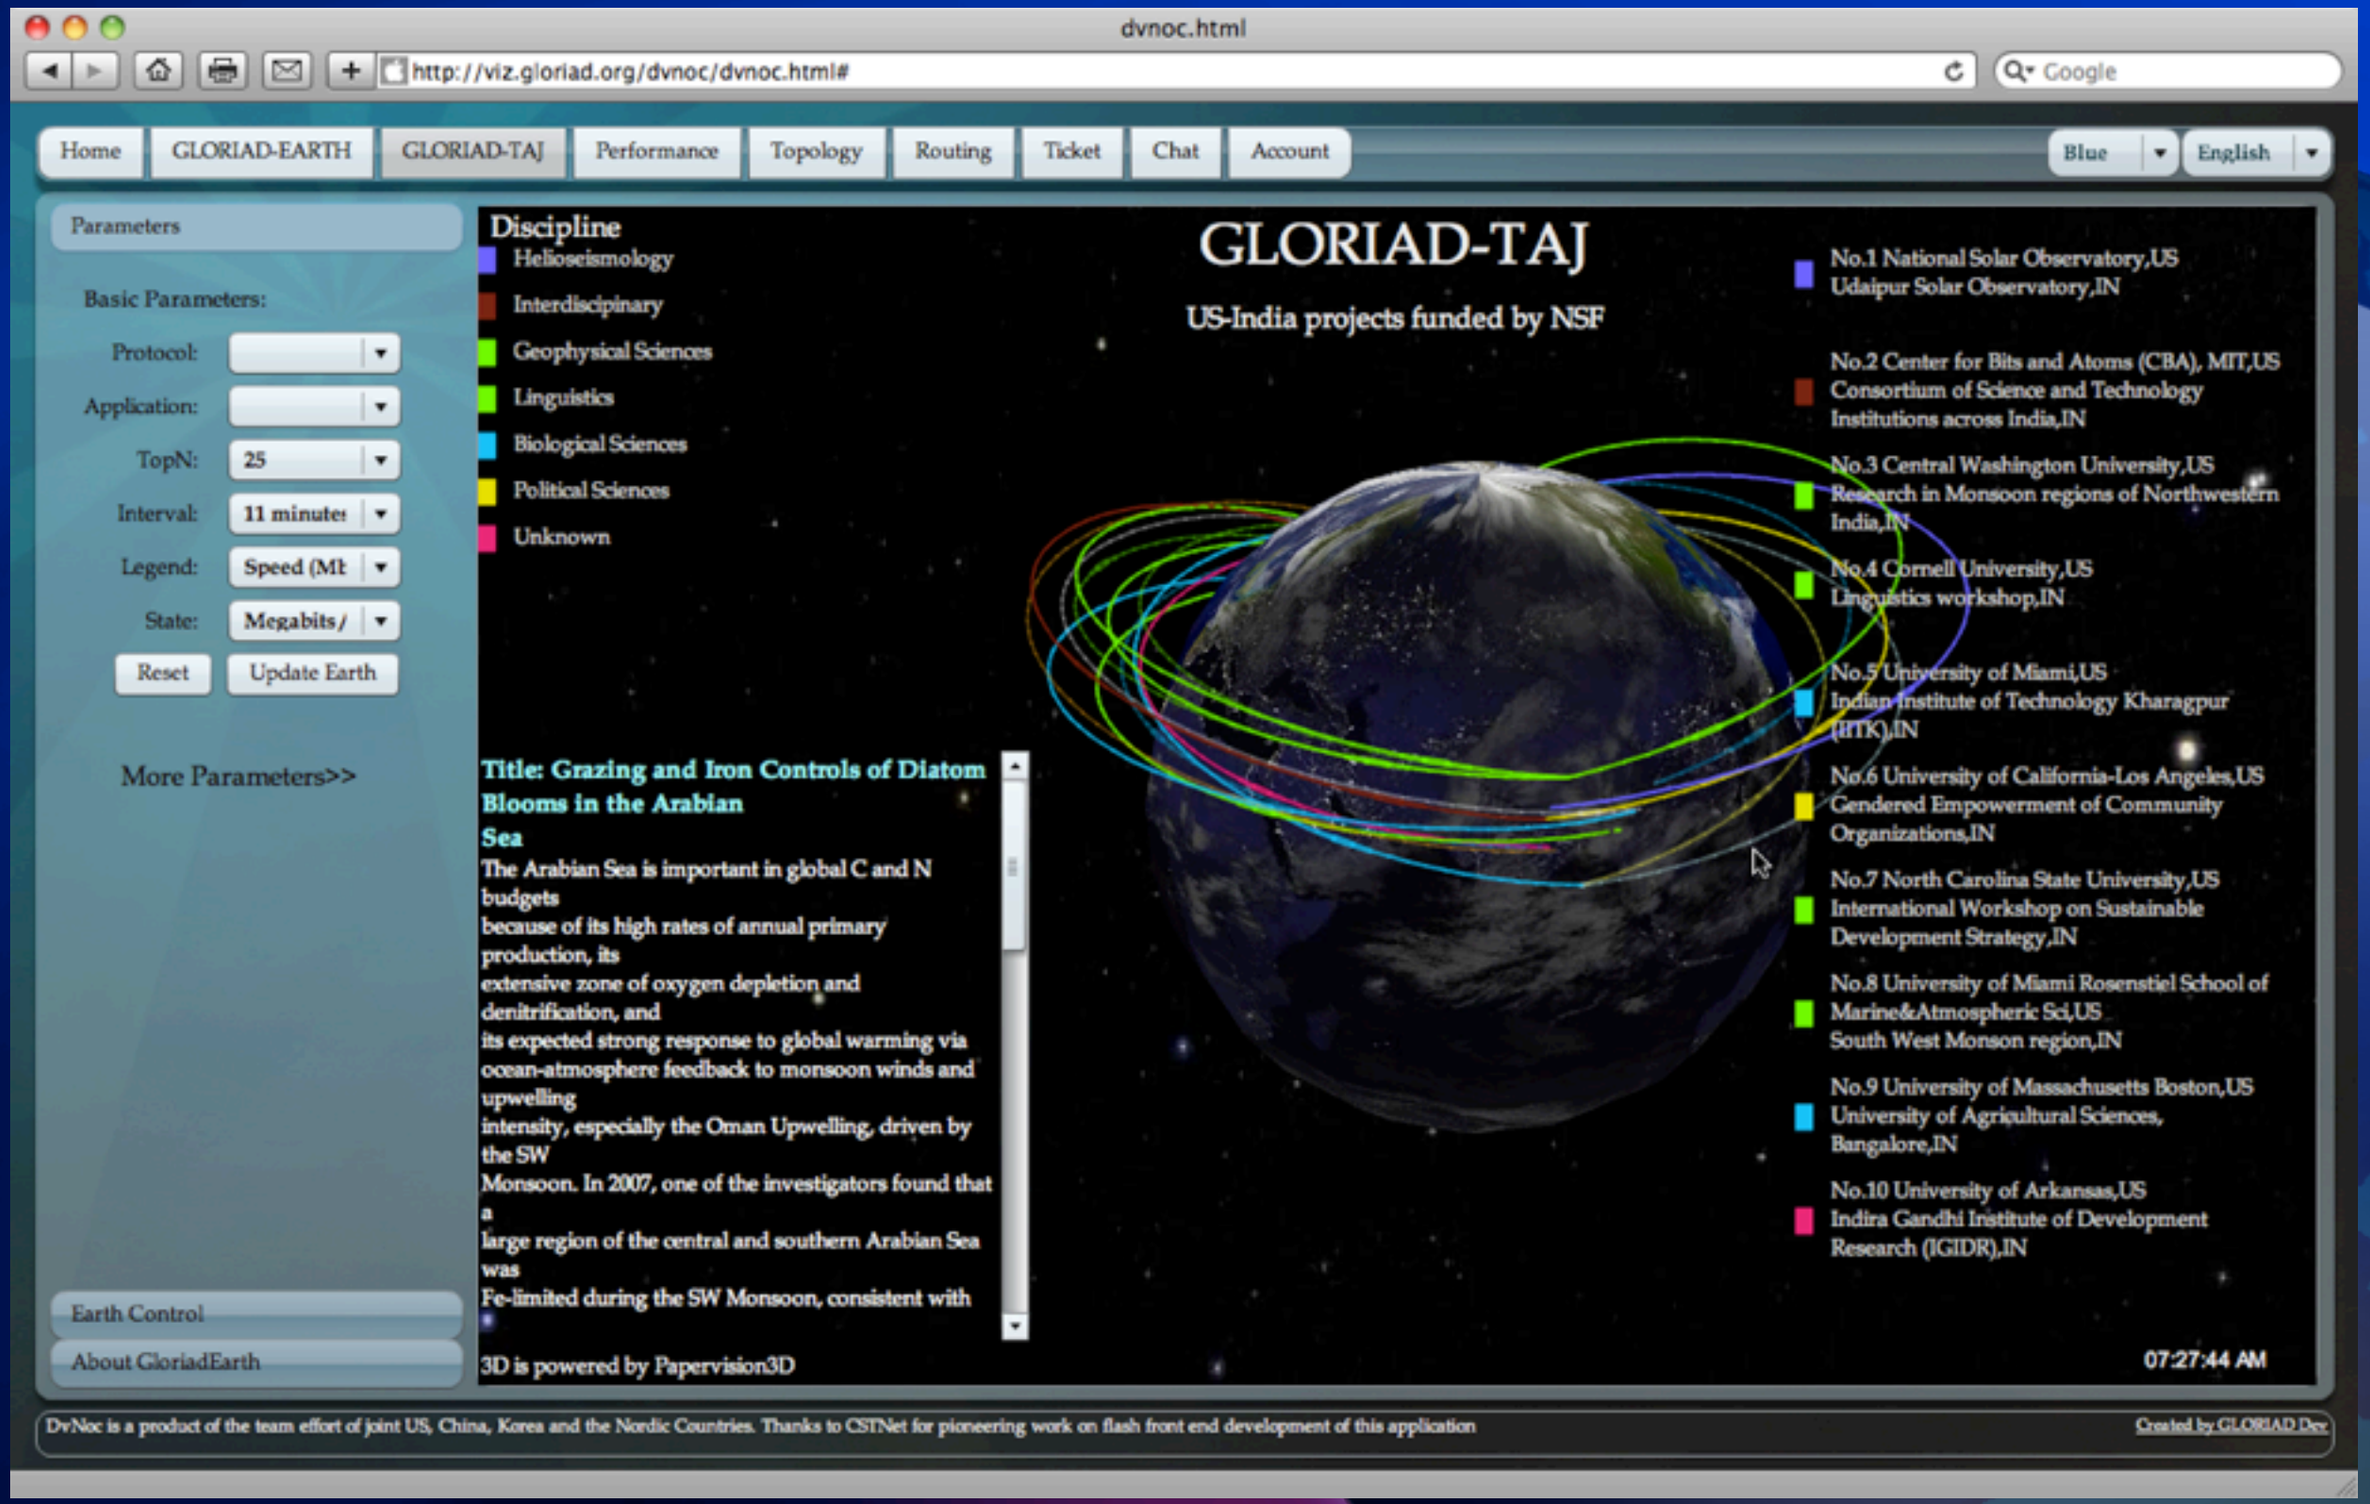

Performance Measurement

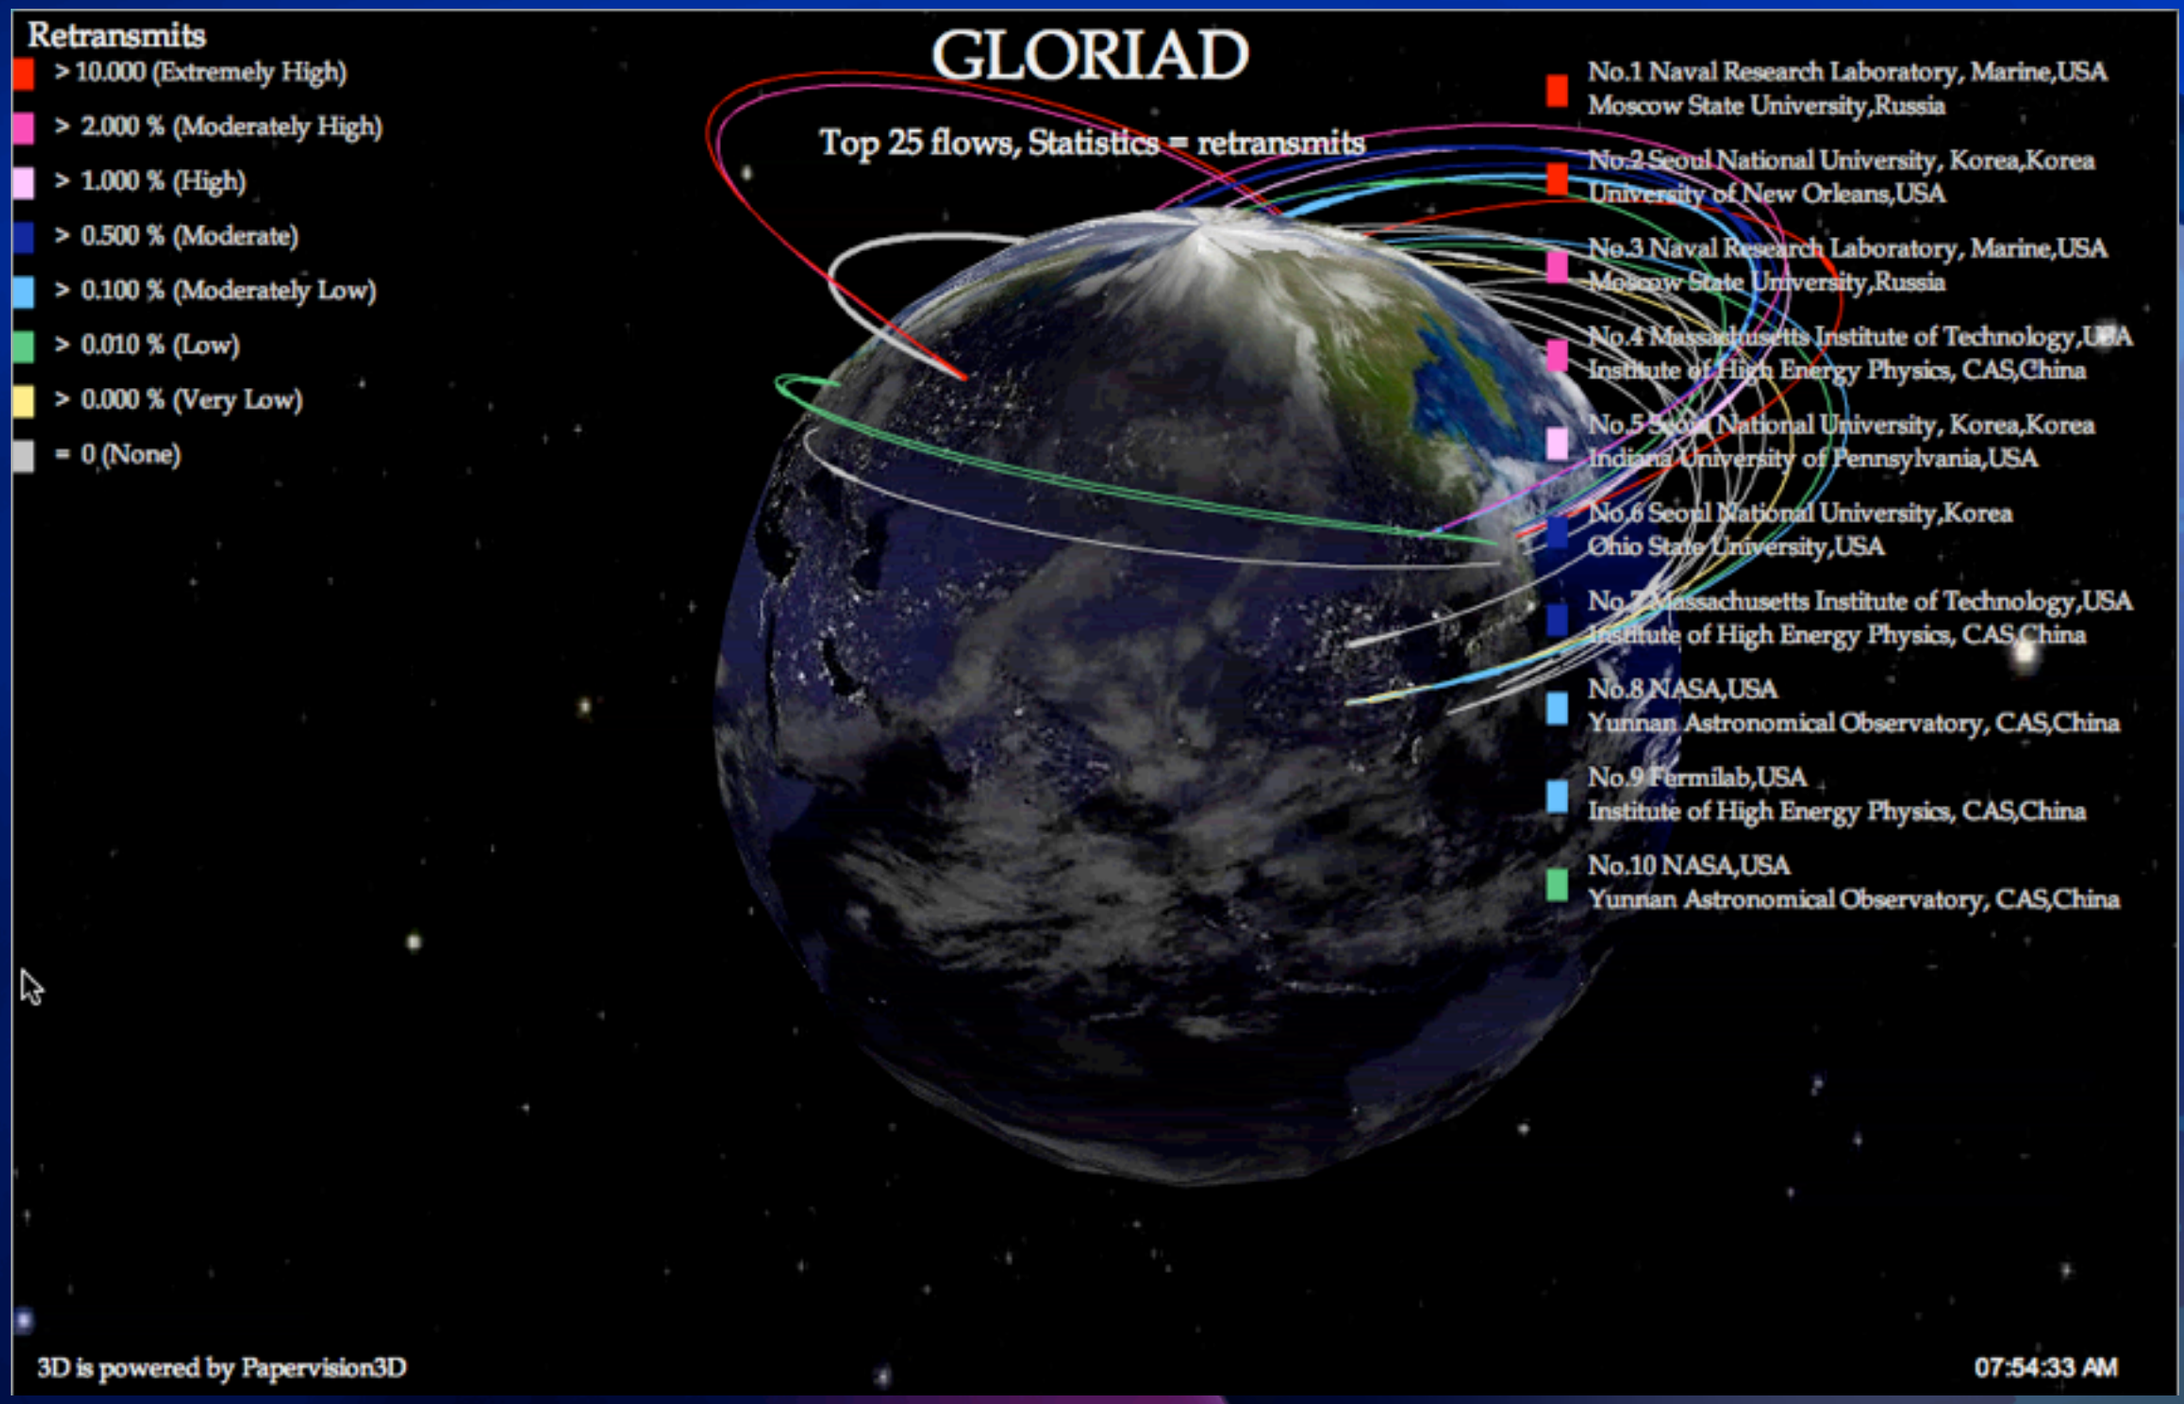

We're trying to shift towards "customer-based performance" in all areas of cyberinfrastructure deployment

X-axis: %loss

Y-axis: RTT(ms)

Z-axis: throughput in bits/sec

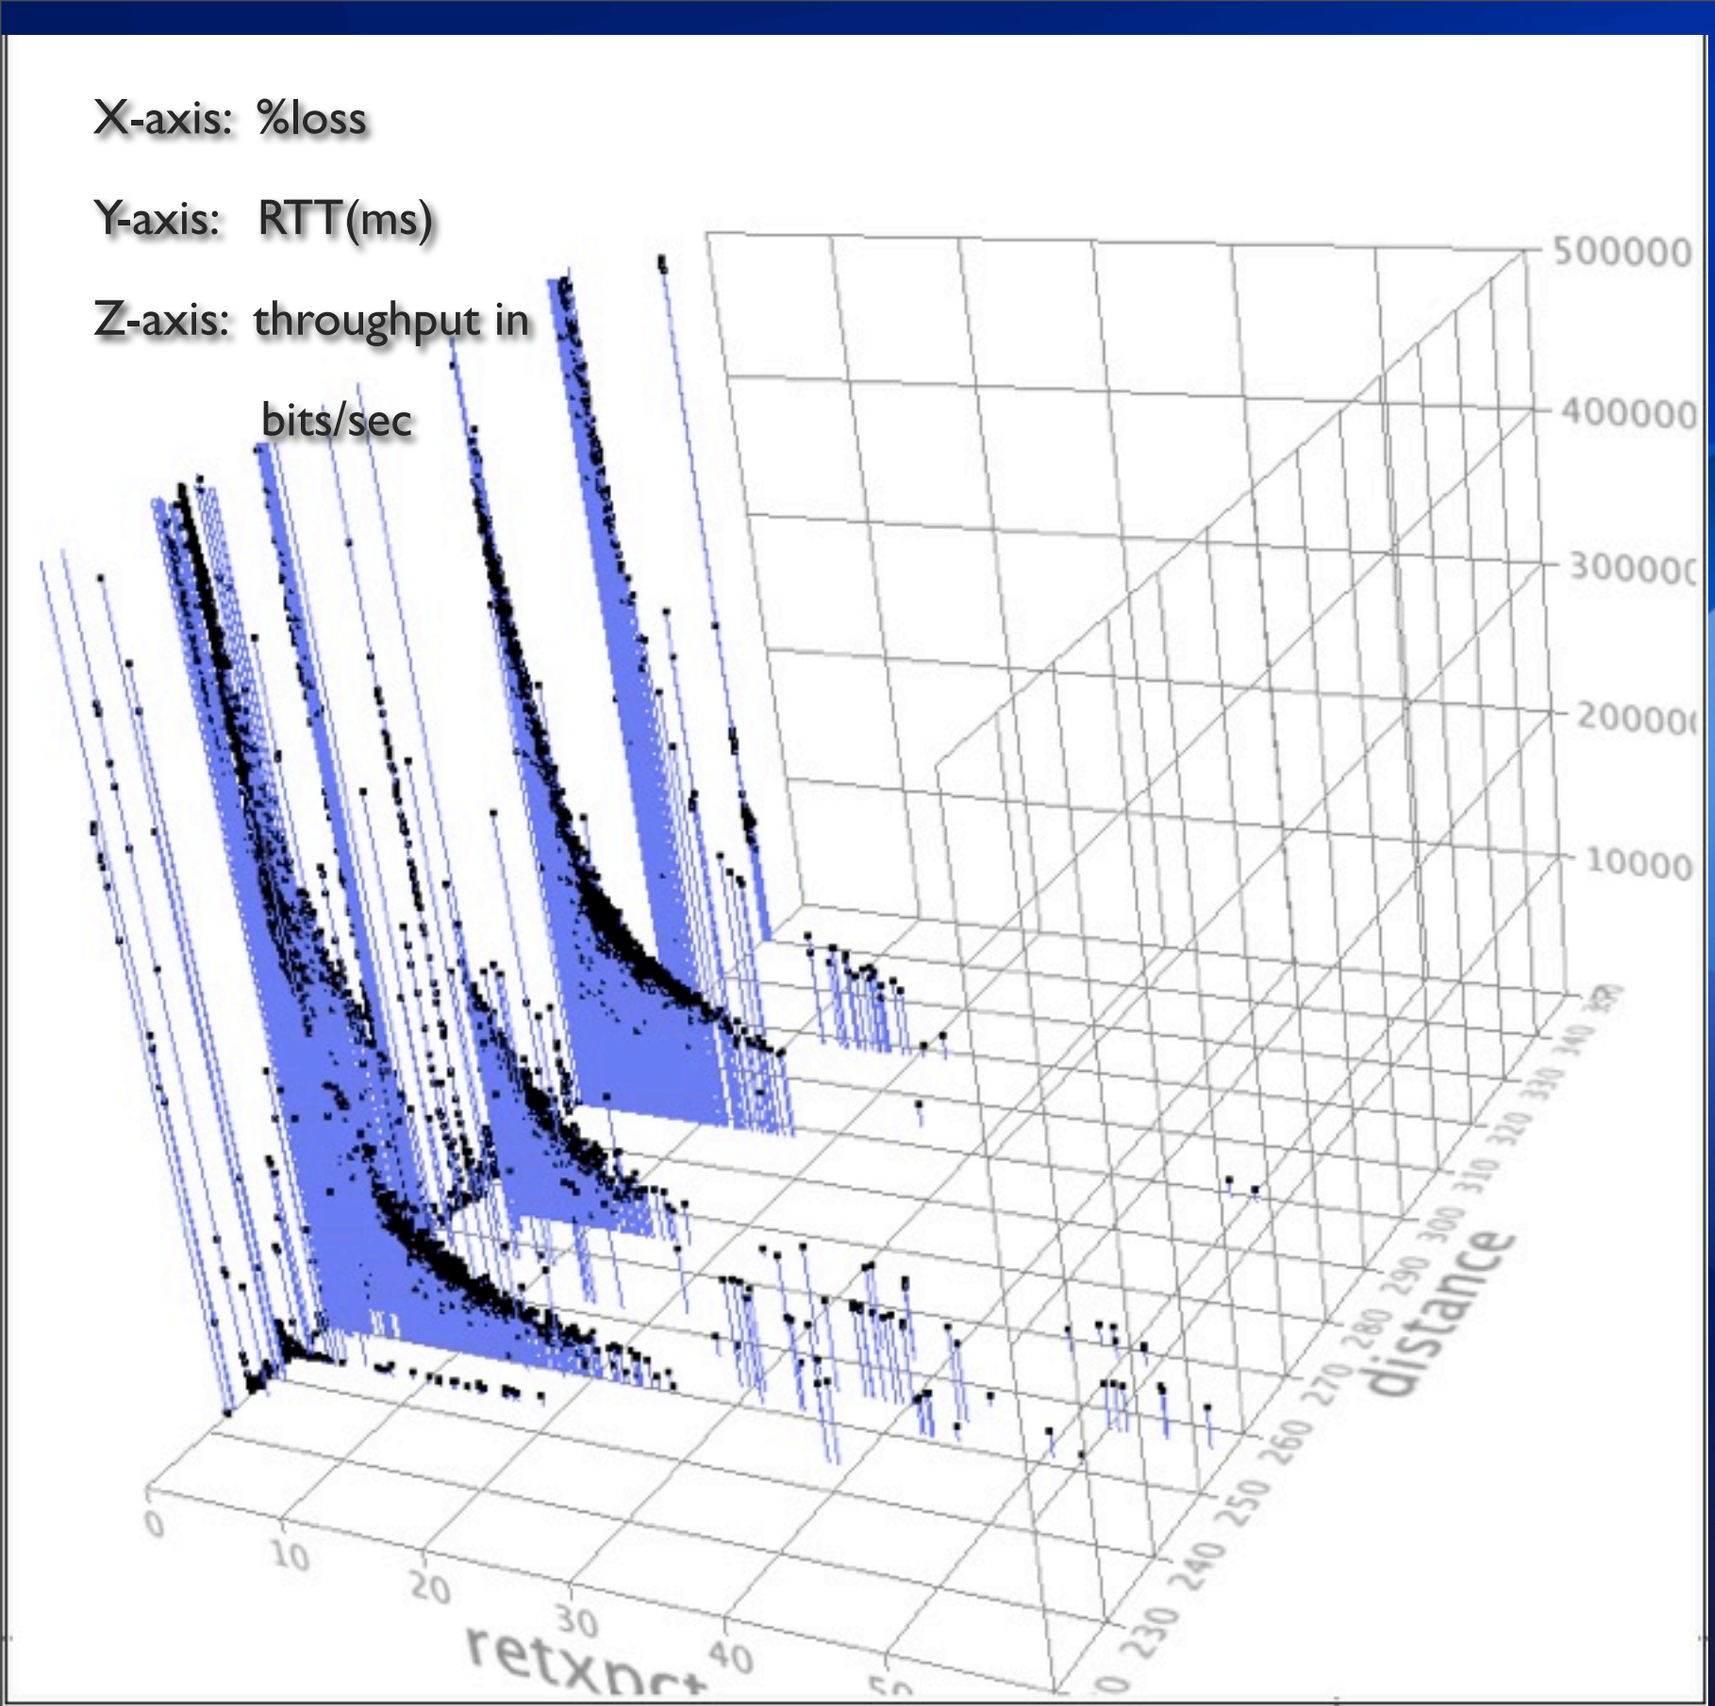

"Needle" chart i.e., a blue needle (topped by a black marker) illustrates one flow

3-D plot of throughput , loss & RTT using flow data from US to CSTNET over a 24hr period on GLORIAD network

Identifying Problem Areas in Global, National, Regional, Local, Campus Networks

- Problem: network operators have insufficient knowledge of nor relationship with each other (local/campus, regional, national, global operators) (and R&E customers less so)
- Solution: encourage common view towards customer-based performance, lead effort towards community-developed shared performance measurement instrumentation and tools for joint engineering management (dvNOC)
- (we will realize many other benefits from this community-building exercise)

Emphasis on Customer Performance

- We wish to know of individual customer-based performance problems before customer can call
- We're developing statistically important base of information about where there are weaknesses in our global/regional/ regional/local networks
- Based primarily (at moment) on measurements of packet retransmits

Automated system to debug under-performing flows in wide area networks

Throughput vs Loss (contd..)

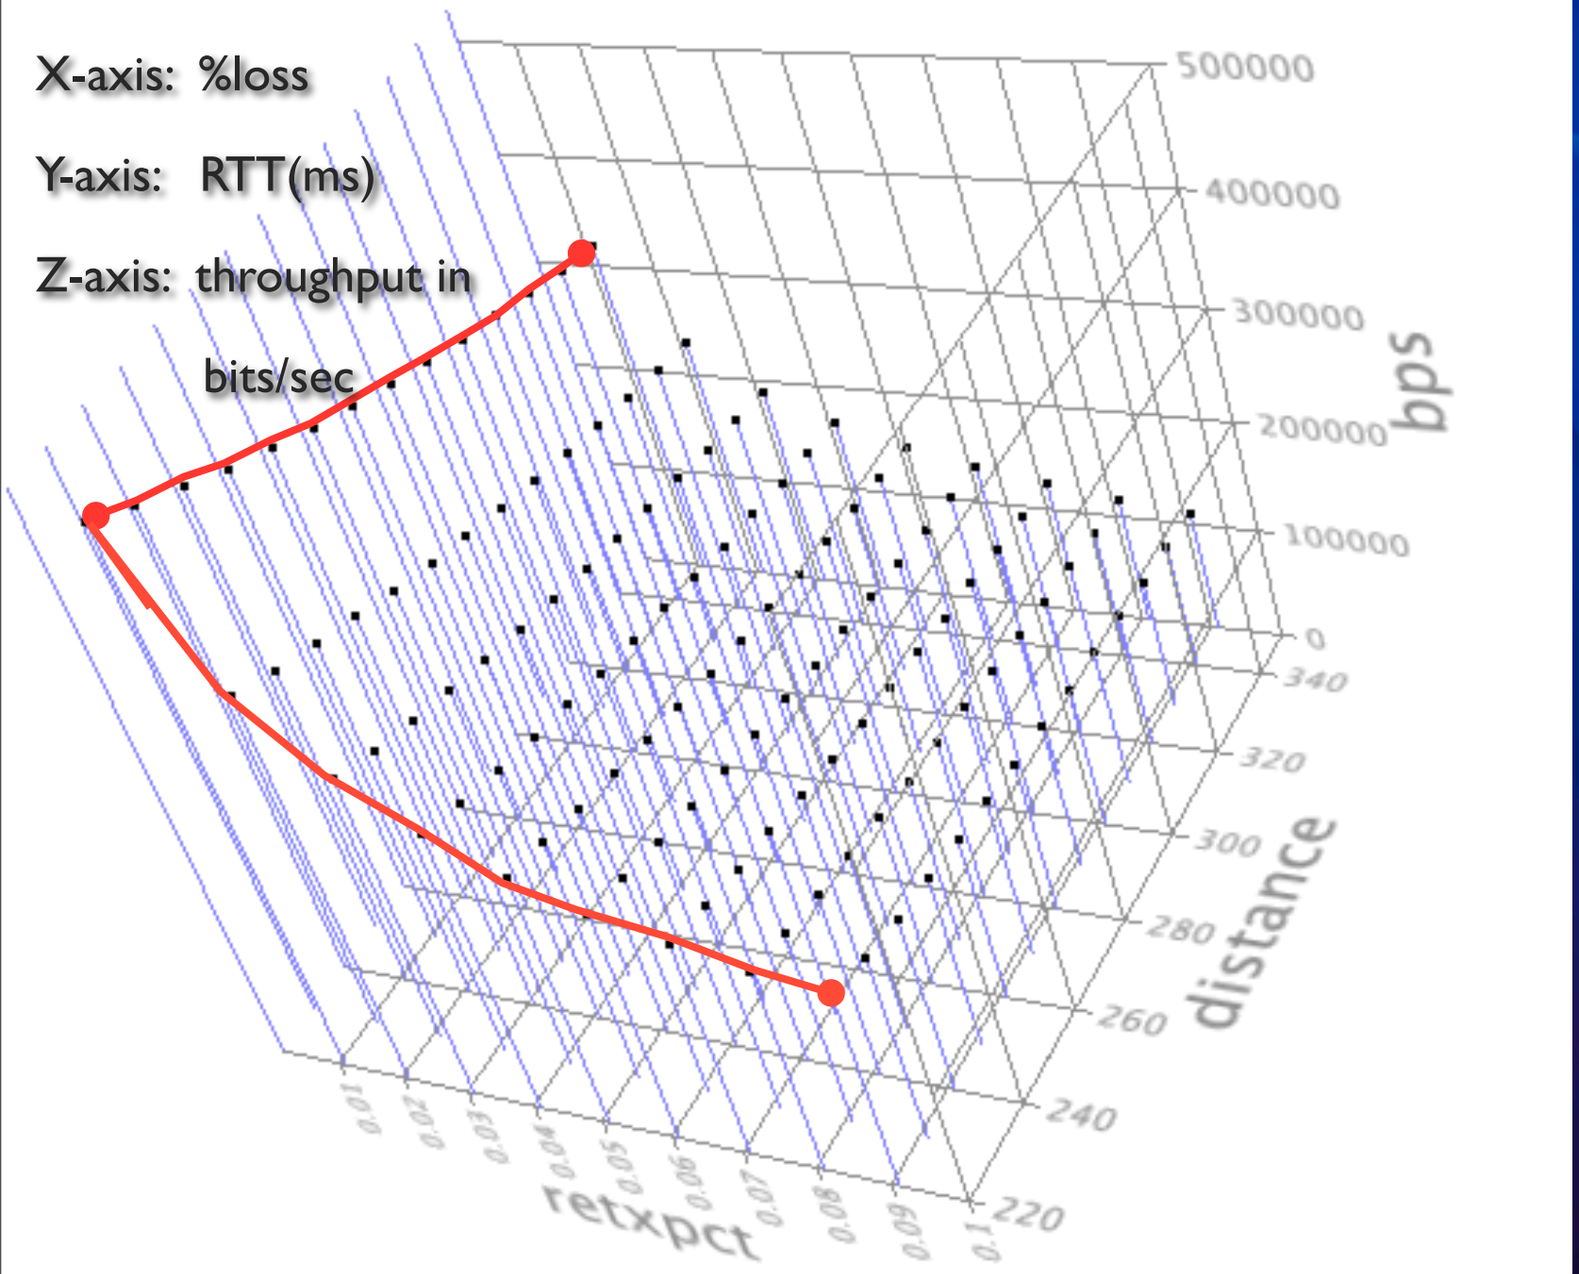

3-D plot of throughput derived from loss & RTT using Mathis formula

- We can see that the decrease in rate is steeper with the increase in loss than the increase in RTT
- Half the loss rate gives throughput increase of ~41%

Hybrid monitoring/data collection system

1. **Passive monitoring sub-system:** Filters network flow data to identify under-performing flows
2. **Active monitoring sub-system:** Collects performance statistics of individual routers

**All the IPs are anonymized in the following slides

Passive monitoring sub-system : Flow filter

- % retransmissions per bytes transfered > 0.01
- Bytes transfered > 5 MB
- Frequency > 4 hours. Same (ip_s, ip_d) pair is not labeled as under-performing for the minimum time period set by the frequency parameter

Passive monitoring sub-system

- Filter the netflow records to identify under-performing flows

ip_src	ip_dst	MB	%rtpct	starttime	endtime
xx.xx.77.70	xx.xx.138.244	0.2137	0	2009-10-18 20:53:43	2009-10-18 20:58:50
xx.xx.16.49	xx.xx.4.71	0.2101	0	2009-10-18 20:53:42	2009-10-18 20:58:51
xx.xx.189.65	xx.xx.224.75	213.3897	0	2009-10-18 20:15:22	2009-10-18 20:58:51
xx.xx.3.226	**xx.xx.244.210**	**7.3098**	**0.9866**	**2009-10-18 20:56:50**	**2009-10-18 20:58:51**

MB - MBytes transfered, %rtpct - Percentage retransmissions per byte

Active monitoring sub-system

- For each under-performing flow identified, MTR runs are triggered to source and destination IPs
- Triggered in **near-real-time** to the flow detected. Thus, test packets are triggered in network conditions similar to those seen by the real traffic
- Combining the two gives approximate end-to-end performance

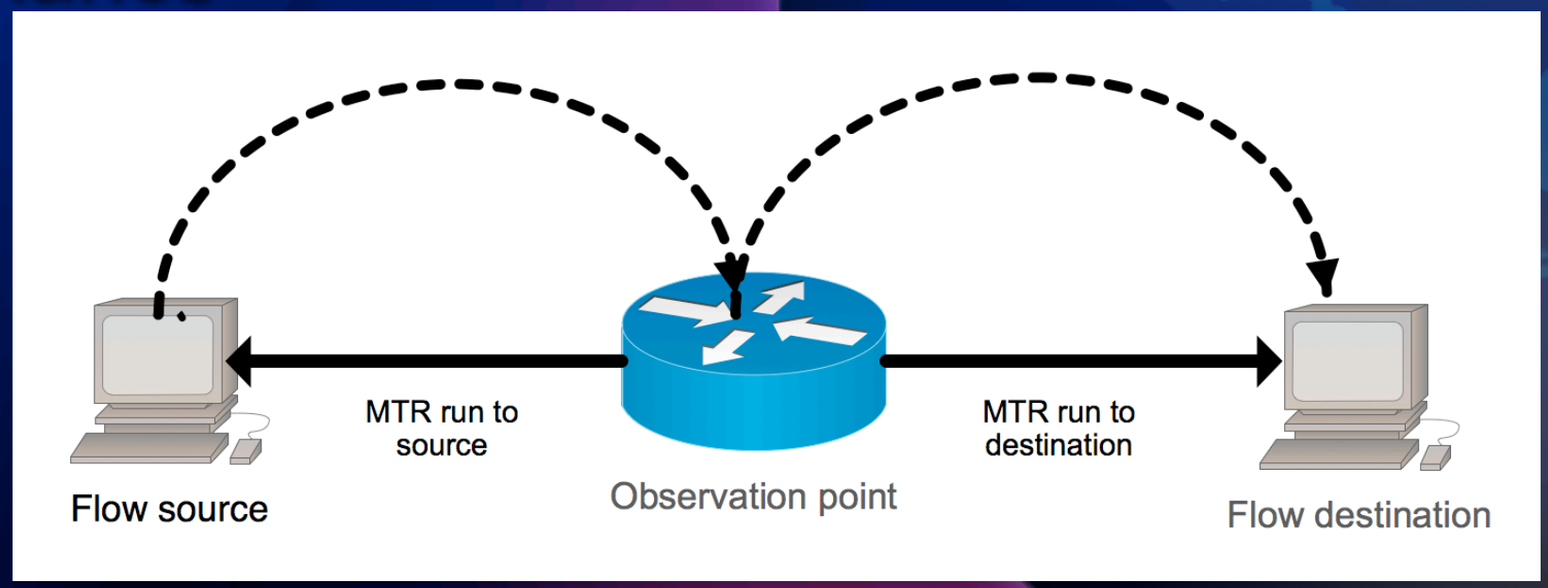

Data collected

ip_s	ip_d	MBytes	rtpct	starttime	endtime	keyid
xx.3.226	xx.244.210	7.31	0.987	2009-10-18 20:56:50	2009-10-18 20:58:51	28995

serial_n	node_ip	loss_pct	packets_s	avg_rtt	best_rtt	wrst_rtt	target_ip	masterkeyid	target_lbl
1	192.31.99.97	0	50	4.7	0.4	12.8	xx.244.210	28995	Destination
2	192.31.99.146	0	50	2.3	1.3	17.7	xx.244.210	28995	Destination
3	216.24.186.5	0	50	28.8	27.2	49.3	xx.244.210	28995	Destination
4	192.43.217.137	0	50	28.6	26.3	62.8	xx.244.210	28995	Destination
5	192.43.217.114	0	50	27.2	27.1	27.7	xx.244.210	28995	Destination
6	128.117.243.75	0	50	27.8	27.2	41.7	xx.244.210	28995	Destination
7	???	100	50	0	0	0	xx.244.210	28995	Destination
1	192.31.99.97	0	50	5.3	0.4	13.5	xx.3.226	28995	Source
2	192.31.99.166	0	50	189.9	189.8	196.3	xx.3.226	28995	Source
3	159.226.254.165	0	50	190.2	189.9	203.6	xx.3.226	28995	Source
4	159.226.254.253	0	50	228.8	228.8	229.1	xx.3.226	28995	Source
5	159.226.254.29	0	50	230.8	228.9	317.8	xx.3.226	28995	Source
6	159.226.254.190	2	50	229.6	229	254.1	xx.3.226	28995	Source
7	159.226.254.170	4	50	229.4	229.2	229.8	xx.3.226	28995	Source
8	159.226.46.230	2	50	229.6	229.4	230.3	xx.3.226	28995	Source
9	???	100	50	0	0	0	xx.3.226	28995	Source

Resultset

Result of MTR runs to source and destination of an under-performing flow

Data interpretation

- Network graphs show individual router behavior cutting across several MTR runs, to different target IPs
- Thus, giving a snap shot of network router topology seen by the under-performing flows

Example network graphs for a few end hosts in U.S.

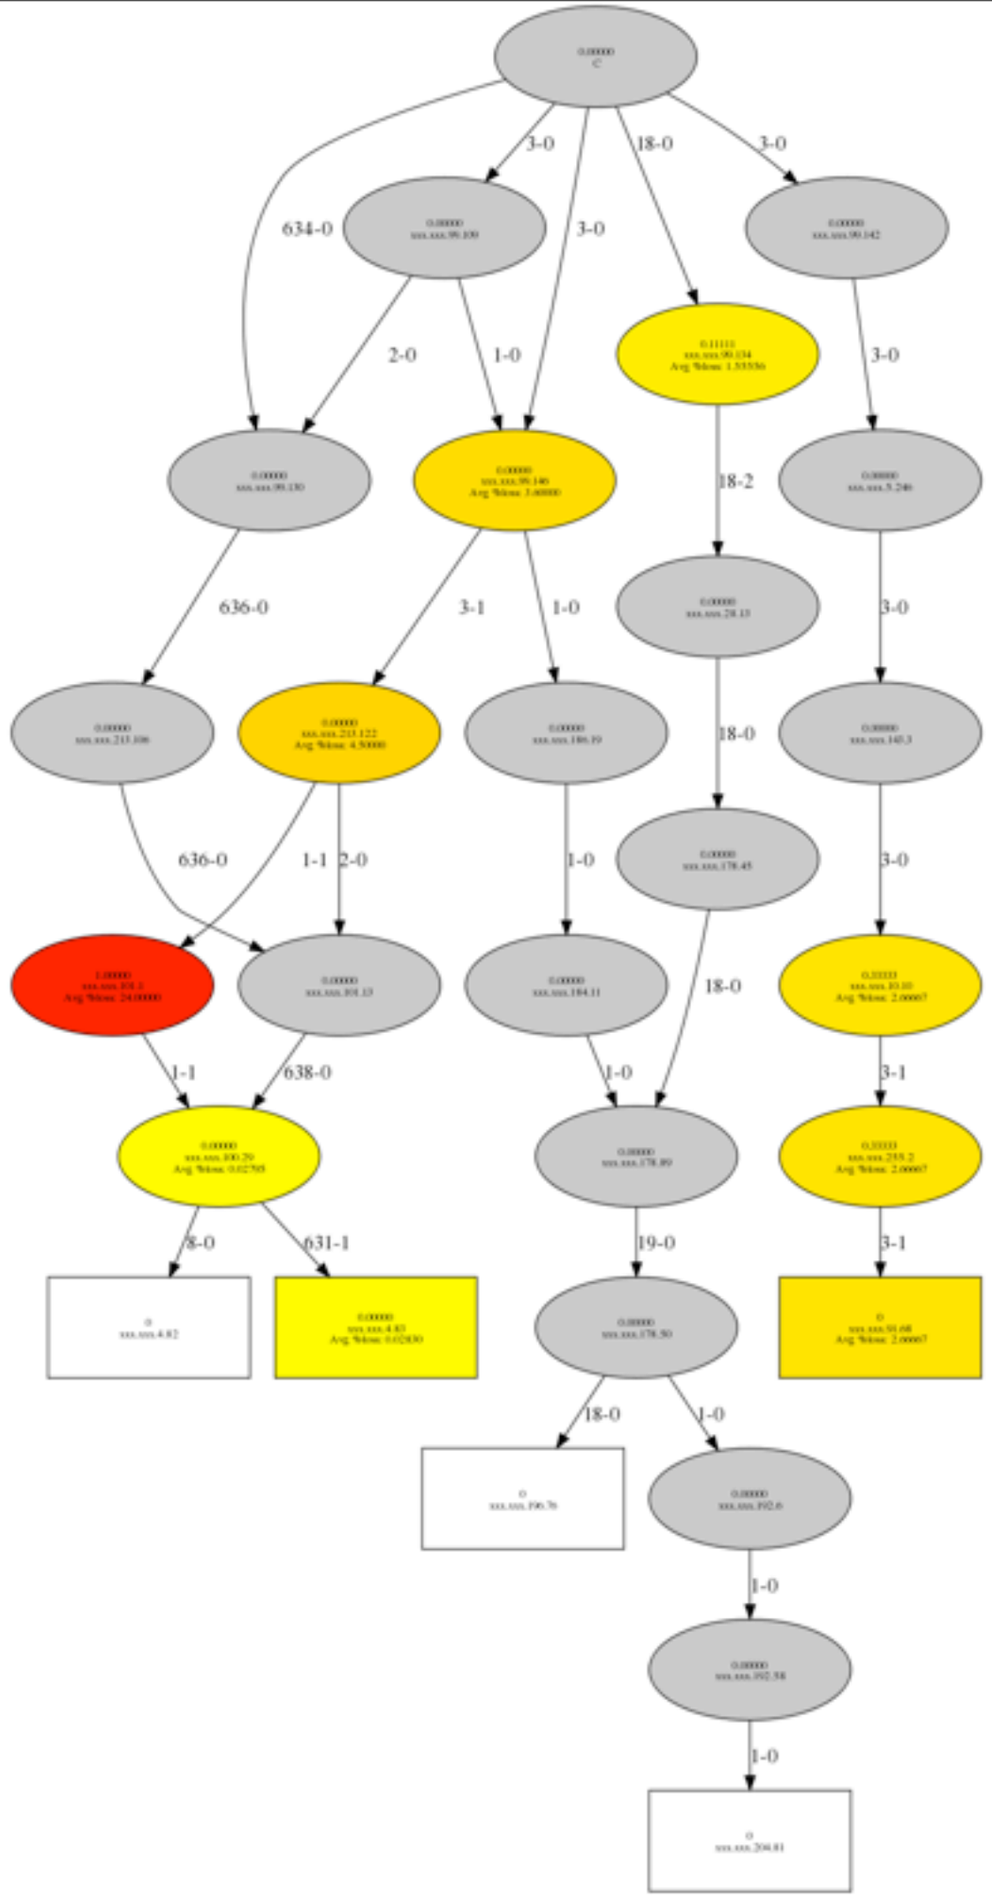

A faulty node

- r_2 is defined as a faulty node if
  - probability of loss (l_i/t_i) is high and is uniformly distributed across all its branches

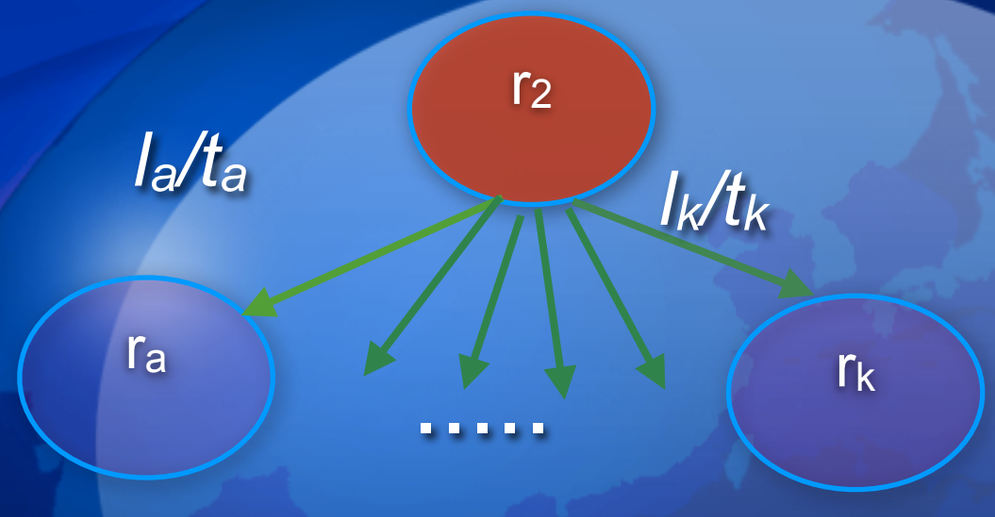

l_i = # of runs via r_2 to r_i seeing loss

t_i = total # of runs via r_2 to r_i

Network Graph analysis

- Developed cost functions to learn the probability of each node being faulty
- Supervised pattern classification algorithms are used to learn the accuracy of the cost functions

Example network graphs for a few end hosts in China

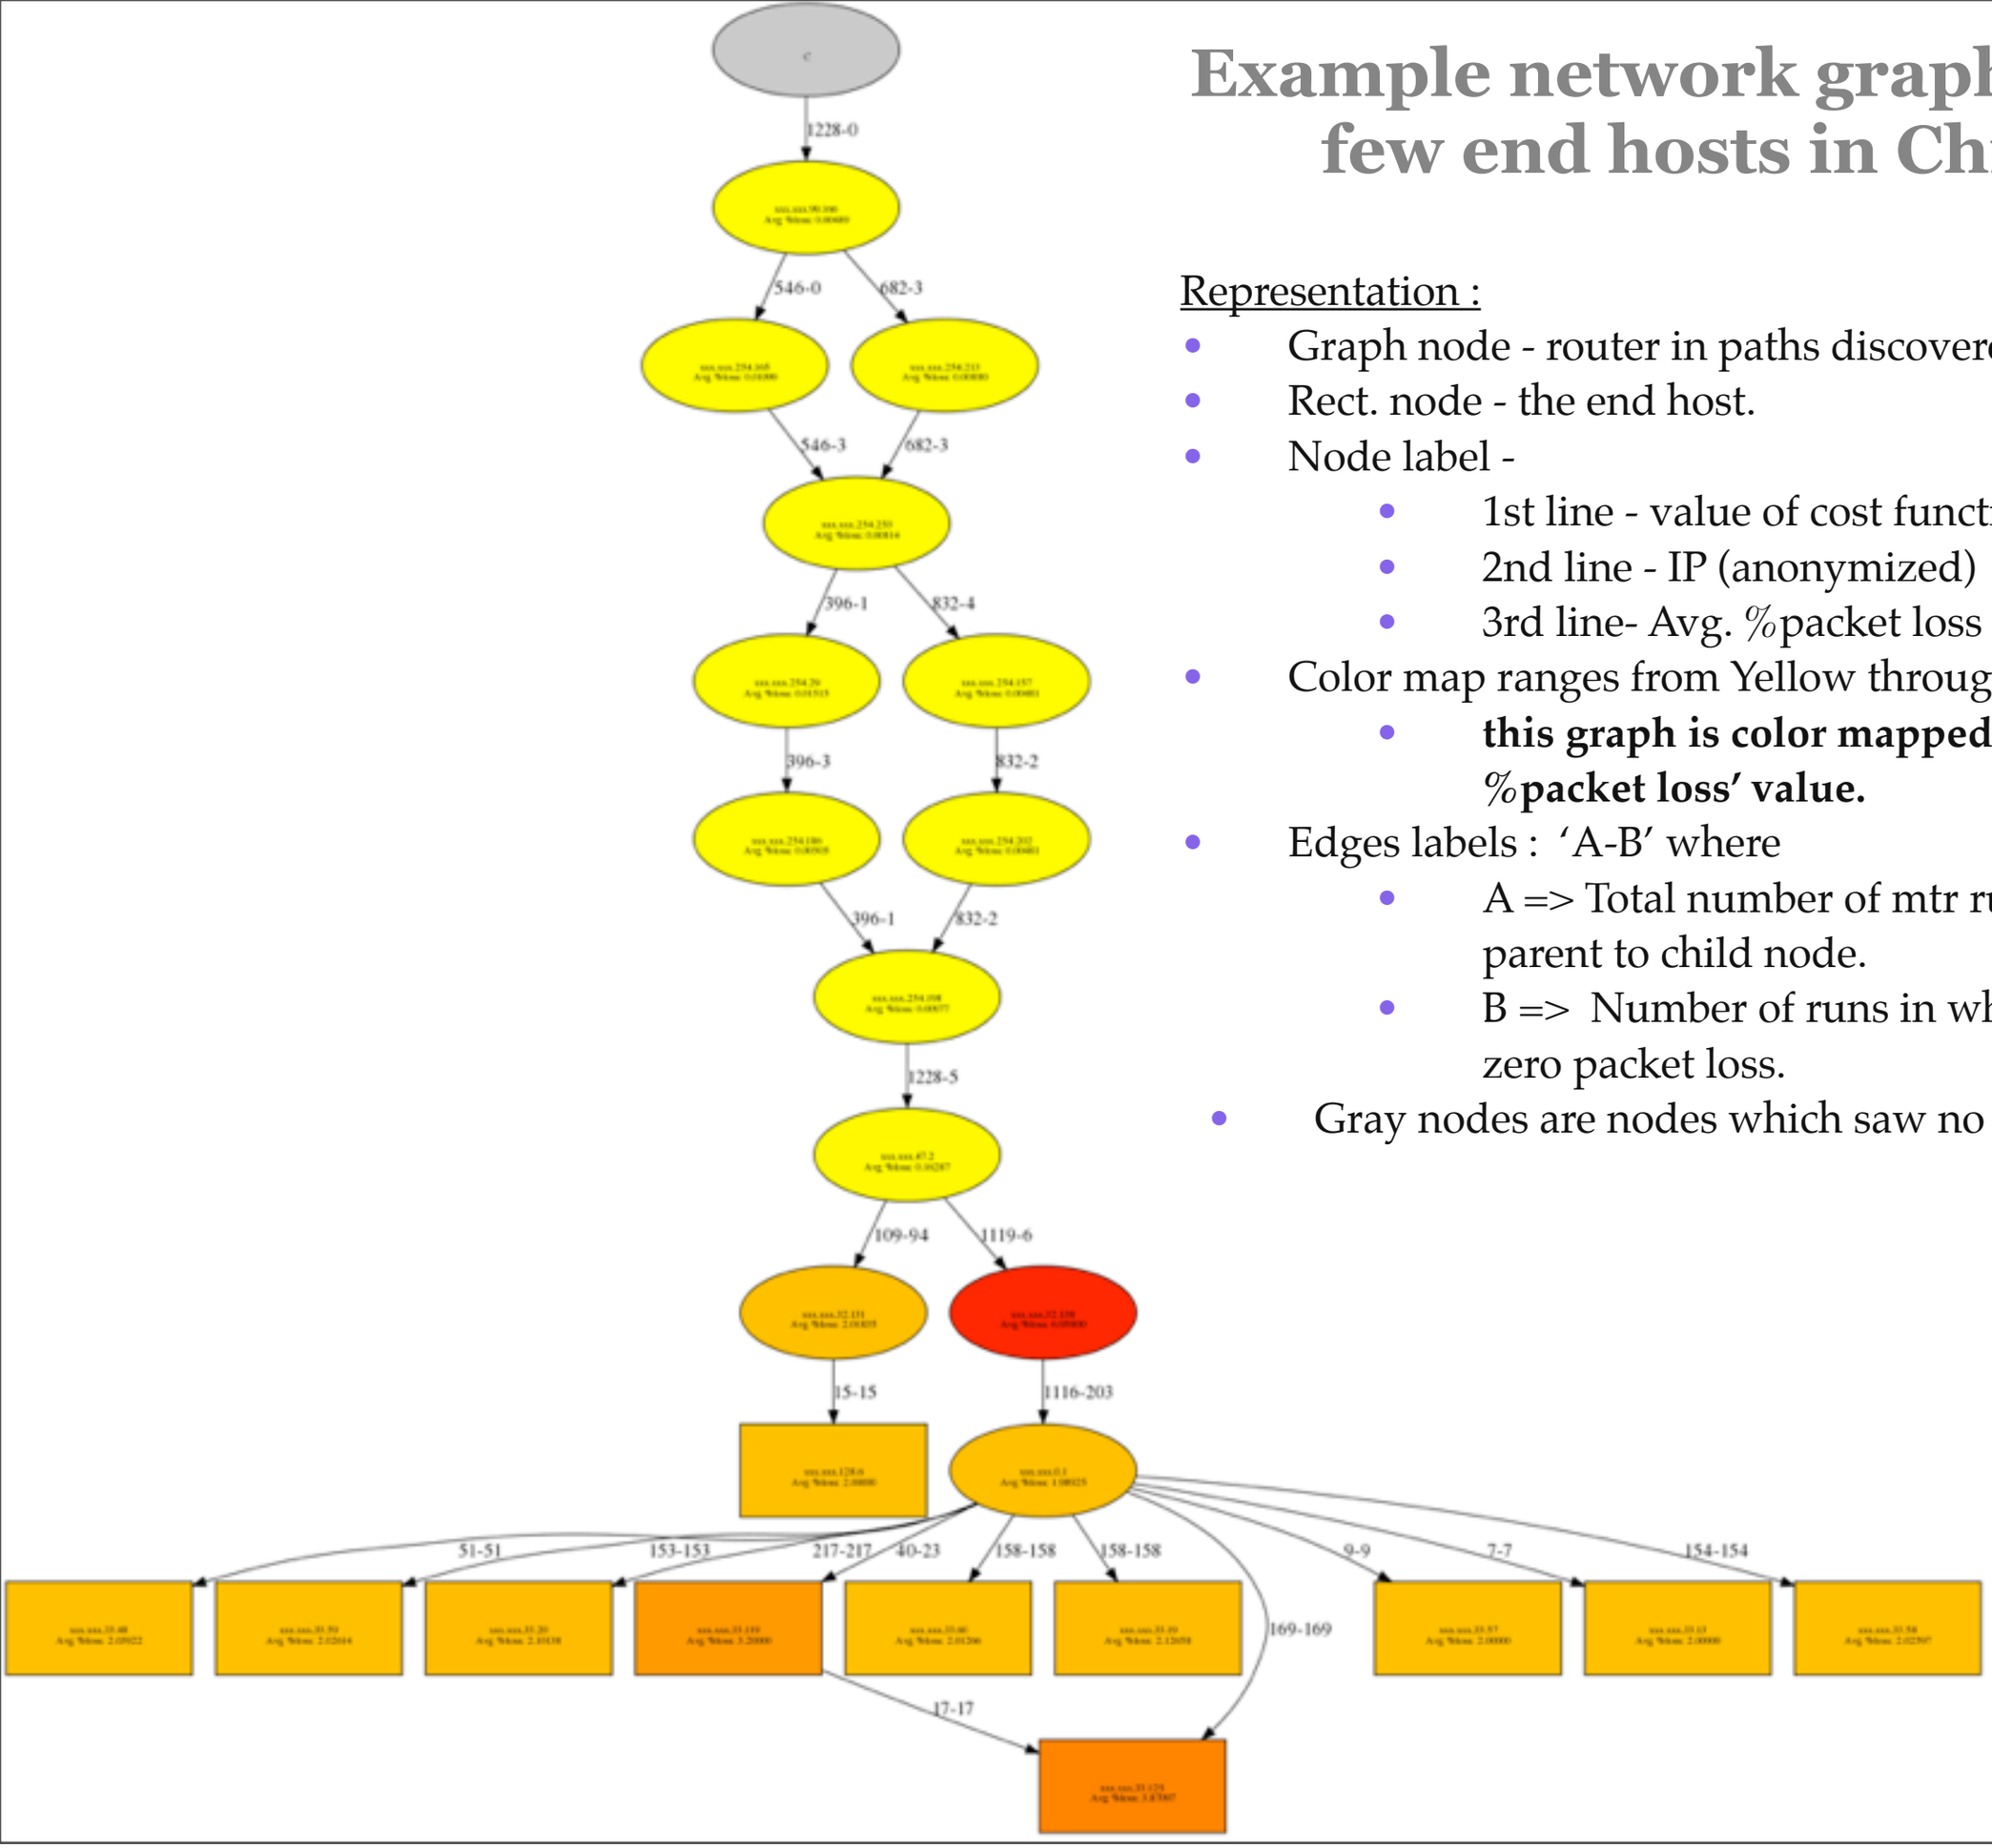

Representation :

- Graph node - router in paths discovered by MTR.
- Rect. node - the end host.
- Node label -
  - 1st line - value of cost function
  - 2nd line - IP (anonymized)
  - 3rd line- Avg. %packet loss at the node.
- Color map ranges from Yellow through orange to red.
  - **this graph is color mapped based on the 'Avg. %packet loss' value.**
- Edges labels : 'A-B' where
  - A => Total number of mtr runs through the parent to child node.
  - B => Number of runs in which there was non-zero packet loss.
- Gray nodes are nodes which saw no packet loss.

Network-monitoring data collection

Packeteer box at Chicago

- Passively monitors traffic to/from GLORIAD router in Chicago
- Exports extended Netflow records
  - Bytes retransmitted
  - Application classification
- **Replacing** Packeteer with open source monitoring box
  - Commercial box
  - Limited to 1G line speed

Nprobe Monitoring box

GOALS

- Network utilization and performance measurement box - running at least at 10G line speed
- Emit extended netflow records including retransmissions, application classification and more

HARDWARE

- Dell PowerEdge R410 Server - 8 core intel processor
- 10GE Intel Fiber Card

Nprobe software

- Nprobe is open source software developed by *Luca Deri* (http://www.ntop.org/nProbe.html)
- Development effort is in progress with help of Luca Deri and CSTNet (GLORIAD-China partners)
  - Current version exports retransmissions data
  - Next steps: Better application classification

Integrating data from other tools

GLORIAD Perfsonar nodes

- Currently deployed at Seattle, Chicago and Singapore
- Soon nodes will be installed in Amsterdam and Hong Kong
- Looking for ways to integrate/visualize perfsonar data in DVNOC

Conclusion

- Common platform to share network operations, utilization, performance and security data
- Addresses “disconnect” between all the different levels of network operators

Thank you.