GSFC FY04 IRAD "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Summary Accomplishments and Future Plans of GSFC's 10-Gbps Lambda Network Project

J. Patrick Gary
Network Projects Leader/606.1
NASA Goddard Space Flight Center
August 8, 2005

Information Complementing Demonstration for NASA AA Al Diaz

GSFC FY04 IRAD "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Project Goals

- **"…establish a "Lambda Network" (in this case using optical wavelength technology and 10 Gbps Ethernet per wavelength) from GSFC's Earth science Greenbelt facility in MD to the Scripps Institute of Oceanography (SIO) through the University of California, San Diego (UCSD) facility over the National Lambda Rail (NLR), a new national dark optical fiber infrastructure."**
- **"…make data residing on Goddard's high speed computer disks available to SIO with access speeds as if the data were on their own desktop servers or PC's."**
- **"…enable scientists at both institutions to share and use compute intensive community models, complex data base mining and multi-dimensional streaming visualization over this highly distributed, virtual working environment."**

Science-Enabling Network Technology Highlights of GSFC's 10-Gbps Lambda Network Project

Key Accomplishments (partial list)

- Partnered with NSF-funded DRAGON & OptIPuter Projects
- First site on DRAGON's regional optical network
- Enabling e-VLBI real-time data flows from GGAO to MIT/Haystack
- One of first 10 users on the National LambdaRail (NLR)
- Demoed real-time data flows at SC04 with OptIPuter Project
- Successfully using pre-COTS 80-km DWDM 10-Gbps Ethernet optics from Finisar
- ~9-Gbps dual stream UDP/IPv4 across DRAGON's ITU-compliant Channel 49 & NLR/WASH-STAR lambdas
- >5-Gbps "user payload" coast-to-coast single stream TCP/IPv4
- 3D HDTV multi-Gbps real-time data streaming from College Park to holographic display at GSFC
- First 10-Gbps network within GSFC
- Leading NASA's way in NLR use for ARC's Project Columbia
- W. Kuang et al., "High Speed Networking and Large-Scale Simulation in Geodynamics", abstract/poster, Fall AGU 2004
- S. Zhou et al., "High-Speed Network and Grid Computing for High-End Computation: Application in Geodynamics Ensemble Simulations", 13th Annual Mardi Gras Conference, February 2005

The DRAGON Testbed

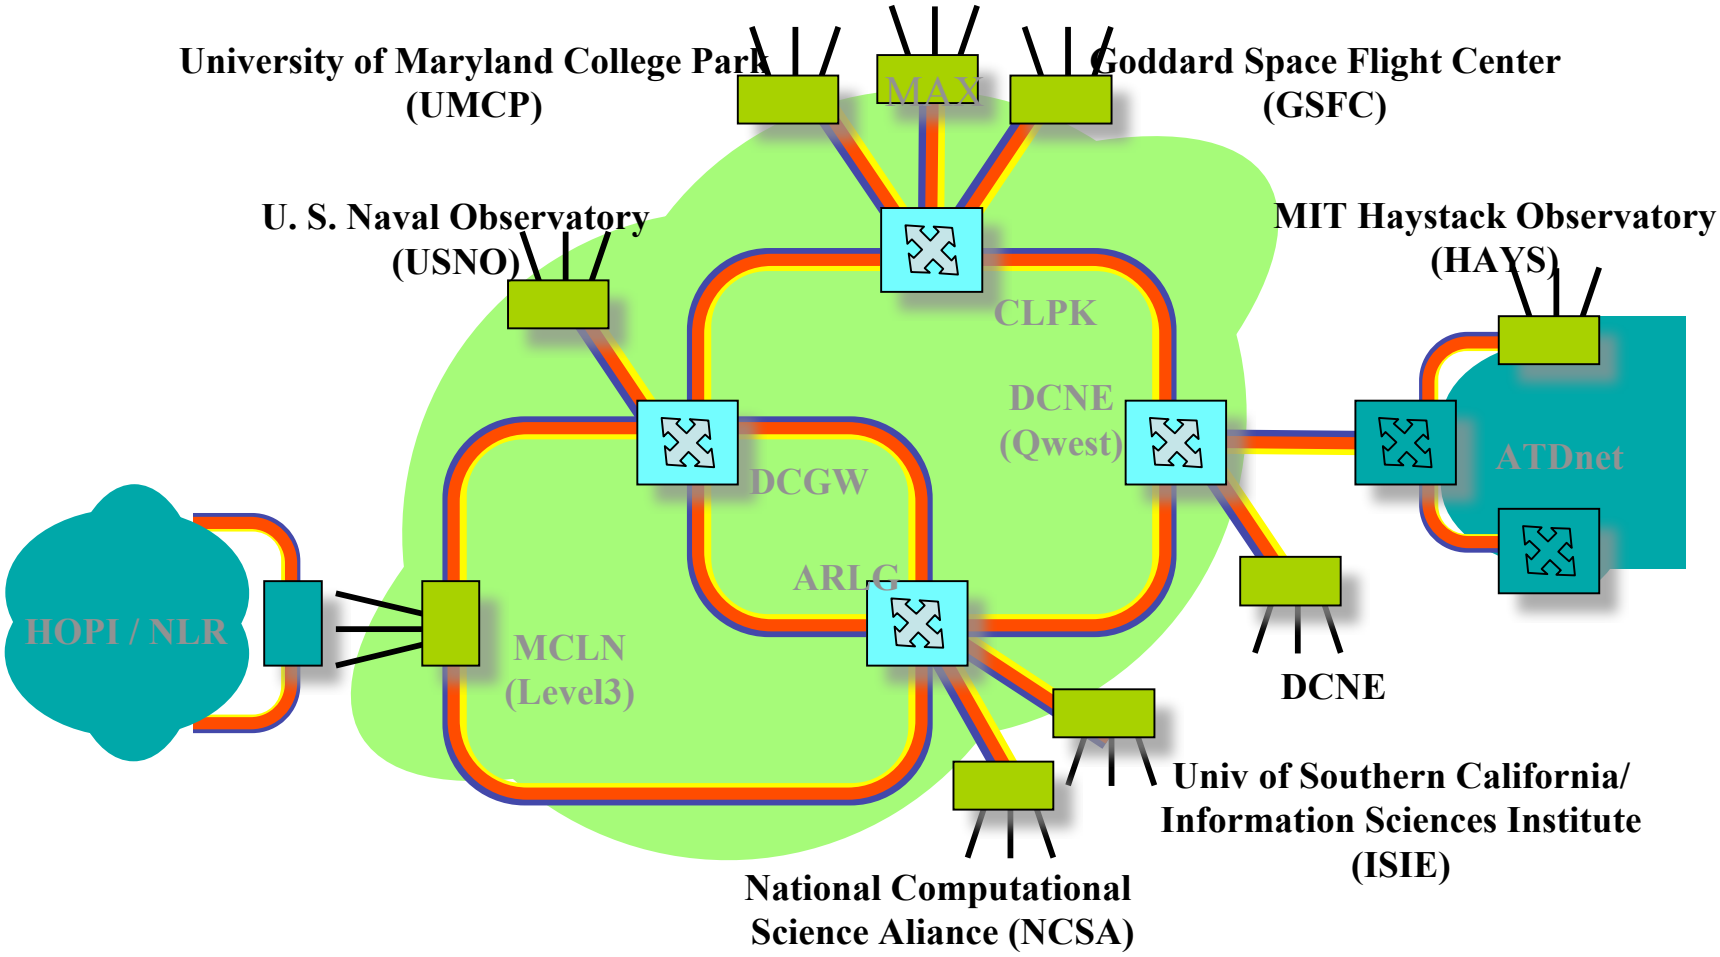

DRAGON eVLBI Experiment Configuration

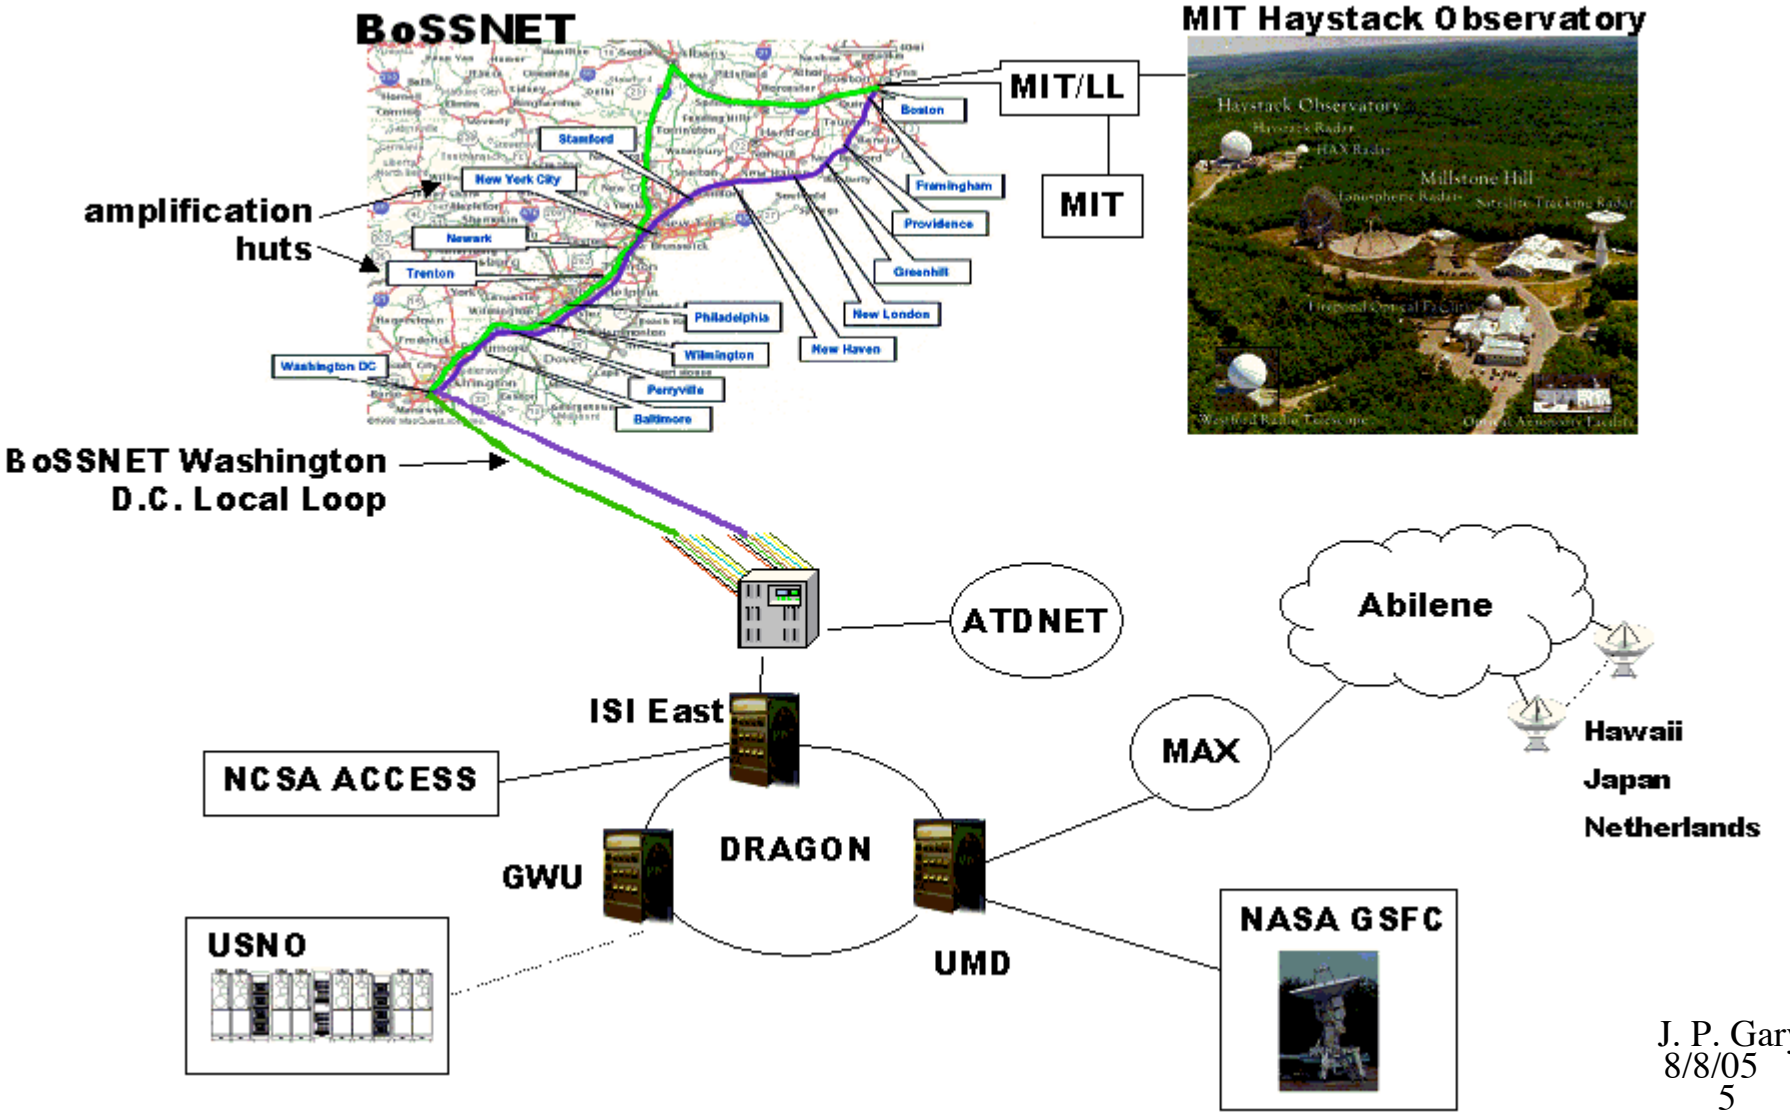

NASA GSFC in the NLR booth with the OptIPuter-provided 15-screen tiled display cluster during SC2004

NLR booth at SC2004 with OptIPuter-provided 15-screen tiled display cluster.

Eric Sokolowsky (GST, Inc.) of GSFC's SVS interactively views model and observation data (set 1) from NASA's Animated Earth project with hyperwall paradigm.

Eric Sokolowsky (GST, Inc.) of GSFC's SVS with model and observation data (set 2) from NASA's Animated Earth project in hyperwall paradigm.

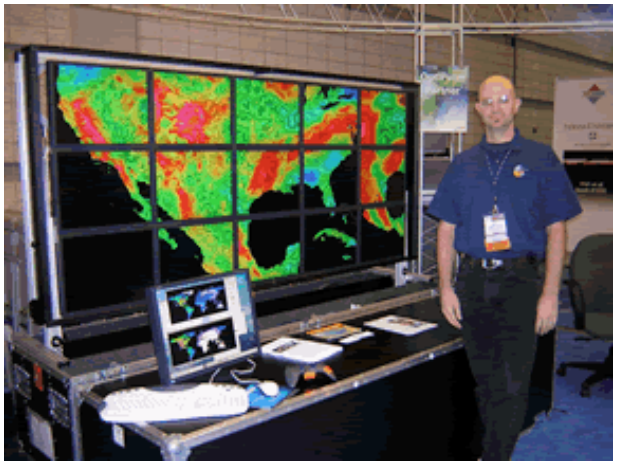

Randall Jones (GST, Inc.) of GSFC's SVS with model data from NASA's Land Information System in OptIPuter's display paradigm.

Various visitors to the NLR booth being briefed by Tom West, president and CEO of the NLR.

Rear view of the OptIPuter-provided 15-screen tiled display cluster.

L-Net SC2004 Photo Gallery: http://esdcd.gsfc.nasa.gov/LNetphoto3.html

Photo Sources: Randall Jones, NASA GSFC

NASA GSFC Tests with OptIPuter Across the National LambdaRail

August 2005

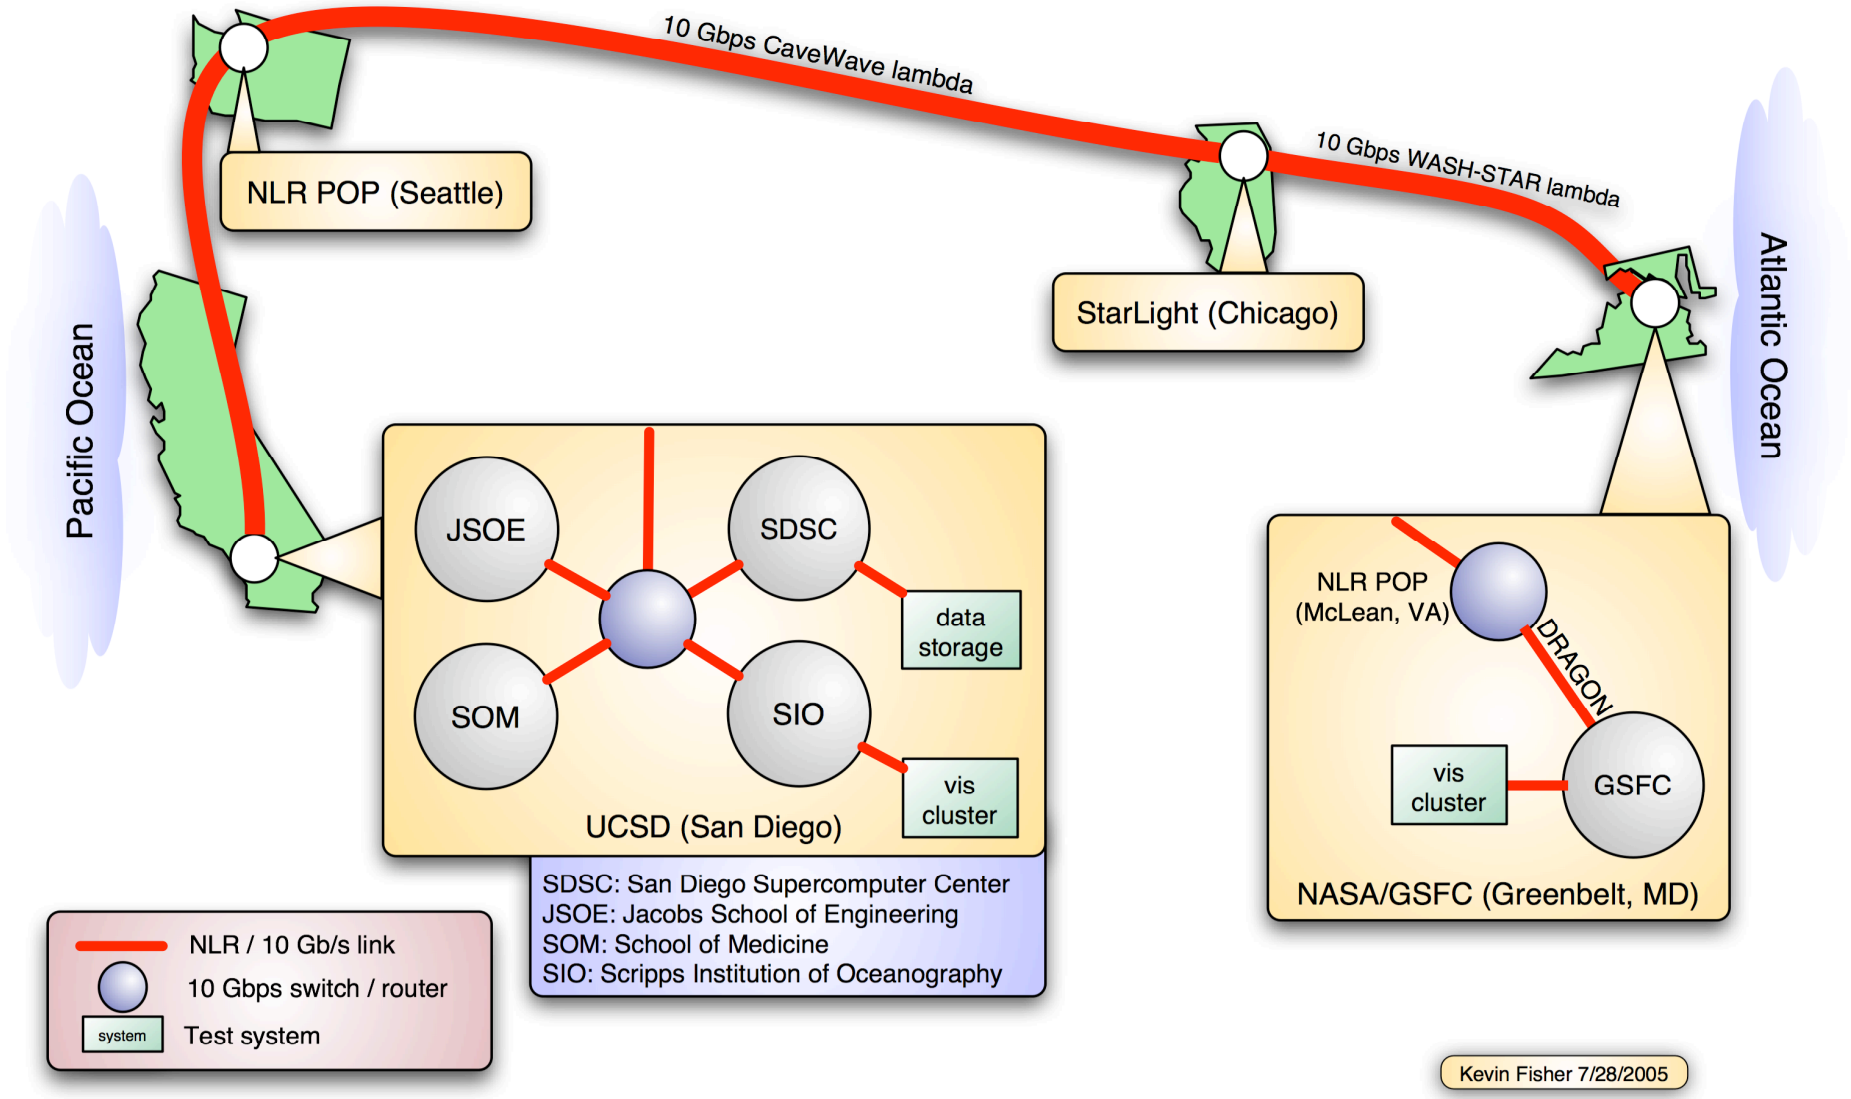

GSFC L-Net Configurations at McLean and Greenbelt

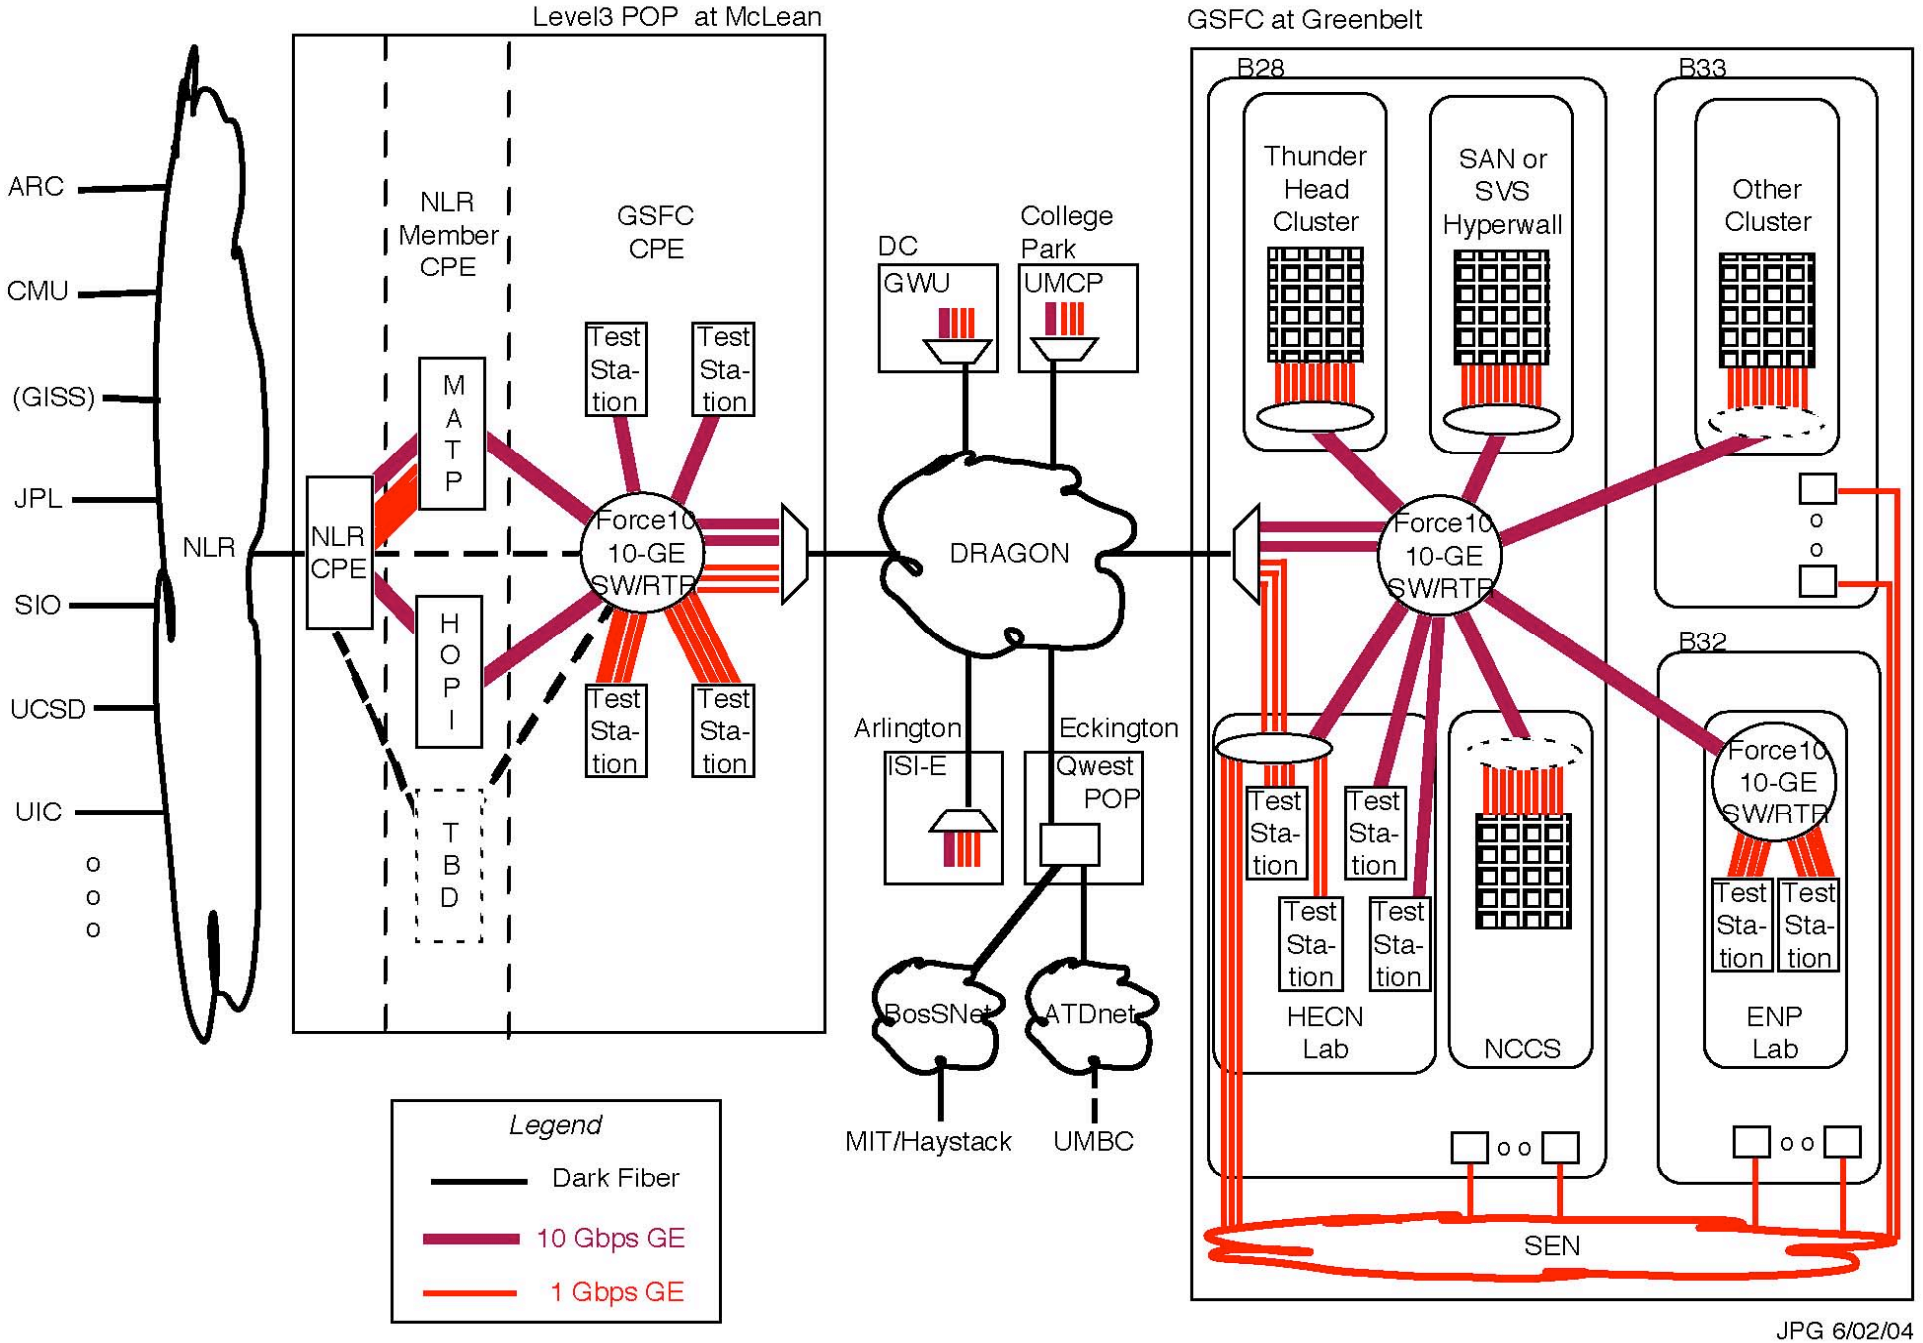

JPG 6/02/04

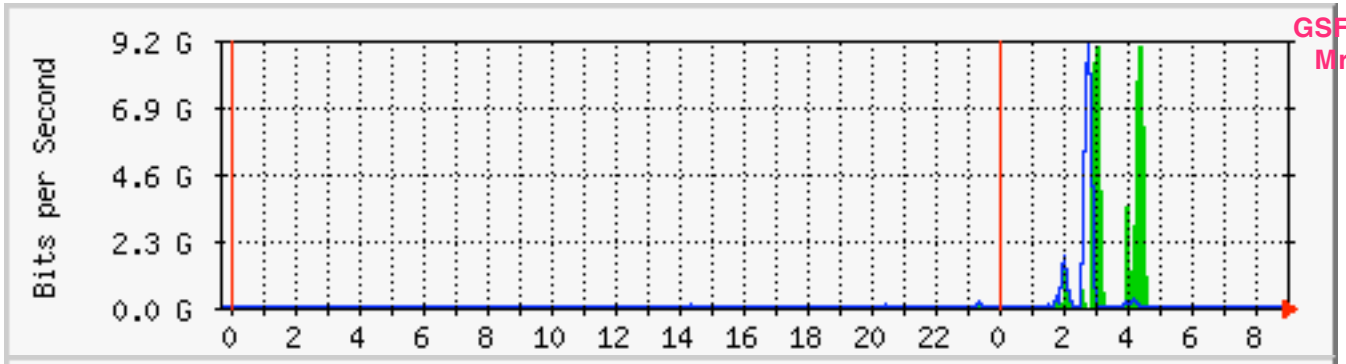

GSFC Scientific and Engineering Network (SEN)
Mrtg-based `Daily' Graph (5 Minute Average)
Bits per second In and Out
On Selected Interfaces

10 GigE from McLean to Chicago via OptIPuter Lambda
5 August 2005

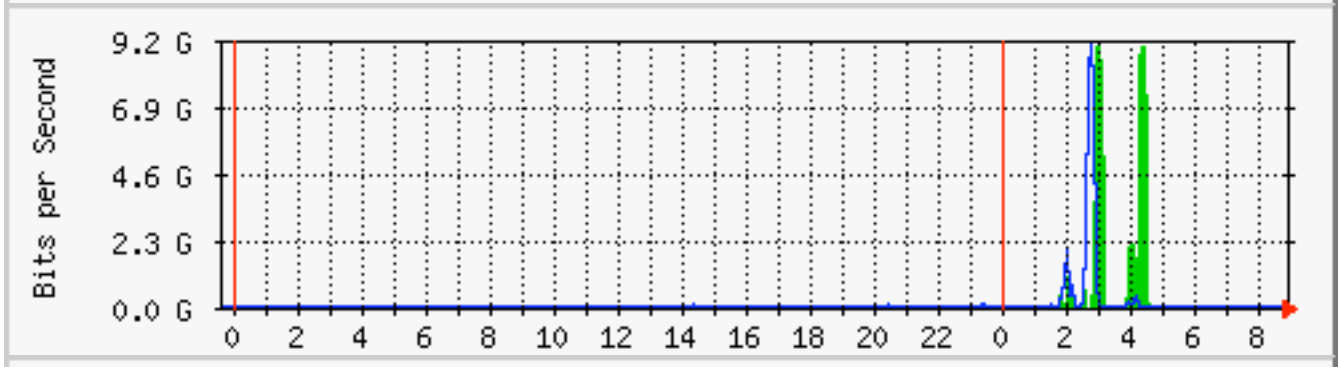

DRAGON 10Gig DWDM XFP
5 August 2005

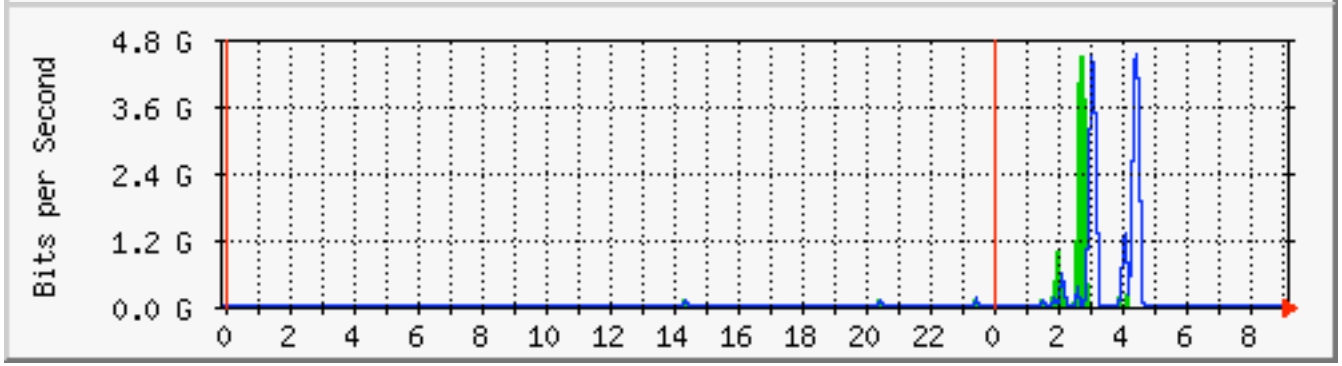

chance1 10Gig (eth1 Intel Pro/10GbE)
5 August 2005

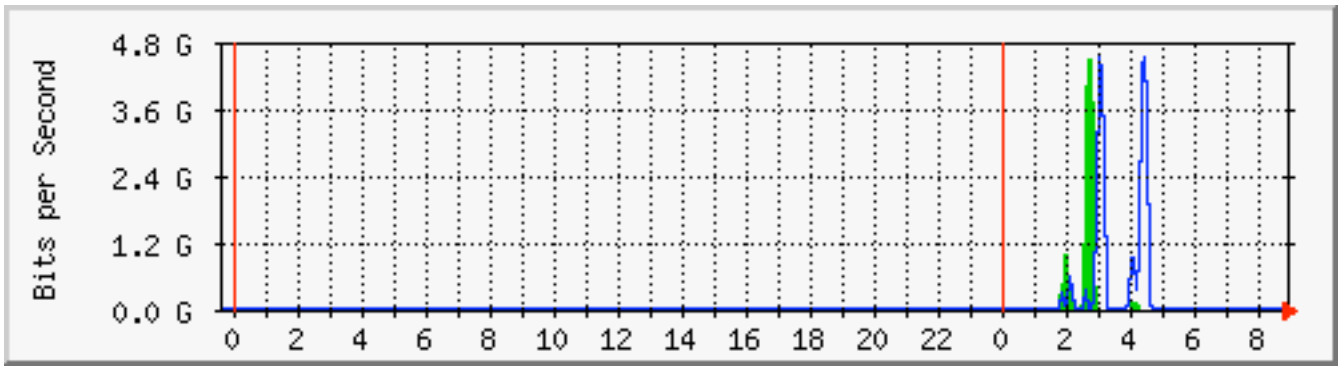

chance2 10Gig (eth1 Intel Pro/10GbE)
5 August 2005

Science-Enabling Network Technology Highlights of GSFC's 10-Gbps Lambda Network Project

Future Plans (partial list)

- Support NASA NRA Proposals
  - "MAP Core Integration LambdaGrid Infrastructure" by Smarr (UCSD) et al to NASA's MAP NRA
  - "Brokering and Chaining Distributed Services and Data Using OptIPuter and the National Lambda Rail" by Ramapriyan (GSFC) et al to NASA's ROSES NRA
  - "Enabling NASA Applications Across Heterogeneous High Performance Networks" by Habib (CUNY) et al to NASA NNH05ZDA001N-Applied Information Systems Research (a.k.a. ROSES:D3)
- Leverage existing 10-Gbps connection with Internet2's NLR/HOPI lambda
- 3D HDTV multi-Gbps real-time data streaming from GSFC to holographic display at UCSD for iGrid 2005
- Support new 10-Gbps connection with NREN's NLR lambda once available
- Extend GSFC's existing10 Gbps L-Net to additional GSFC buildings, computers, and users
- Increase the number and type of GSFC science/exploration research projects that benefit from the increased throughput performance that multi-wavelength optical networking can provide

Next Step: OptIPuter, NLR, and Starlight Enabling Coordinated Earth Observing Program (CEOP)

Source: Milt Halem, NASA GSFC

Accessing 300TB's of Observational Data in Tokyo and 100TB's of Model Assimilation Data in MPI in Hamburg -- Analyzing Remote Data Using GRaD-DODS at These Sites Using OptIPuter Technology Over the NLR and Starlight

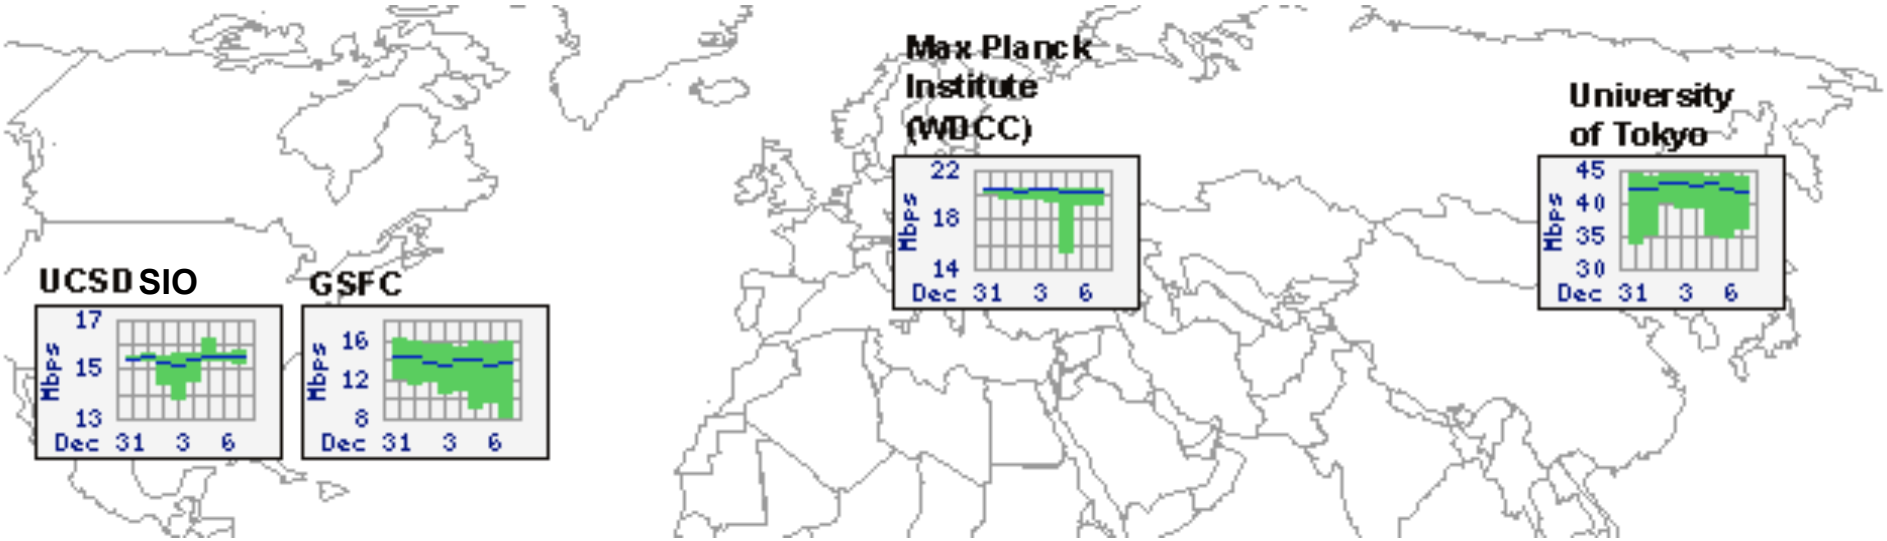

Note Current Throughput 15-45 Mbps: OptIPuter 2005 Goal is ~10 Gbps!

OptIPuter

http://ensight.eos.nasa.gov/Organizations/ceop/index.shtm

it2

OptIPuter and NLR will Enable Daily Land Information System Assimilations

- The Challenge:
  - More Than Dozen Parameters at ~ 50 GB per Parameter , Produced Six Times A Day, Need to be Analyzed
- The LambdaGrid Solution:
  - Sending this Amount of Data to NASA Goddard from Project Columbia at NASA Ames for Human Analysis Would Require < 15 Minutes/Day Over NLR
- The Science Result:
  - Making Feasible Running This Land Assimilation System Remotely in Real Time

OptIPuter

Source: Milt Halem, NASA GSFC

it²

NLR/GSFC Applications: Hurricane Prediction

- The NASA Finite-Volume General Circulation Model (fvGCM) has been producing real-time, high-resolution (~25 km) weather forecasts focused on improving hurricane track and intensity forecasts.
- During the active 2004 Atlantic hurricane season, the fvGCM provided landfall forecasts with an accuracy of ~100 km up to 5 days in advance.
- The 50–100 Mbps throughput available between fvGCM users at GSFC and the Columbia supercomputer at ARC greatly hindered carrying out time-critical simulations of the hurricanes that devastated Florida.
- The 10 Gbps NLR access will enable remote, 3D visualization analysis as soon as forecast variables become available.
- Key Contacts: Ricky Rood, Bob Atlas, Horace Mitchell, GSFC; Chris Henze, ARC.

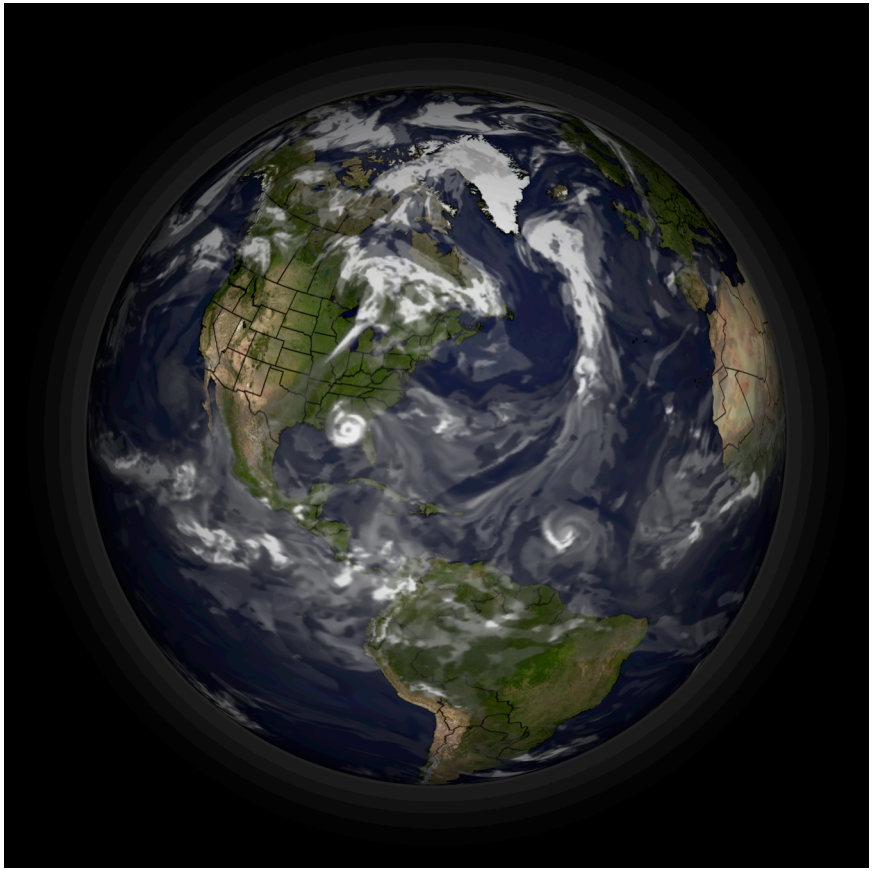

In an fvGCM forecast, Hurricane Frances makes landfall on the Gulf Coast of Florida while Hurricane Ivan intensifies in the tropical Atlantic. Visualization by J. Williams, GST.

http://fvnwp.gsfc.nasa.gov

NLR/GSFC Applications: Global Aerosols

- Project Atmospheric Brown Clouds (ABC) is an international effort to discover and analyze areas of brown colored atmosphere to learn how dust and pollution particles are transported and what impacts they have on the environment, climate, agricultural cycles, and quality of life.
- GSFC and the Scripps Institution of Oceanography (SIO) are planning a collaboration to predict the flow of aerosols from Asia across the Pacific to the U.S. on timescales of days to a week.
- GSFC will provide an aerosol chemical tracer model (GOCAR) embedded in a high-resolution regional model (MM5) that can assimilate data from Indo-Asian and Pacific ground stations, satellites, and aircraft.
- Remote computing and analysis tools running over the NLR will enable acquisition and assimilation of the Project ABC data.
- Key Contacts: Yoram Kaufman, William Lau, GSFC; V. Ramanathan, Chul Chung, SIO

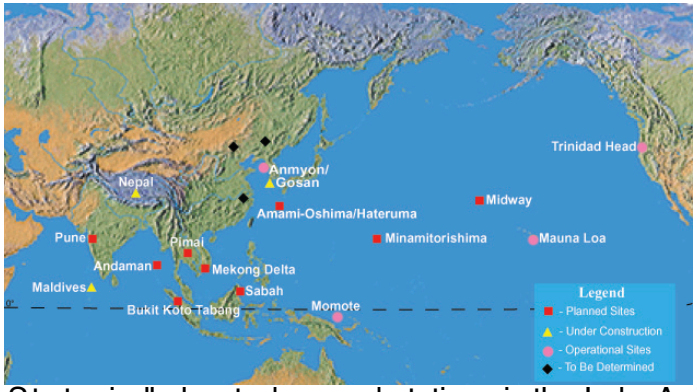

Strategically located ground stations in the Indo-Asian and Pacific regions monitor atmospheric pollution.

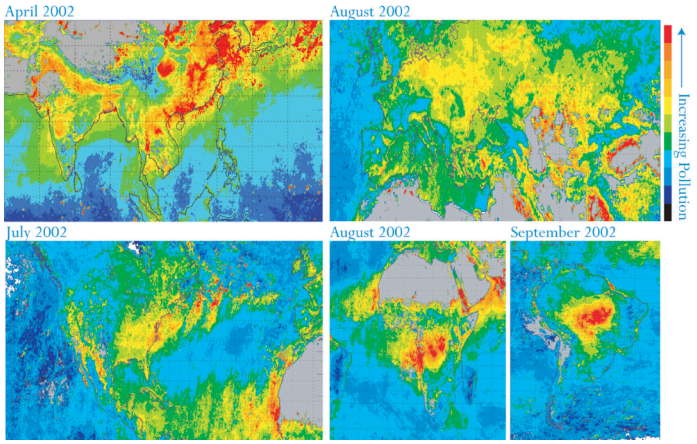

The global nature of brown clouds is apparent in analysis of NASA MODIS Data. Research by V. Ramanathan, C. Corrigan, and M. Ramana, SIO.

http://www-abc-asia.ucsd.edu

NLR/GSFC Applications: Remote Viewing and Manipulation of Large Earth Science Data Sets

- Remote viewing and manipulation of data sets at GSFC and JPL is needed to support EOSDIS and Earth system modeling.
- GSFC's EOSDIS Clearing House (ECHO) and JPL's GENESIS prototype science analysis system (iEarth) will become connected over the NLR. The link will enable comparison of hundreds of terabytes of data, generating large, multi-year climate records.
- Initial work will focus on the Estimating the Circulation and Climate of the Ocean (ECCO) modeling team. Besides ready access to the NLR, the team will need versatile subsetting and other data manipulation functions to reduce compute and bandwidth requirements as well as a set of Grid-accessible statistical analysis and modeling operators to refine and validate the ECCO models.
- Key Contacts: ECHO metadata gateway team, GSFC; GENESIS team, led by Tom Yunck, JPL

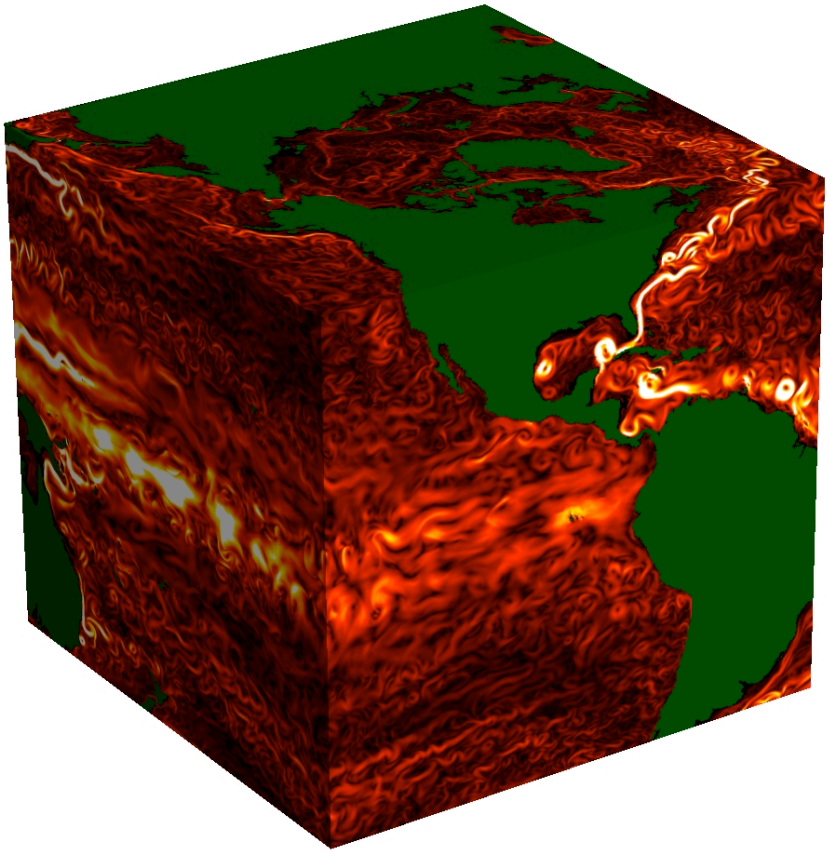

Near-surface (15-m) ocean current speed from an eddy-permitting integration of the cubed-sphere ECCO ocean circulation model. Research by JPL and MIT. Visualization by C. Henze, Ames.

http://www.ecco-group.org

NLR/GSFC Applications: Integration of Laser and Radar Topographic Data with Land Cover Data

- NASA has executed two advanced missions to create an accurate high-resolution topographic model of the Earth: the Shuttle Radar Topography Mission (SRTM) and ICESat, with its Geoscience Laser Altimeter System (GLAS).
- The agency now has the opportunity to merge the two data sets, using SRTM to achieve good coverage and GLAS to generate calibrated profiles. Proper interpretation requires extracting land cover information from Landsat, MODIS, ASTER, and other data archived in multiple DAACs.
- Use of the NLR and local data mining and subsetting tools will permit systematic fusion of global data sets, which are not possible with current bandwidth.
- Key Contacts: Bernard Minster, SIO; Tom Yunck, JPL; Dave Harding, Claudia Carabajal, GSFC.

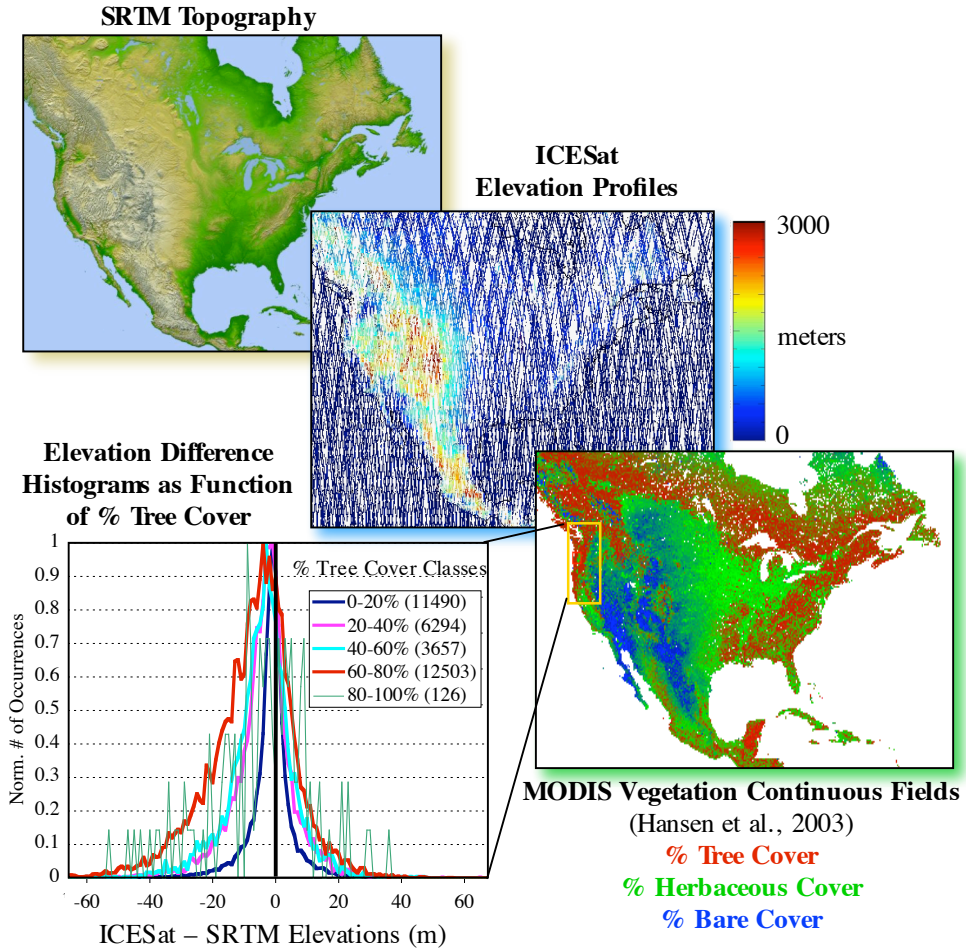

http://icesat.gsfc.nasa.gov
http://www2.jpl.nasa.gov/srtm
http://glcf.umiacs.umd.edu/data/modis/vcf

High speed networking and Grid computing for large-scale simulation in geodynamics

W. Kuang[1], W. Jiang[2], S. Zhou[3], P. Gary[1], M. Seablom[1], W. Truszkowski[1], J. Odubiyi[4], D. Liu[2], J. Palencia[5], G. Gardner[6]

[1]NASA Goddard Space Flight Center, [2]JCET, UMBC, [3] Northrop Grumman IT/TASC, [4]Bowie State University, [5] Raytheon ITSS,[6] INDUSCORP

Introduction

Now large-scale simulation has been wide-spread in many disciplines of solid Earth science research. A typical numerical test in the simulation can easily reach 10^{13} flops and beyond.

One such research problem that we are working on now is to establish a framework on predicting geomagnetic secular variation on decadal and longer time scales, utilizing surface geomagnetic/paleomagnetic records and our MoSST core dynamics model (Figure 1). In this approach, model forecast results and observations are weighted to provide initial state for assimilation (Figure 2). Typically 30 independent numerical tests are necessary for a reasonable ensemble size. This could easily require a computing cycle on orders of petaflops and larger.

A single super-computing facility for such studies is not an optimal choice, due to many limitations, in particular those on user management and administration. But it is relatively easy for users (researchers) to manage because of a unified system environment.

Grid computing can be a much better choice so that independent numerical tests can be carried out independently on different systems. However, researchers (users) have to deal with heterogeneous systems and other problems, such as those on network communication.

In this poster, we discuss our activities in GSFC on application of grid computation to geodynamics modeling.

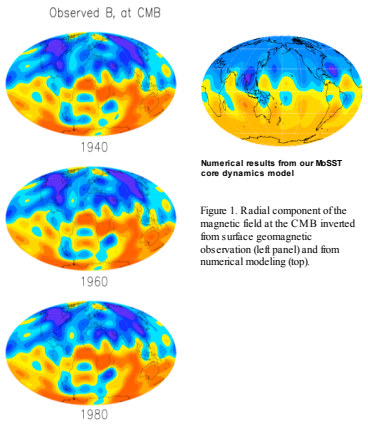

Figure 1. Radial component of the magnetic field at the CMB inverted from surface geomagnetic observation (left panel) and from numerical modeling (top).

Geomagnetic data assimilation

$$\mathbf{x}^a = \mathbf{x}^f + \mathbf{K}\left(\mathbf{x}^o - \mathbf{H}\mathbf{x}^f\right)$$

$\mathbf{X}^a$: Assimilation solution
$\mathbf{X}^f$: Forecast solution
$\mathbf{X}^o$: Observation data

Figure 2. Mathematical foundation of data assimilation. The common gain **K** depends on knowledge of error statistics of observations and of models. If ensemble Kalman-filter approach is applied. An ensemble size of at least 30 (i.e. independent tests) is required.

Prototype on MoSST simulation with independent systems

The objective of this prototype work is to test operability of executing our MoSST core dynamics model on independent computing systems. Individual computing units are slated out from selected components of our beowulf system to mimic independent computing environment. The prototype program for grid computing is built upon xcat3 framework (based on java/python). See Figure 3 for conceptual layout of our prototype experiment.

The sample script and the execution process are shown in Figure 4.

Our prototype experiment is very successful. With this experiment, we can proceed further our test on real remote systems. Also with this experiment, we can identify the needs from the user's considerations on supporting environment and other middleware that makes grid computing "friendly".

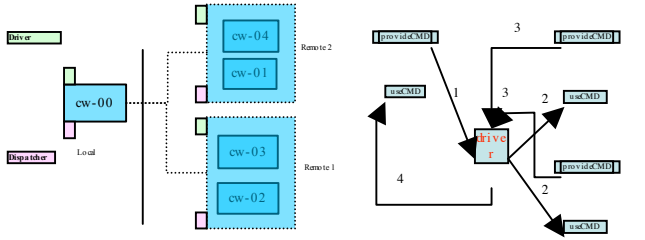

System configuration

OS: Fedora core 2; MPICH-1.2.5.2; Intel Fortran Compiler; Java 2
PE: Dual Intel Xeon, 2.4 Ghz, 1 GB, 1 GigEthernet

Figure 3. Prototype layout

```
import sys
import cca

from jarray import zeros

from java.lang import System
from java.lang import String, Object
# get the absolute location for XCAT

# create the TypeMap for the user component
# create component wrappers
geo1Provides = cca.createComponentWrapper("geo1Provider",
geo1ProviderMap)
geo2Provides = cca.createComponentWrapper("geo2Provider",
geo2ProviderMap)
#uses = cca.createComponentWrapper("user", userMap)
# assign a machine name
cca.setMachineName(uses, "cw-00")
cca.setMachineName(provides, "cw-00")
cca.setMachineName(geo1Provides, "cw-03")
cca.setMachineName(geo2Provides, "cw-04")

# set a creation mechanism to in-process
cca.setCreationMechanism(uses, "local")
cca.setCreationMechanism(provides, "local")
cca.setCreationMechanism(geo1Provides, "ssh")
cca.setCreationMechanism(geo2Provides, "ssh")

# connect their ports
cca.connectPorts(uses, "dispatchUsesPort", provides,
"dispatchProvidesPort")
cca.connectPorts(uses, "geo1UsesPort", geo1Provides,
"geo1ProvidesPort")
cca.connectPorts(uses, "geo2UsesPort", geo2Provides,
"geo2ProvidesPort")
# invoke the method
cca.invokeMethodOnComponent(uses,

                            portClassName,
                            providesPortName,
                            methodName,
                            methodParams)
```

Figure 4. Prototype Operation Script (left) and Screen Caption (right)

Discussions

1. Our research on geomagnetic data assimilation can greatly benefit from grid computing.
2. Our prototype experiment is successful and can be readily expanded to systems with identical settings and SSH communication protocol.
3. Our prototype experiment is limited in many areas, such as handling network communication between independent systems (e.g. instant feedback of remote systems to host systems), heterogeneous environment (e.g. prior knowledge on participating systems is necessary), authentication (e.g. prototype cannot handle high level access security requirement). Therefore, further experiment is needed to improve our work, such as integrating our work with other (developed and developing) middleware handling the problems.

Related work at GSFC

There are parallel, but related research going on in GSFC on networking and software development. These research activities are updated in http://esdcd.gsfc.nasa.gov/LNetImplementhtml. Recent overview of GSFC research activities is given by Dr. M. Halem and can be found in http://esdcd.gsfc.nasa.gov/L-Netpdfs/ESSAAC_MHpres9904.pdf. Some of the activities listed in the report are shown in Figures, 5 and 6. These activities work towards establishing 21st century cyber infrastructure for large-scale scientific teamwork based on fast network.

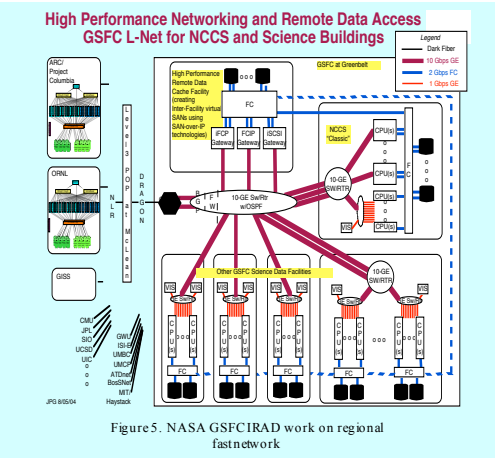

Figure 5. NASA GSFC/IRAD work on regional fast network

An Example of Application Requiring L-NET

- Drs. Paul Houser and Mike Bosilovich of Code 970 are collaborating with Dr. John Roads of SIO on the **Coordinated Earth Observing Program (CEOP)** accessing 300TB's of Observational Data in Tokyo and 100TB's of Model Assimilation Data in MPI in Hamburg Germany and analyzing remote data using GRaD-DODS at these sites over the NLR and Starlight.

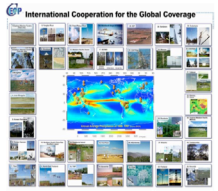

Figure 6.

APPLICATIONS -
Future GRID on 10-GE Network

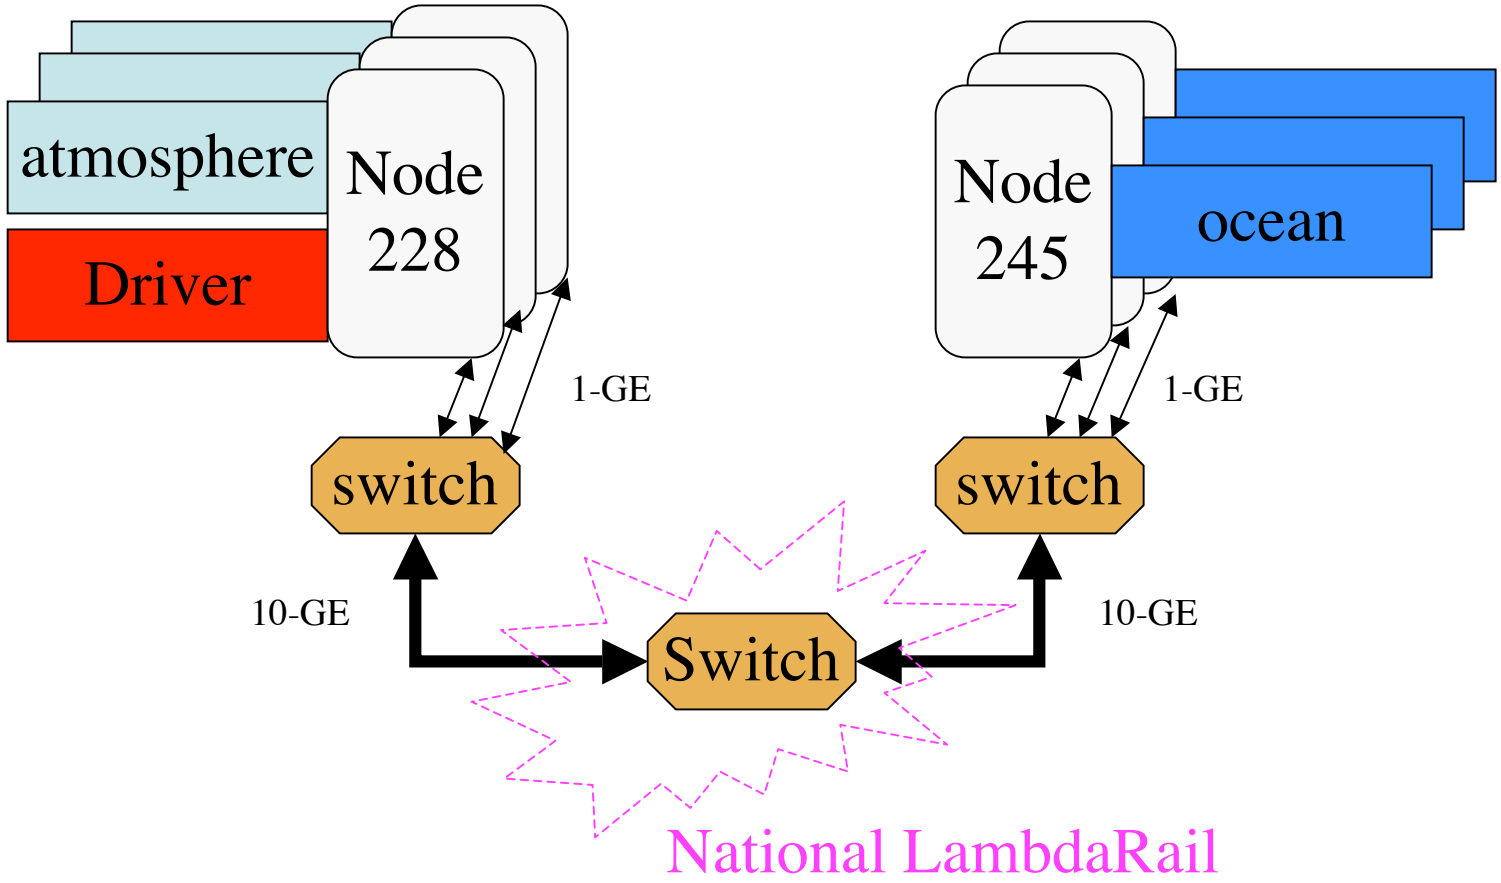

Dr. Zhou is working on applying Grid Computing and High-Speed Network to large-scale distributed computing in Earth and Space Science. More details can be found at http://esto.nasa.gov/conferences/estc2004/papers/a4p1.pdf.

"Brokering and Chaining Distributed Services and Data Using OptIPuter and the National Lambda Rail" by Ramapriyan (GSFC) et al to NASA's ROSES NRA

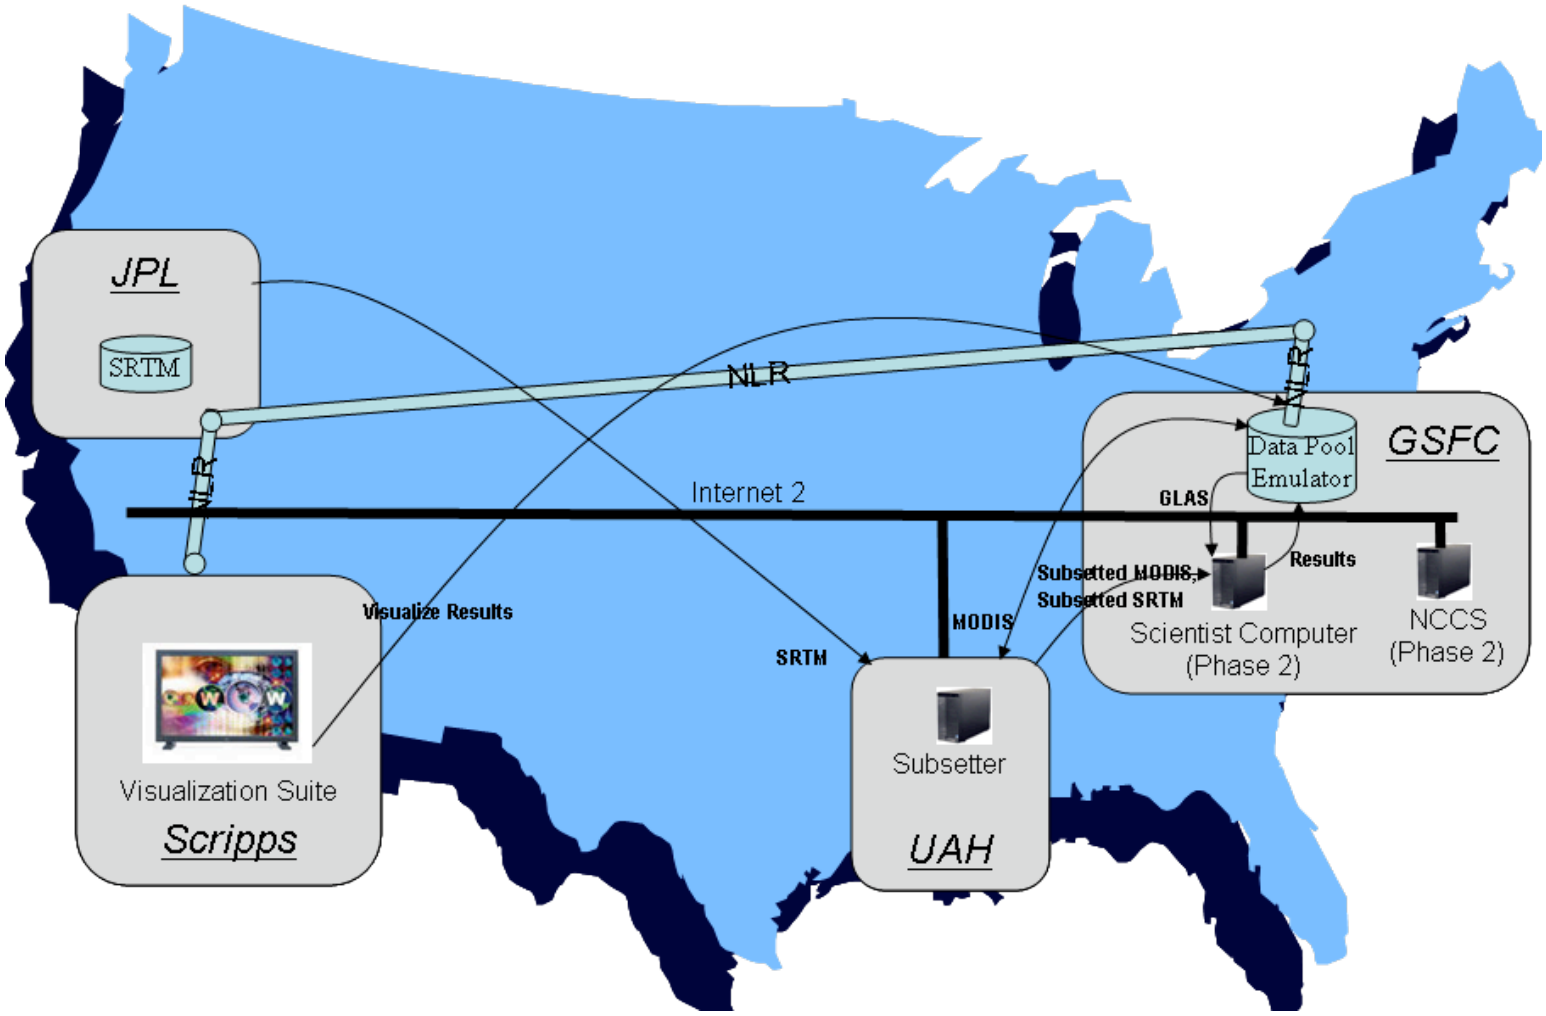

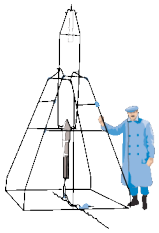

"Enabling NASA Applications Across Heterogeneous High Performance Networks" by Habib (CUNY) et al to NASA NNH05ZDA001N-Applied Information Systems Research (a.k.a. ROSES:D3)

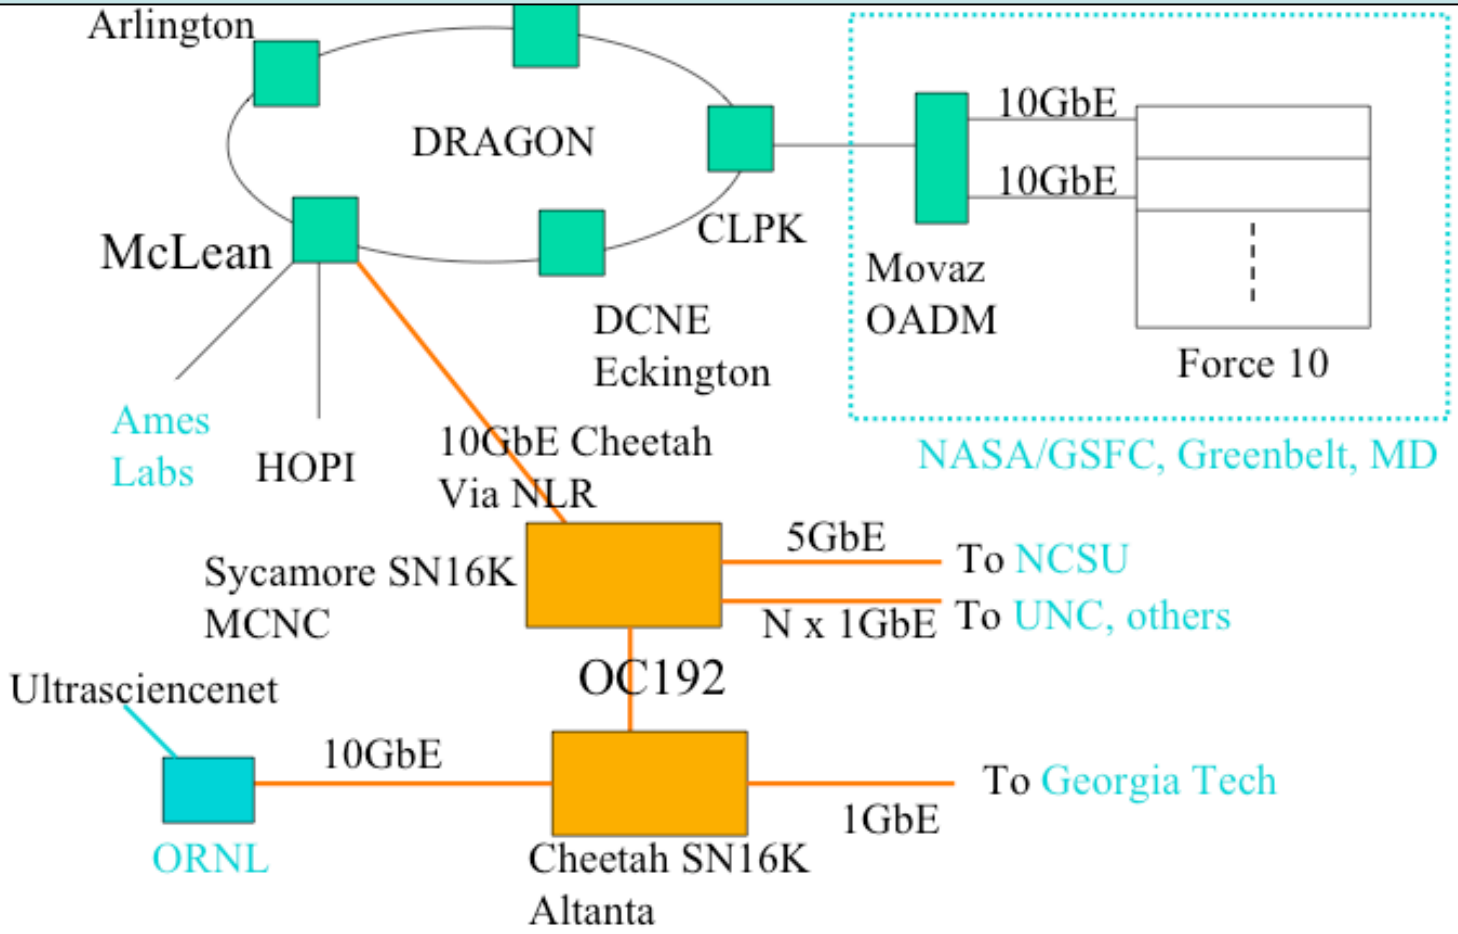

Fig. 1: Overall Proposed Network Connectivity to Cheetah

HOPI Topology

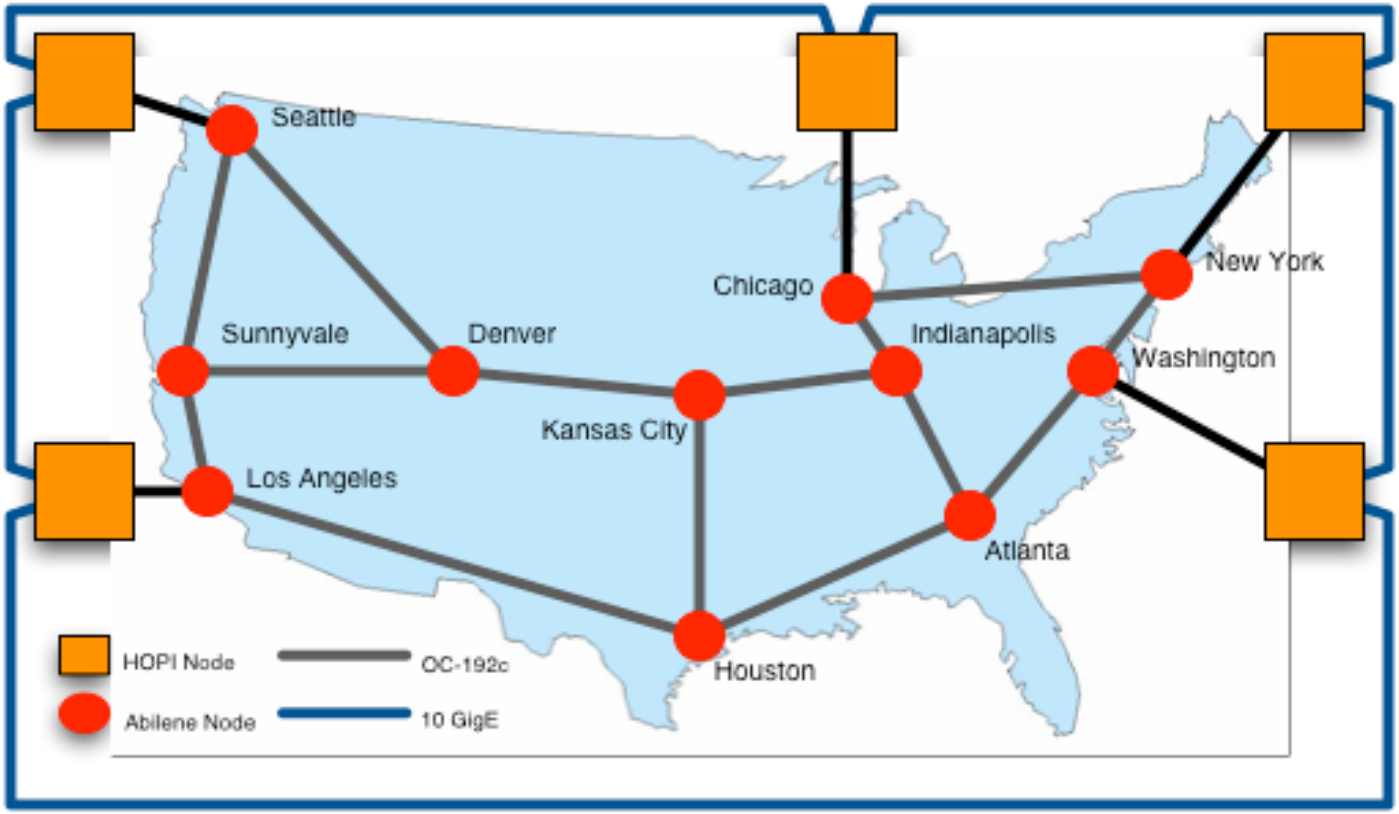

Future NREN Over NLR

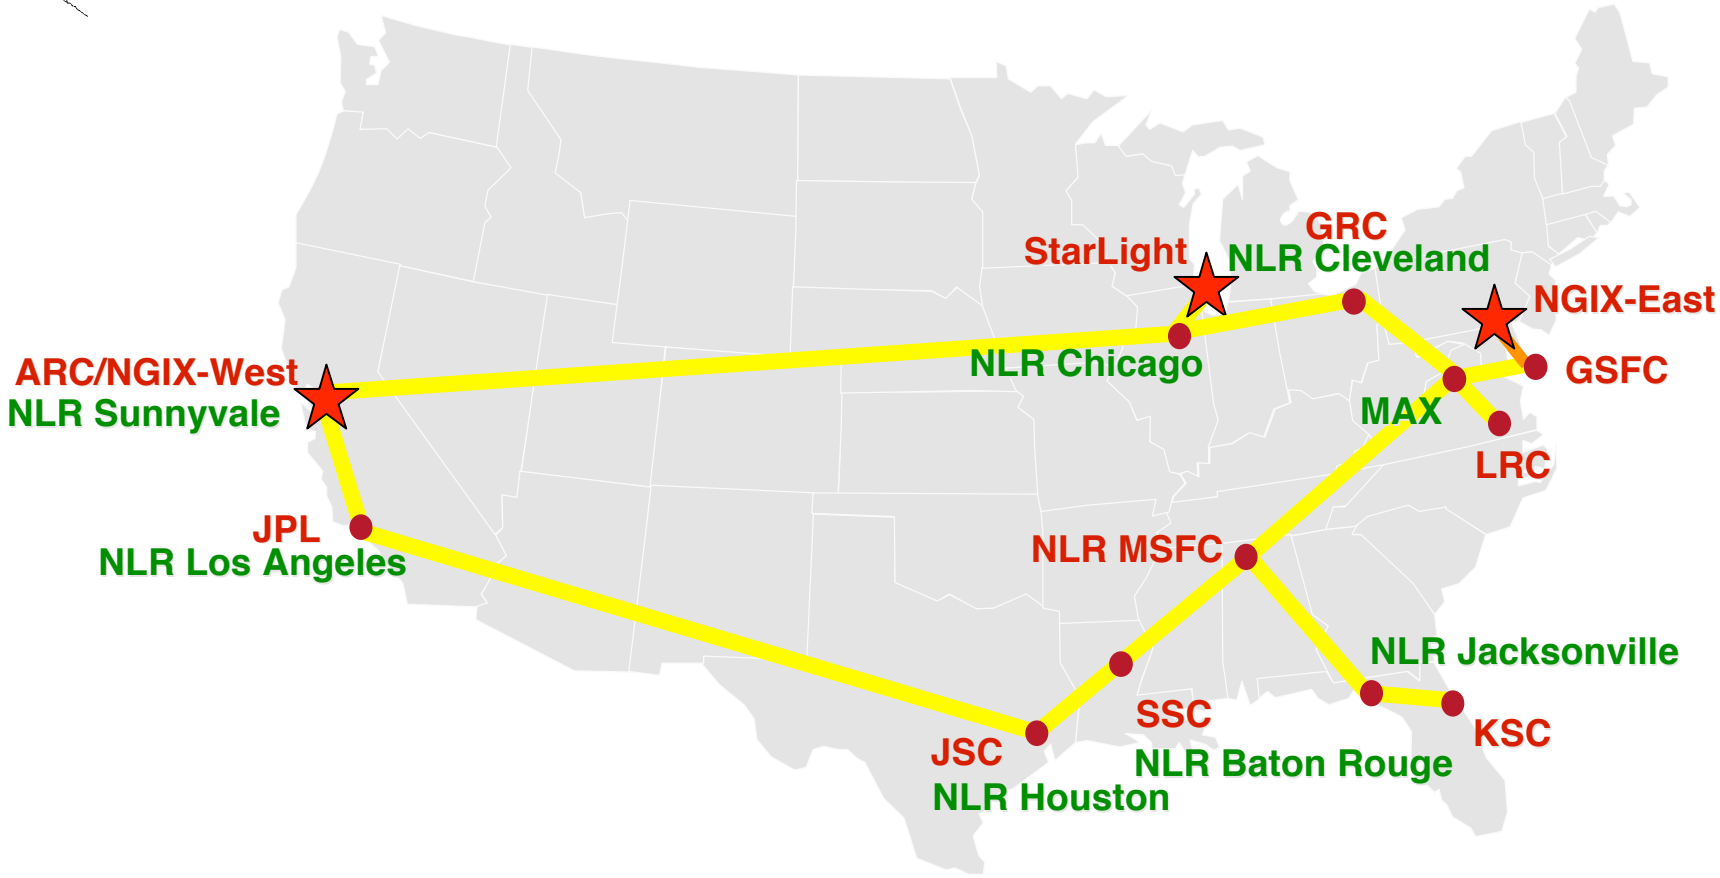

High Performance Networking and Remote Data Access
GSFC L-Net for NCCS and Science Buildings

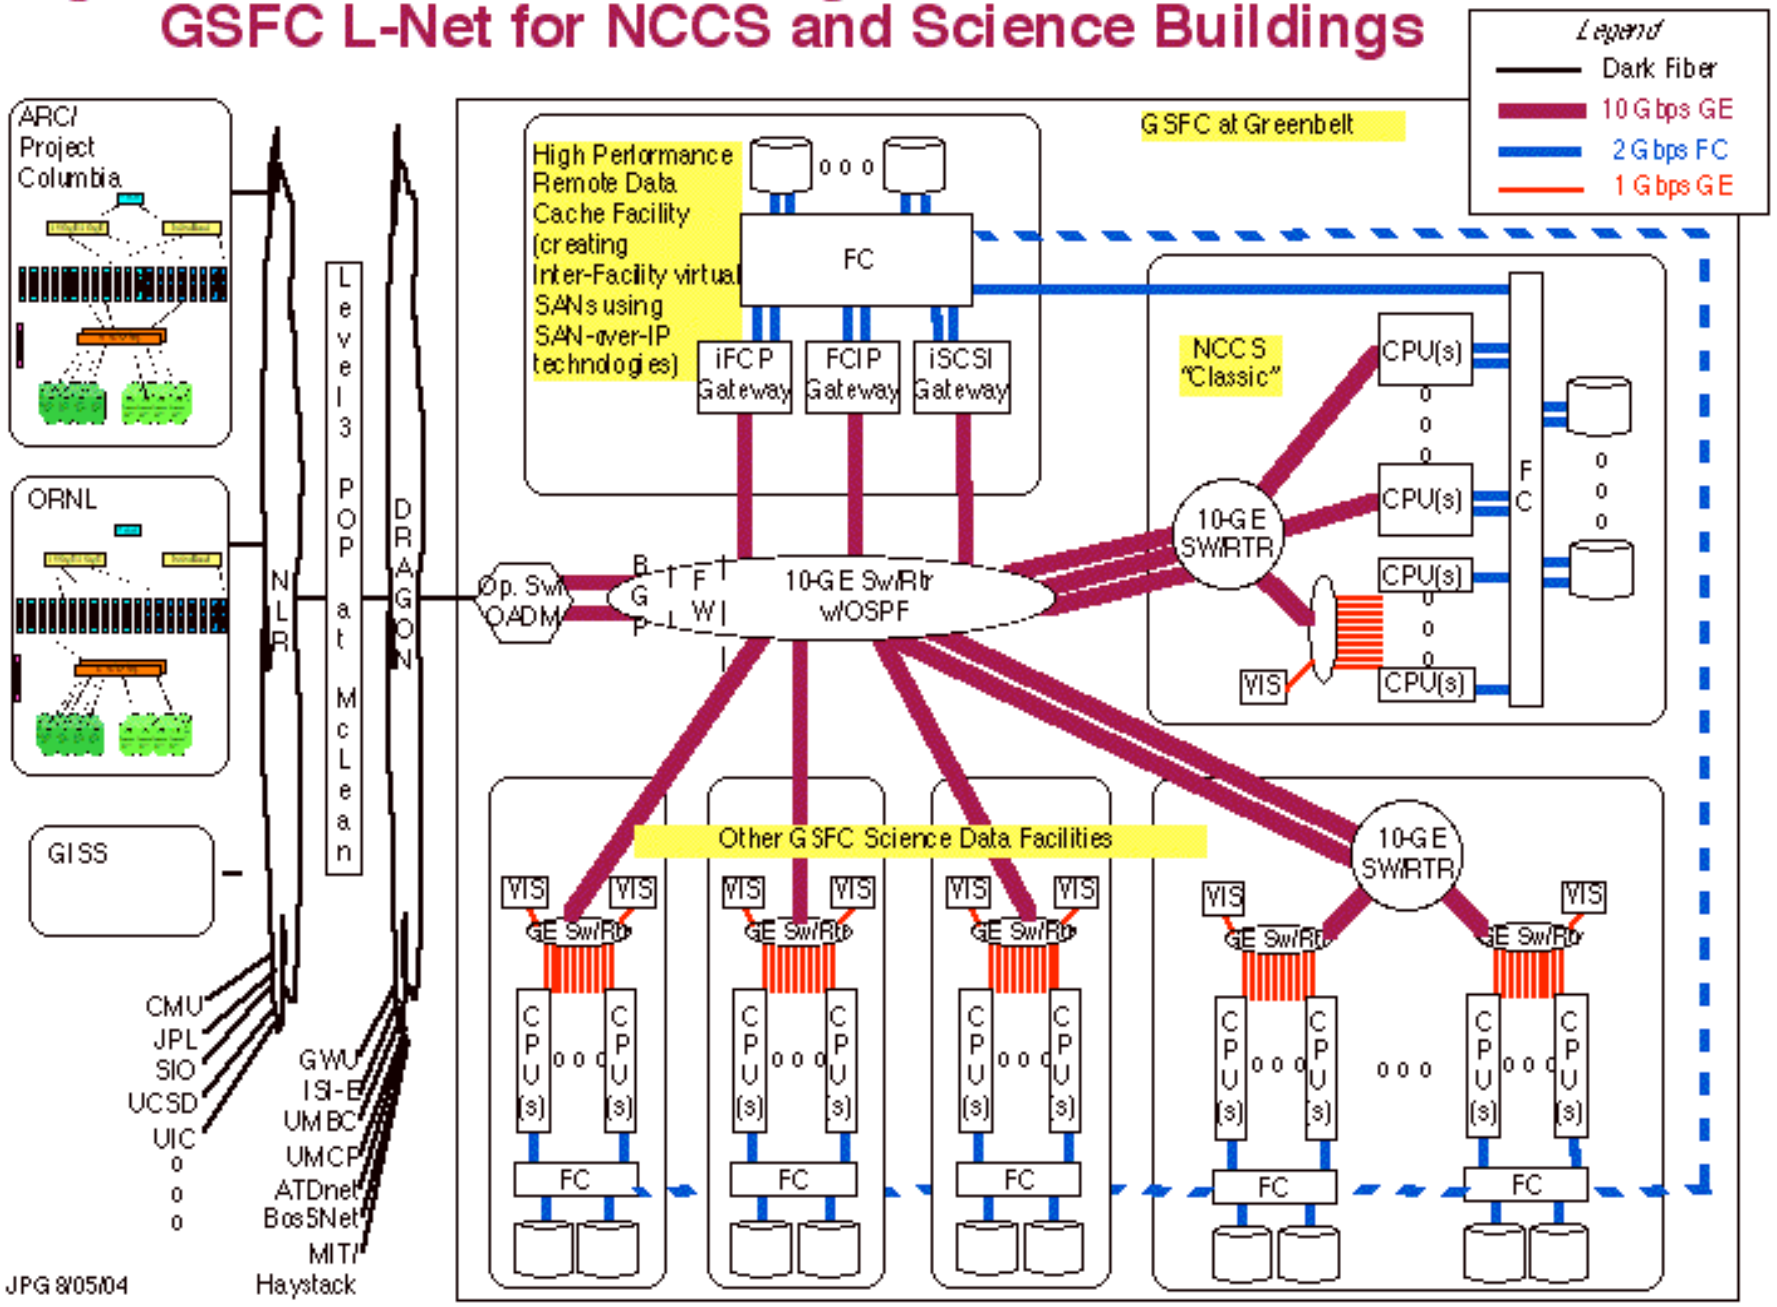

Science-Enabling Network Technology Highlights of GSFC's 10-Gbps Lambda Network Project

Special Acknowledgements

GSFC Internal

- **High End Computer Network Team**
  - **Bill Fink/606.1**
  - **Kevin Kranacs/585**
  - **Paul Lang/ADNET/606.1**
  - **Aruna Muppalla/ADNET/606.1**
  - **Jeff Martz/CSC/606.2**
  - **Mike Steffenelli/CSC/606.2**
  - **Kevin Fisher/586/UMBC coop**
- **ESDIS Network Prototyping Lab**
  - **George Uhl/SWALES/423**
- **ESTC Computing Technology Project**
  - **PM: Jim Fischer/606**
- **IT Pathfinder Working Group**
  - **Chair: Dr. Milton Halem/Emeritus & UMBC**
- **Thunderhead Cluster**
  - **John Dorband/696**
  - **Josephine Palencia/RSTX/606.2**

GSFC External

- **Nationa LambdaRail**
  - **CEO: Tom West**
  - **Net Eng Lead: Debbie Montano**
- **OptIPuter Project (NSF-funded)**
  - **PI: Dr. Larry Smarr/UCSD**
  - **Co-PI: Dr. Tom DeFanti/UIC**
  - **PM: Maxine Brown/UIC**
  - **UCSD Net Eng: Greg Hidley, Arron Chin, Phil Papodopolos**
  - **UIC Net Eng: Alan Verlo, Linda Winkler**
- **DRAGON Project (NSF-funded)**
  - **PI: Jerry Sobieski/UMCP**
  - **Co-I: Tom Lehman/USC-ISI/E**
  - **Net Eng: Chris Tracy/UMCP**
- **NASA Research and Education Network**
  - **DPM: Kevin Jones/ARC**

Science-Enabling Network Technology Highlights of GSFC's 10-Gbps Lambda Network Project

GSFC Lambda Network Project Website

- http://cisto.gsfc.nasa.gov/IRAD_Lambda.html
  - **Designs**
    - GSFC Local Network Part (i.e., within GSFC)
    - Regional Network Part (i.e., between GSFC in Greenbelt, MD, & Level3 POP in McLean, VA, typically involving the DRAGON optical network)
    - Transcontinental Network Part (i.e., use of NLR, GSFC 10-GE switch & workstations in the Level3 POP in McLean, VA, & remote end users/sites)
  - **Implementation Status**
    - GSFC Local Network Part
    - Regional Network Part
    - Transcontinental Network Part
  - **Presentations/Events in the News**
    - Eg: P. Gary's 18Feb05 presentation at GSFC's FY04 IRAD Colloquium <http://cisto.gsfc.nasa.gov/L-Netpdfs/FY04IRADGARY.pdf>
    - Live Demonstration of 21st Century National-Scale Team Science <http://www.calit2.net/articles/article.php?id=660>
  - **Related Links (e.g., DRAGON, HOPI, NLR, OptIPuter, …)**