The Neutron Science TeraGrid Gateway: TeraGrid Cyberinfrastructure at ORNL

Presented by

John W. Cobb, Ph.D.

Computer Science and Mathematics

Office of Cyberinfrastructure
National Science Foundation

In collaboration with many teams:

NSTG, SNS, McStas group, Open Science Grid, Earth Systems Grid, TechX Corp, OGCE, UK eScience effort, and the TeraGrid Partners teams.

SC07

TeraGrid

OAK RIDGE National Laboratory

Outline

1. Cyberinfrastructure (CI)
2. TeraGrid
3. The Spallation Neutron Source
4. The Neutron Science TeraGrid Gateway (NSTG)
5. Other collaborations

Outline

1. Cyberinfrastructure (CI)
2. TeraGrid
3. The Spallation Neutron Source
4. The Neutron Science TeraGrid Gateway (NSTG)
5. Other collaborations

Cyberinfrastructure overview (from 2006)

Personal view: A hierarchy of needs†

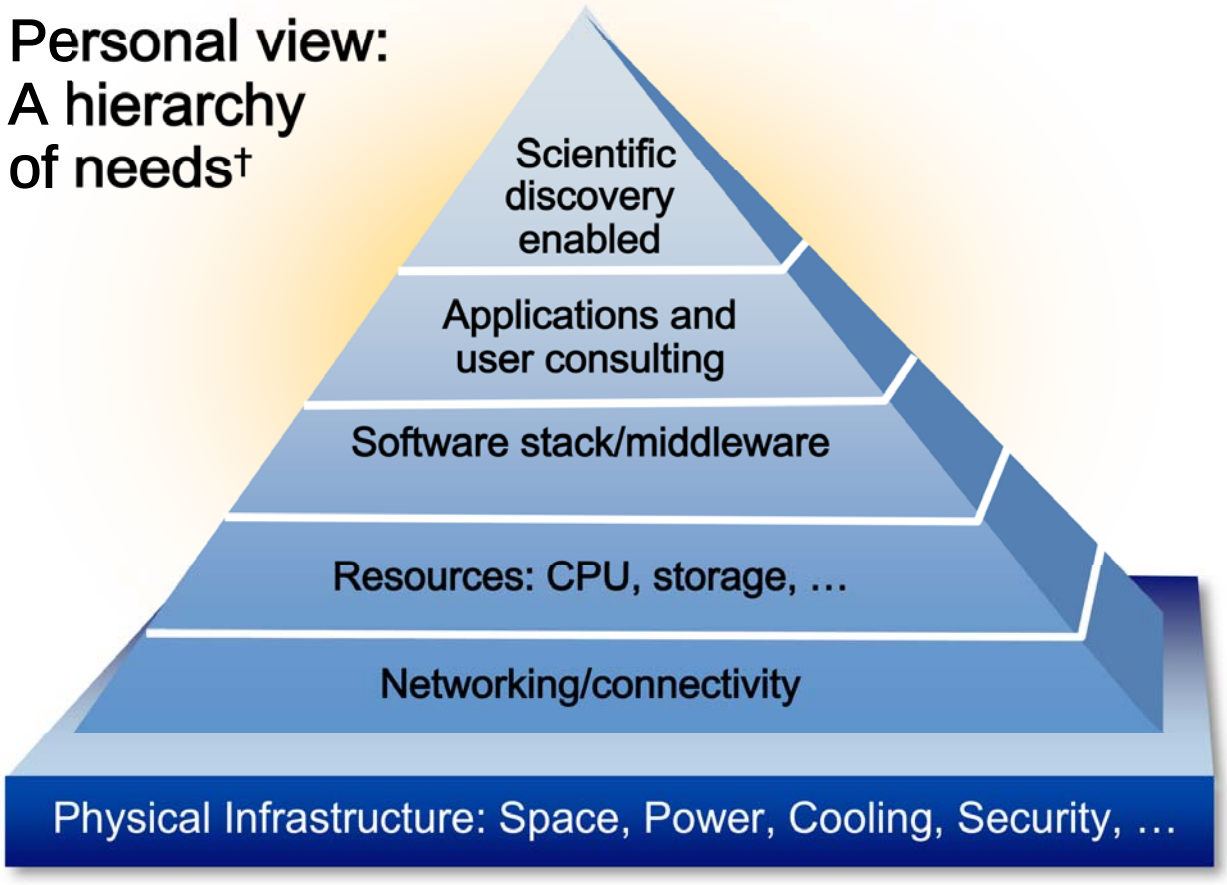

2003 Blue Ribbon Panel: "Atkins Report"

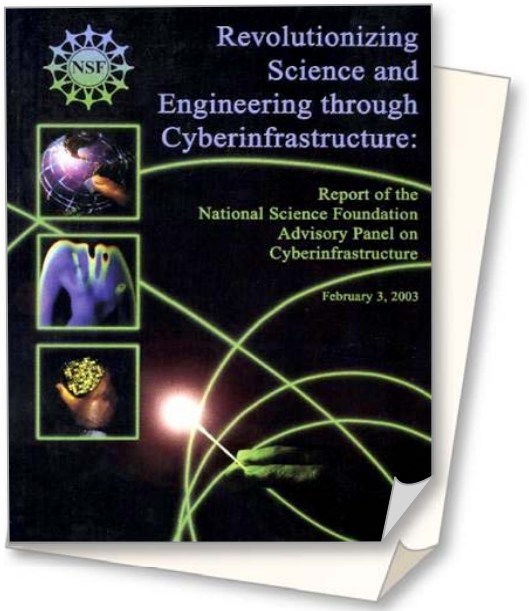

"Like the physical infrastructure of roads, bridges, power grids, telephone lines, and water systems that support modern society, **"Cyberinfrastructure"** refers to the distributed computer, information and communication technologies combined with the personnel and integrating components that provide a long-term platform to empower the modern scientific research endeavor" ‡

†Maslow, A. H. (1943). *A Theory of Human Motivation.* Psychological Review, 50, 370-396.

‡ David Hart in NCSA News release "National Science Foundation Releases New Report from Blue-Ribbon Advisory Panel on Cyberinfrastructure" February 3, 2003 http://access.ncsa.uiuc.edu/Releases/03Releases/02.03.03_National_S.html as quoted on Cyberinfrastructure Wikipedia entry.

Cyberinfrastructure (CI) Maturing

- CI is now more understood
- Infrastructure is recognized as critical to service and research success
- "We believe, we get tremendous competitive advantage by essentially building our own infrastructures." Eric Schmidt, Google CEO*
- National focus now shifting to enabling scientific discovery
  - From data to knowledge: Enhancing human cognition and generating new knowledge from a wealth of heterogeneous digital data
  - Understanding complexity in natural, built, and social systems: Deriving fundamental insights on systems comprising multiple interacting elements
  - Building virtual organizations: Enhancing discovery and innovation by bringing people and resources together across institutional, geographical and cultural boundaries
- Cyber Enabled Discovery is predicated upon advanced, national-scale cyberinfrastructure

* Source The New York Times online edition, Sept. 30, 2007

Outline

1. Cyberinfrastructure (CI)
2. TeraGrid
3. The Spallation Neutron Source
4. The Neutron Science TeraGrid Gateway (NSTG)
5. Other collaborations

TeraGrid overview

"The TeraGrid is the world's largest, most powerful and comprehensive distributed cyberinfrastructure for open scientific research. It currently supports more than 1,000 projects and over 4,000 researchers geographically spanning the entire United States."

National Science Foundation in press release 07-09

- Noun: Cyberinfrastructure
- Adjectives
  - Most powerful
  - Comprehensive
  - Distributed
  - For open research

TeraGrid network

- Dedicated Links
- Resource to Resource
- Bandwidth unit: 10 Gbps
- Substrate for high performance high quality data movement and communications

And …

OAK RIDGE National Laboratory

2 New Resource Providers

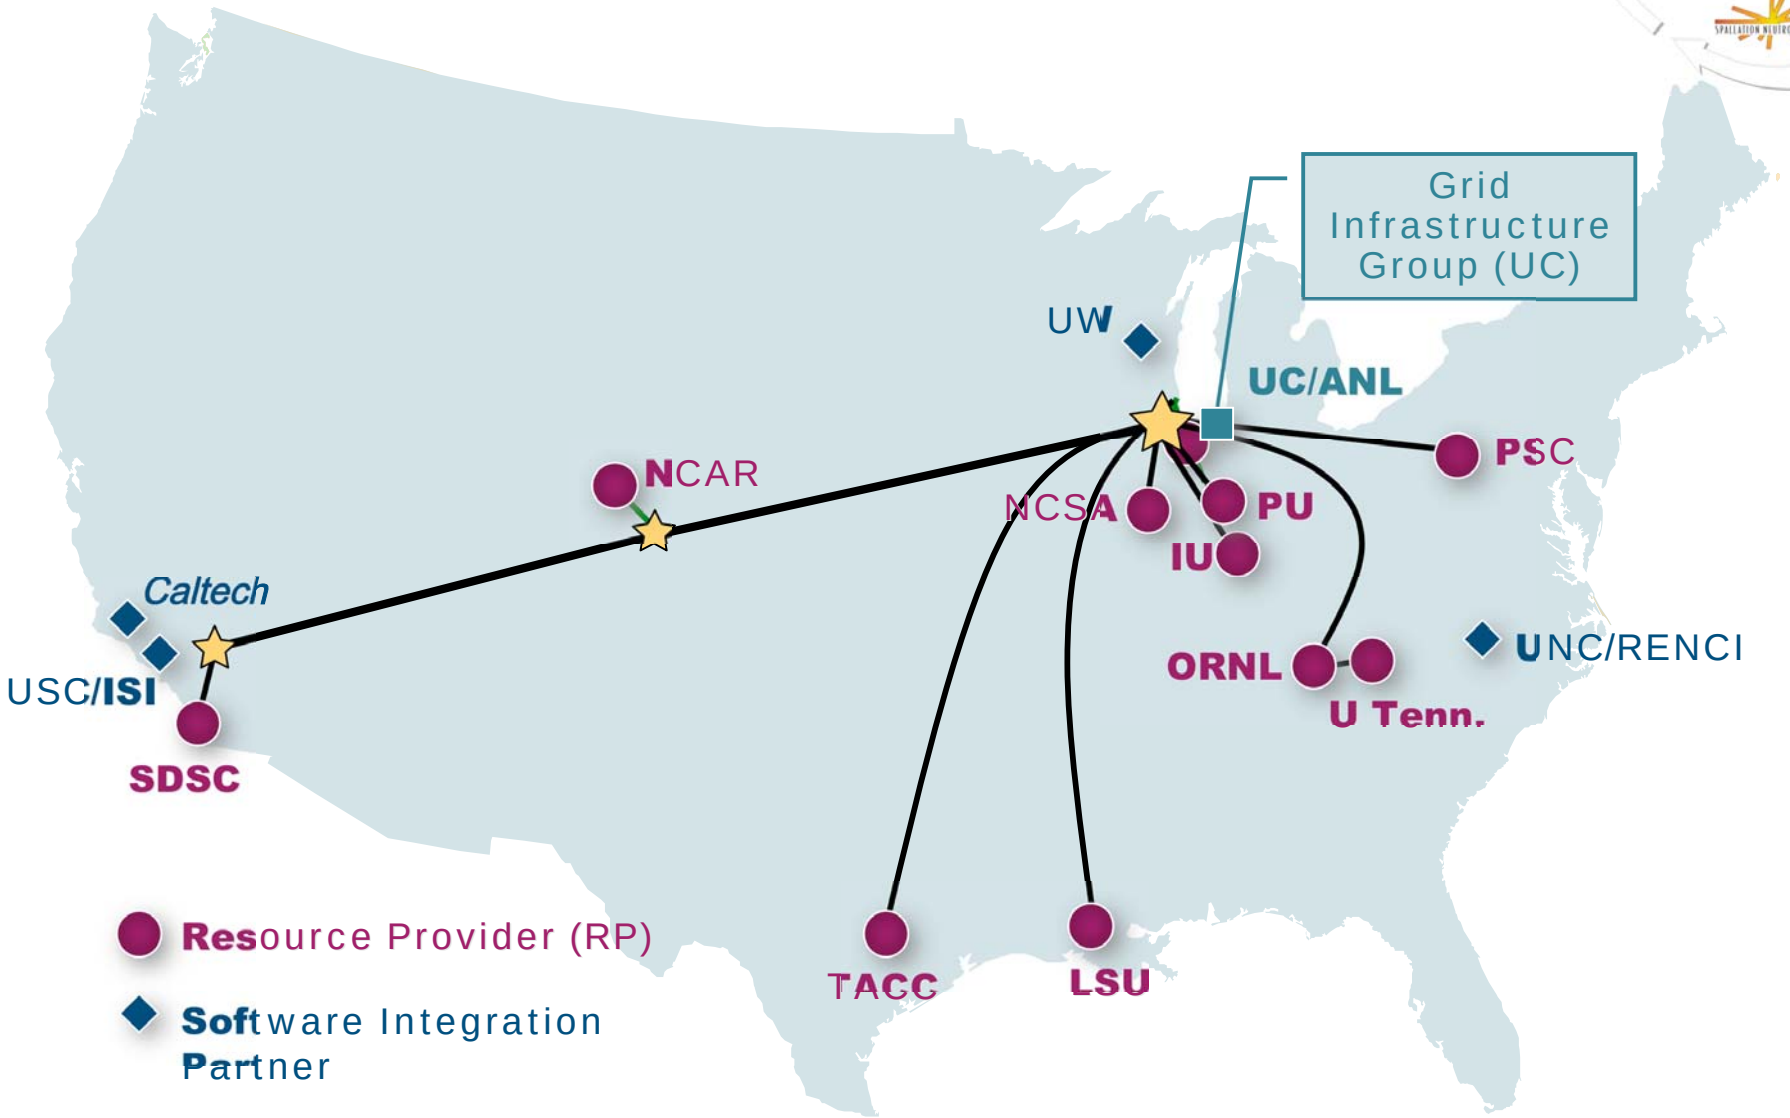

TeraGrid resources

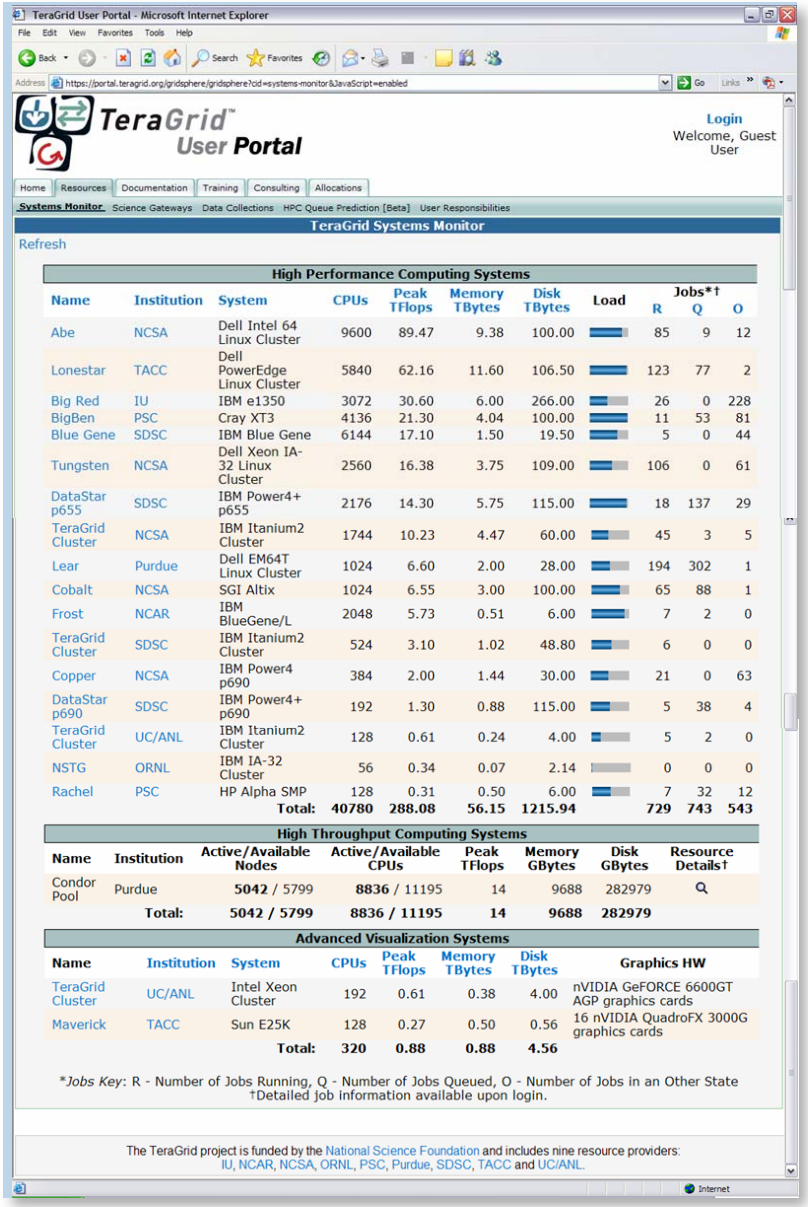

- 9 distributed centers (soon 11)
- > 300 TFlops today
  +400 TF in January
  +1 PF in late 2008
  On path to double every year
- > 1.5 PB disk
- Multi PB archive
- Different Resource types
  - Capability
  - Capacity
  - High throughput
  - Visualization
  - Specialized use
- Common user environments
- Tuned data transfer

TeraGrid distributed and unified management

- All TeraGrid resources are allocated via a peer-reviewed, unified national process, XRAC
- Proposals: POPS See “Allocations” tab on portal
- Single sign-on across TeraGrid (GSI)
- Common software stack across the sites
- TeraGrid portal for computer user access and ease
  - Unified user support TeraGrid-wide
  - Unified ticket system (7x24 coverage)
  - User services
- Advanced support for TeraGrid applications (ASTA)
- Coordinated high performance data movement
  - GridFTP
  - Wide area file systems (GPFS in production, Lustre, pNFS in pilots)

https://portal.teragrid.org

TeraGrid serves NSF users and the open science community

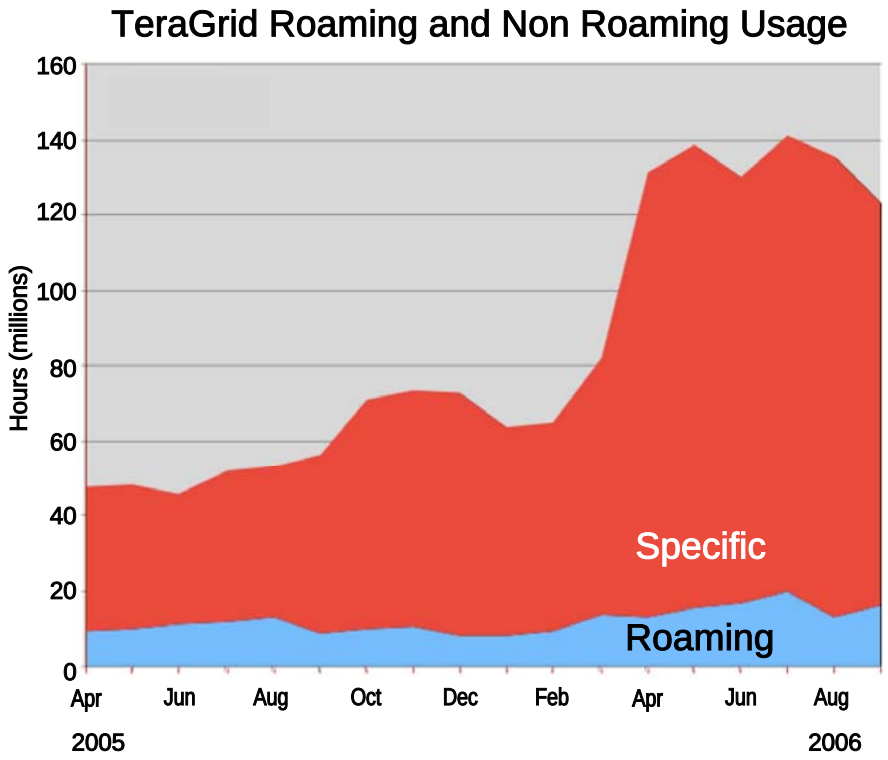

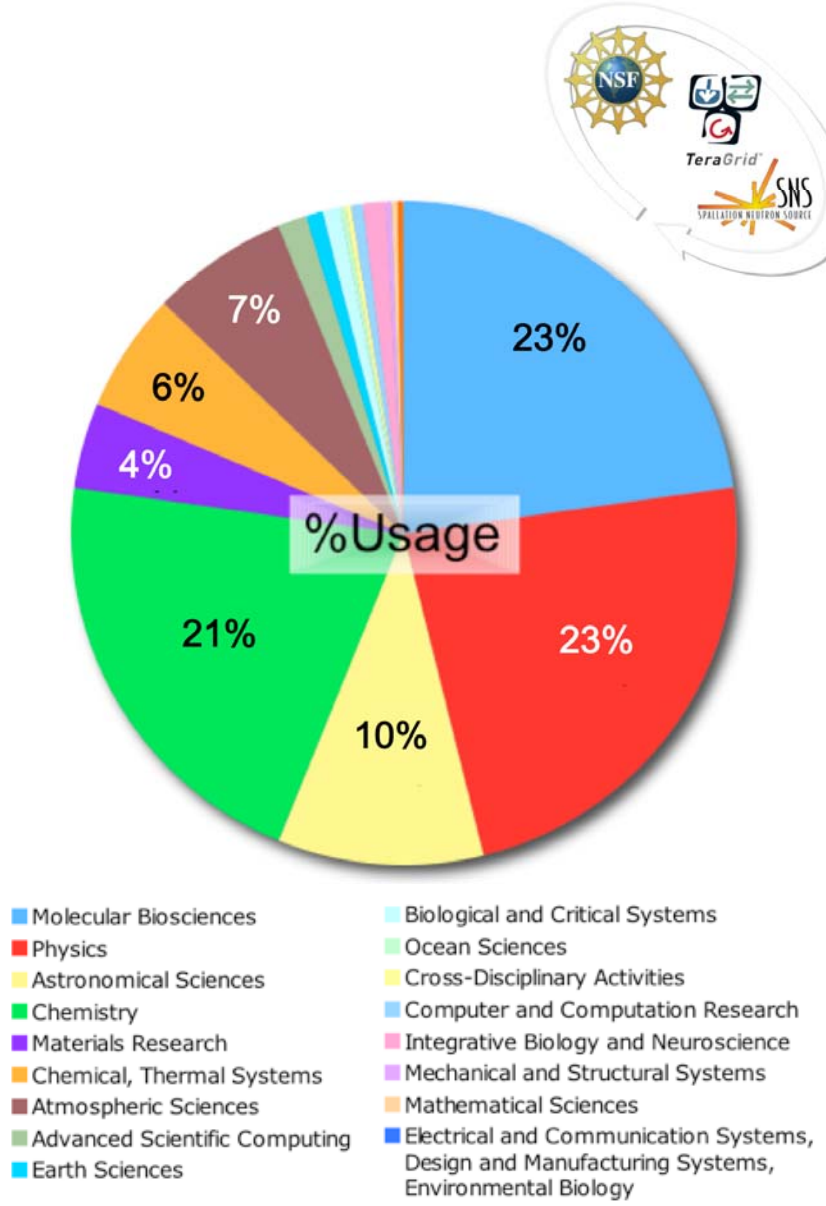

Science gateways

- TeraGrid size: How can we manage it?
  - Today: 4000+ user accounts across 20 resources
  - Target: 10–100× increase in user base; increase in number of resources
- TeraGrid Science Gateways
  1. Web portal with users in front and TeraGrid services in back
  2. Grid-bridging gateways: Extending the reach of a community grid (devoted to a single area of science) so it can use TeraGrid resources
  3. Application programs running on users' machines (e.g., workstations and desktops) that access TeraGrid (and other) services
- The NSTG is one of 20+ existing gateways

Charting new territory for wide cyberinfrastructure services:

- Federated user management
- Callback end-user identification with automated service auditing across federated identity space "gateway community accounts"
- Grid interoperability
- Participation in attribute-based authorization

Outline

1. Cyberinfrastructure (CI)
2. TeraGrid
3. The Spallation Neutron Source
4. The Neutron Science TeraGrid Gateway (NSTG)
5. Other collaborations

Spallation Neutron Source

"The next generation of materials research"

Large world-class user facility	• Neutrons as a probe of matter
Construction complete	• 1998–2006, TPC $1.4 billion
Accelerator power-up	• Ongoing, 183 kW record set in August 2007
instruments	• Commissioning 3 to 4 per year
User program	• 3 allocated instruments in Fall 2007 : – Backscattering spectrometer – Liquids reflectometer – Magnetism reflectometer

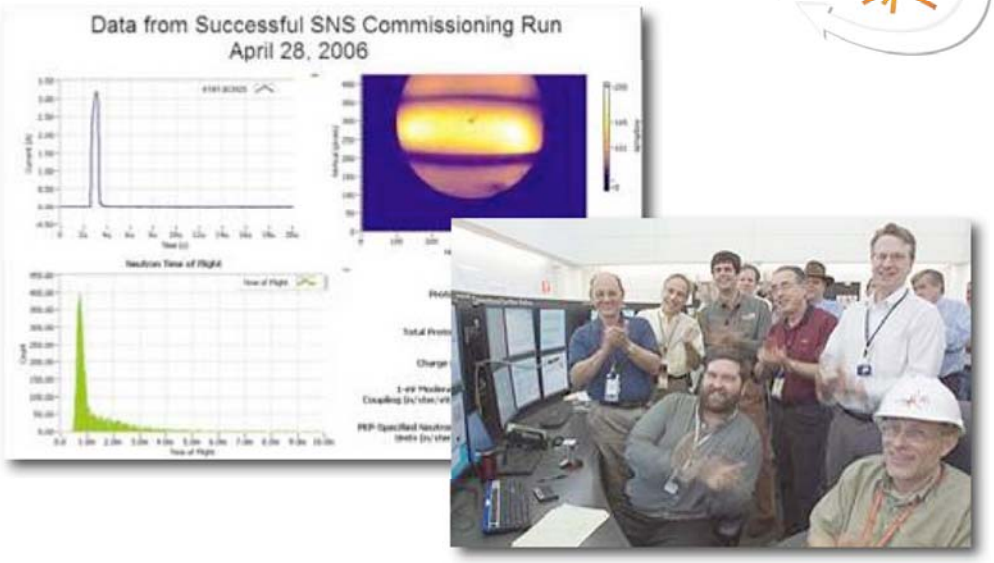

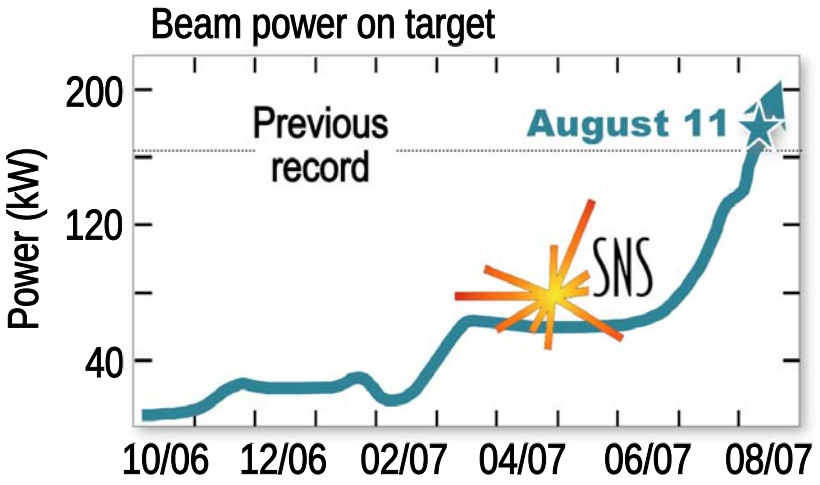

OAK RIDGE National Laboratory

SNS data access via portal
(Reduction, analysis, and simulation as well)

Turnkey solutions are being requested by users and user groups

Problem	SNS approach
Providing advanced tools in an easy fashion while retaining customizability	• Neutron Science User Portal – Totally in Java

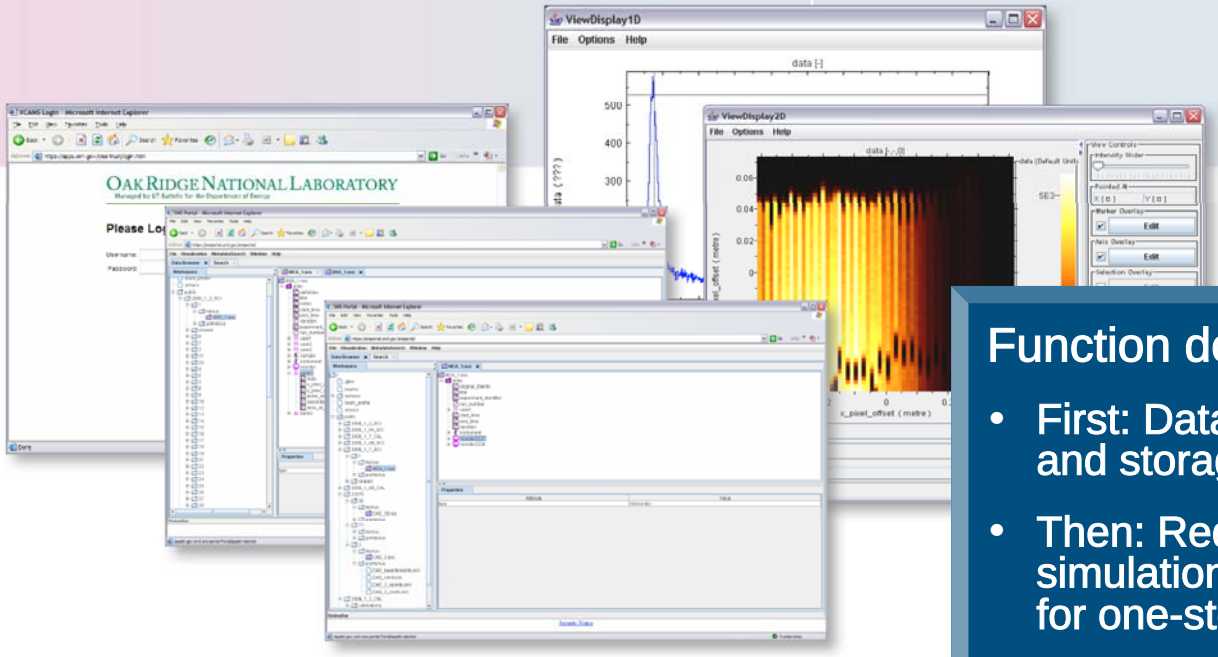

Function deployment order

- First: Data acquisition and storage/stewardship
- Then: Reduction, analysis, simulation, proposals for one-stop access

Acknowledgement: S. Miller and SNS Advisory SW group; J. Kohl, S. Vazhkudai and ORNL/CSM division)

Neutron science portal view

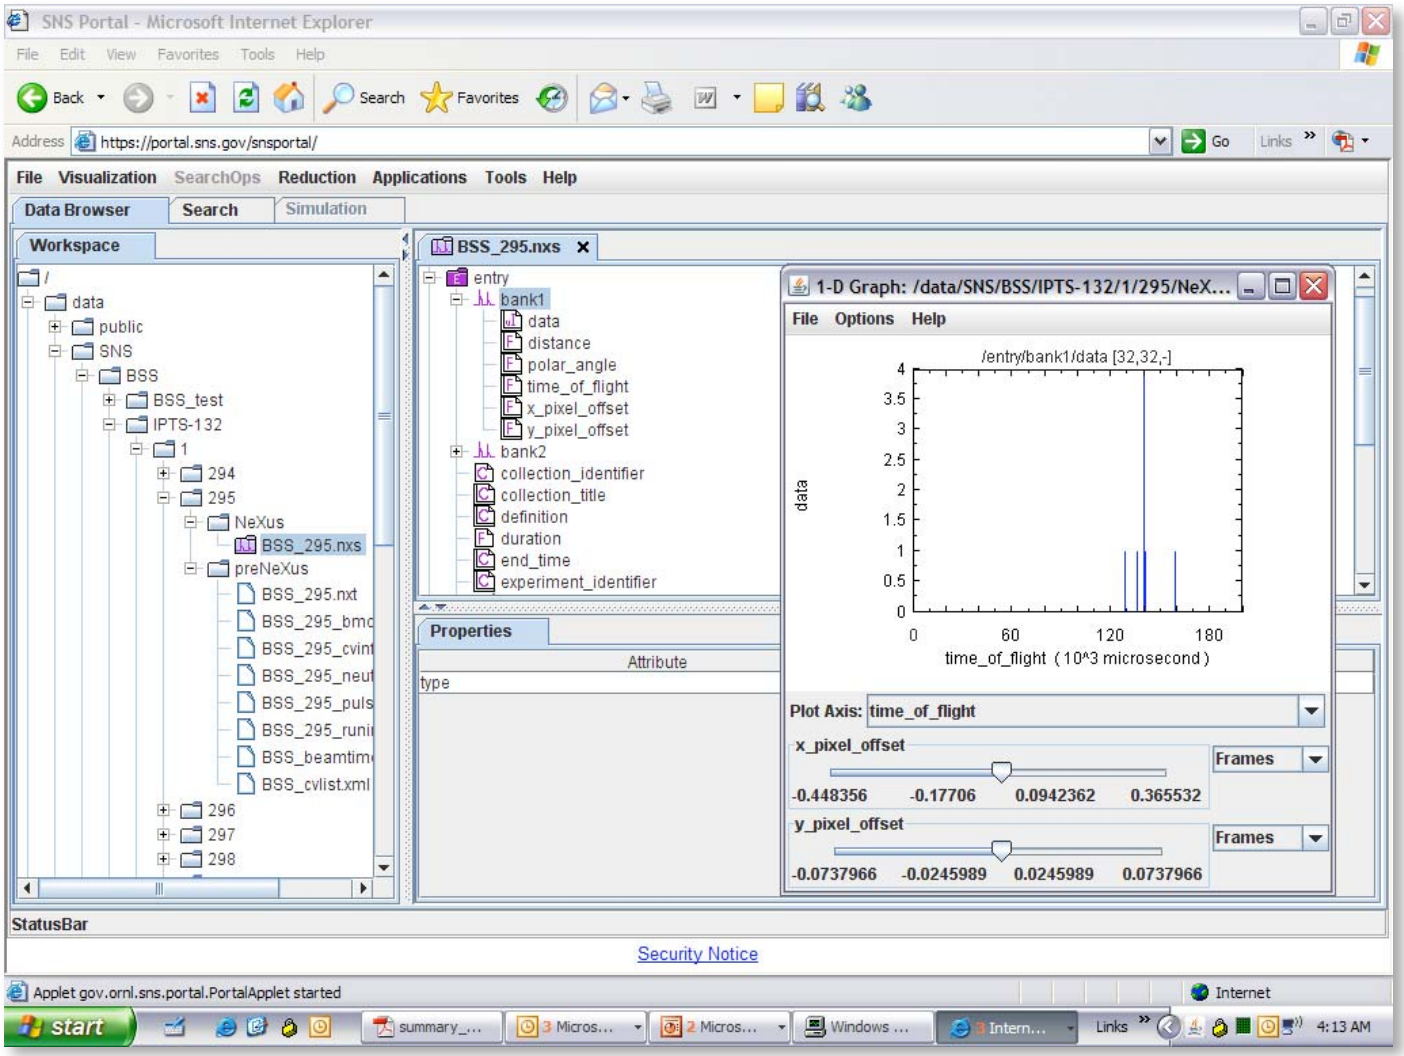

NeXus
A community-driven standards process

- Data format for neutron, X-ray, and muon science
- Based on HDF5
- Self-describing
- Use increasing as new facilities come online
  - SNS
  - OPAL at ANSTO
  - ISIS second target station
  - J-SNS
- McStas (neutron Monte Carlo ray-trace simulation tool) can read and write NeXus format

http://www.nexusformat.org/

SNS cyberinfrastructure to enable the Neutron Science Portal

Software and systems	• Data acquisition system (DAS) at each instrument • iCAT captures raw data – Stores data locally to instrument in a pre-NeXus format – Broadcast data downstream to analysis team – Catalogues raw data • Translation service: Raw data to NeXus formats • Messaging service • Application orchestration • Integrated authentication with facility software XCAMS
Hardware	• Instrument local DAS and analysis machines • Facility centralized services including storage, application hosting, execution servers
TeraGrid	• Opportunity to access HPC resources via the Neutron Science Portal

Outline

1. Cyberinfrastructure (CI)
2. TeraGrid
3. The Spallation Neutron Source
4. The Neutron Science TeraGrid Gateway (NSTG)
5. Other collaborations

Neutron Science TeraGrid Gateway (NSTG)

One of 9 TeraGrid partner Resource Providers

Focus areas	Resources provided
• Neutron science • Connecting facilities with cyberinfrastructure • Bridging cyberinfrastructures • Data movement within and across TeraGrid	• Outreach to a specific science community (neutron science) – Expertise – Technology transfer – Education, in a broad sense – User outreach • A science gateway interface for neutron science • Exploration of large science facility integration with national-scale cyberinfrastructure • Combine TeraGrid computational resources with available neutron scattering datasets for analysis and simulation comparison

OAK RIDGE National Laboratory

NSTG operational components

- Wide area network connections to TeraGrid
- Local connections, at maximum bandwidth, to local facilities (SNS and HFIR)
- Modest local cluster with complete TeraGrid environment for small local TeraGrid jobs
- Long-term HPSS archival storage via ORNL's LCF/NCCS
- System operations
- Application consulting

TeraGrid footprint: ORNL

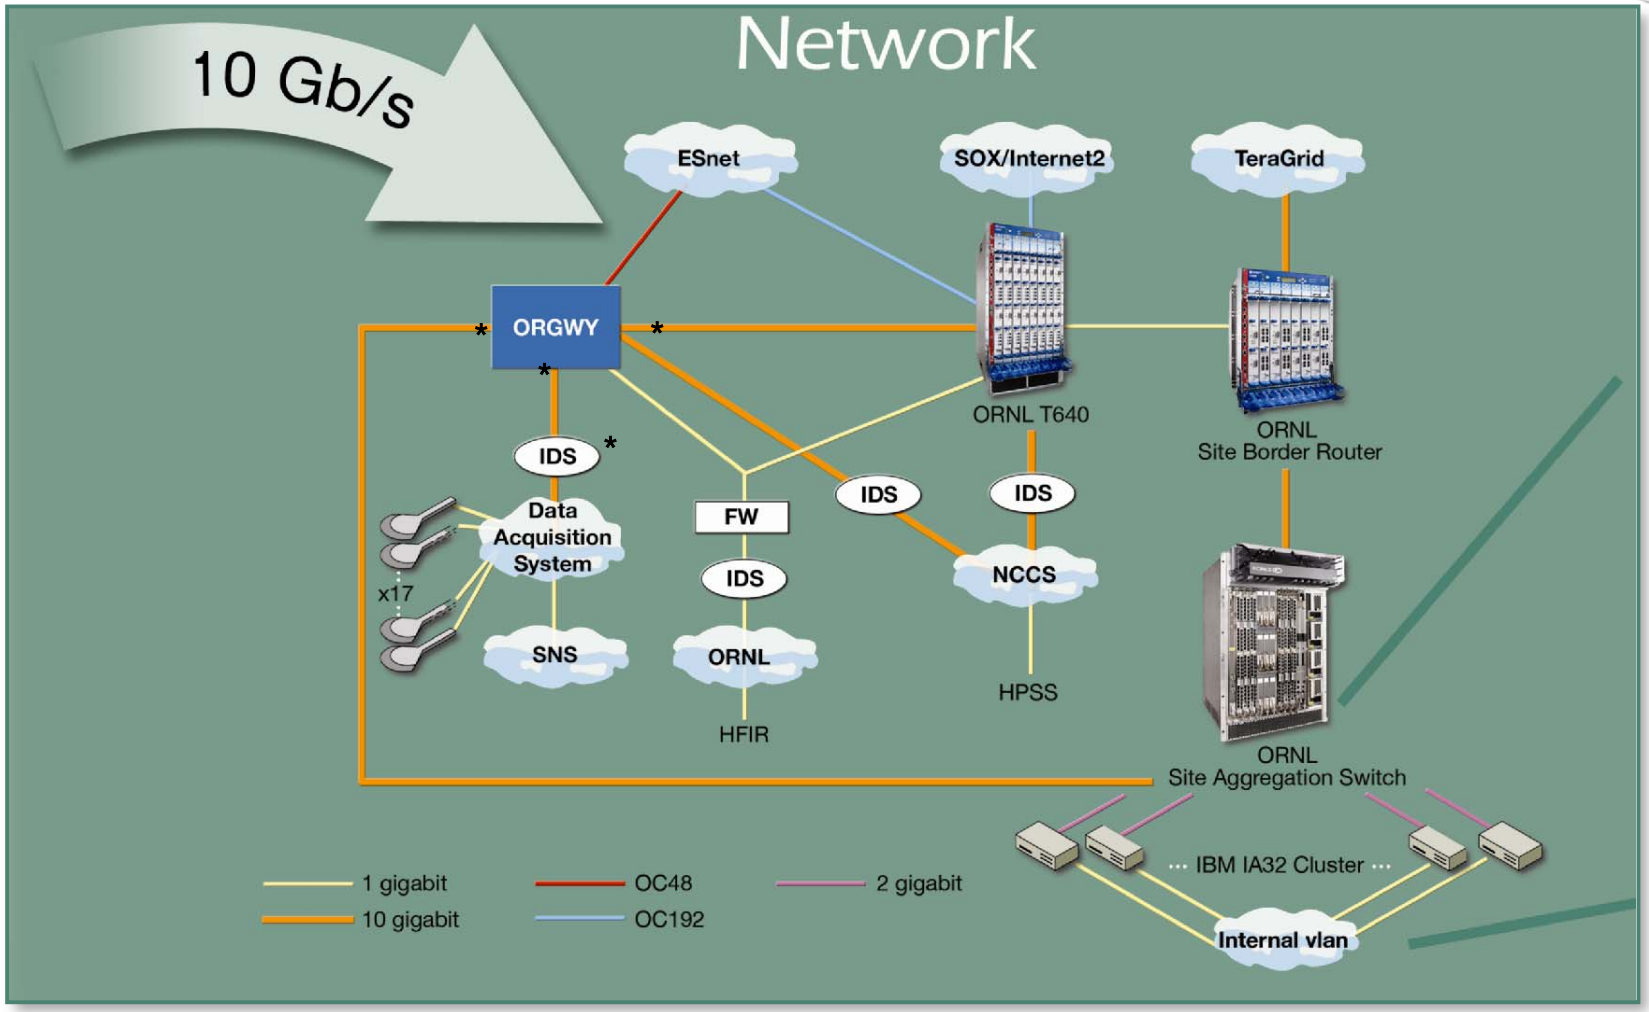

NSTG results: Highlight

First principles sample kernel simulations of experimentally relevant materials (E. Farhi, ILL)

- Ability to launch McStas neutron Monte Carlo ray tracing instrument simulations on the TeraGrid directly from Neutron Science Portal
- Production use of TeraGrid community account “Jimmy Neutron”
  - New TeraGrid science gateway scaling technology
- In development: A generalized fitting service for experimental data integrated into neutron science portal and using TeraGrid computational resources
- Deployment of wide area Lustre pilot client; write speeds of 9.72 Gbs over a 10-Gbs wide area connection
- Testing of GridFTP servers with 10-Gbs end-to-end (NIC-switch-router-WAN-router-switch-NIC); >90% of theoretical maximum

Neutron science simulations

Next step: Provide analysis and relevant simulations in support of neutron science

Instrument design optimization:

Moving simulations to TeraGrid allows larger runs and faster turnaround. Assisted in effort to redesign instrument to lower cost. Simulations provided confidence that resolution would not be sacrificed.

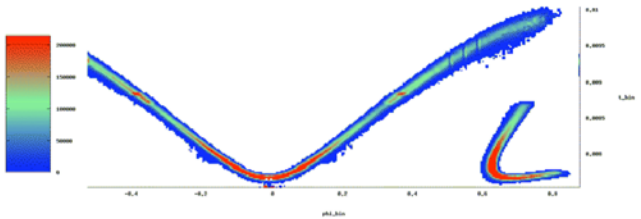

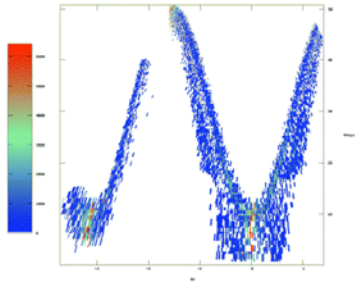

Porting and deployment of simulation tools to TeraGrid cyber-infrastructure: McStas and IDEAS

Many codes currently I/O limited. Next step is to modify I/O strategy in codes to take advantage of I/O architectures.

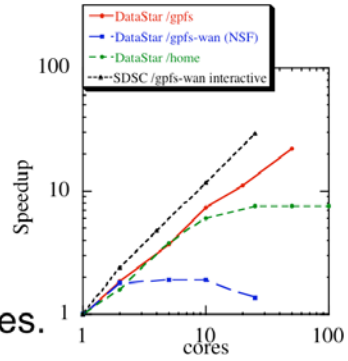

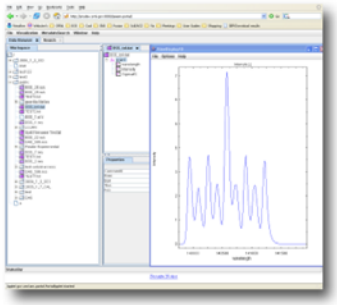

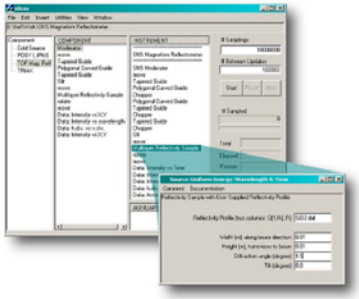

Improved and faster data analysis: Reflectometry

Faster and better χ^2 minimization.

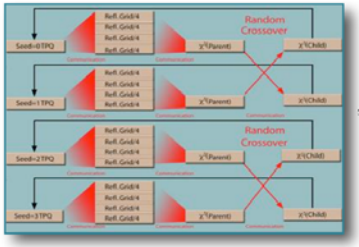

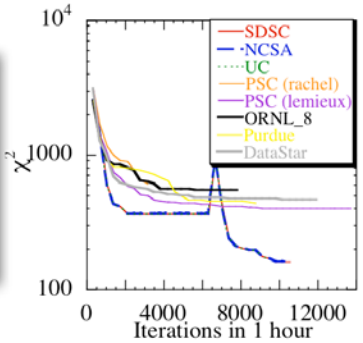

OAK RIDGE National Laboratory

Portal-initiated simulations . . . under the covers!

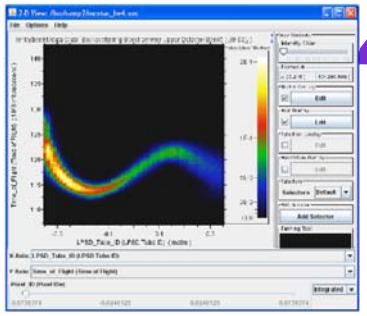

Interactive User Inputs

Portal Applet (Simulation GUI)

https

Back-End Web Server

Simulation Servlet

"runas", params

Job Information Service

Authorization As User

Globus Front-End

SNS-user

globus-job-status

user "Jimmy Neutron"

ssh

globus-job-submit

grid-proxy-init

Community Certificate

Results!

private key

TeraGrid Computational Resources as "Jimmy Neutron" Community Account

GridFTP

globus-url-copy

NSTG was and is "ahead of the curve"

- Proposed a Science Gateway as part of ORNL RP in 2003, before creation of TG science gateways program
- Included a focus on cyber-enabled science from the outset: Enabling neutron science
- Early (but not first) intergrid production operation
- Early (but not first) concentration on intra-TeraGrid high-performance wide area data transfer
- Concentration on high-performance transfer between TeraGrid and external endpoints
- Early adopter and pathfinder postures fit with NSTG RP role

Outline

1. Cyberinfrastructure (CI)
2. TeraGrid
3. The Spallation Neutron Source
4. The Neutron Science TeraGrid Gateway (NSTG)
5. Other collaborations

NSTG participates in several collaborative areas

Neutron science:

- SNS and other facilities
- McStas instrument simulation software
- Portal development tools and deployments
  - OGCE, GumTree Eclipse RCP

Data transport:

- Within TeraGrid partners
- In wider area (REDDnet)

Inter-grid operation:

- Open Science Grid
- Earth Systems Grid

McStas

- Widely used neutron instrument simulation package for instrument design and experiment planning
- Development team at Risø National Laboratory (Denmark) and Institut Laue-Langevin (France)
- Monte Carlo ray tracing of neutron path through beamline elements
- Many advanced features
- Use cases:
  - Instrument design and construction
  - Experiment interpretation
  - Experiment planning
  - Experiment proposal evaluation
- Challenge: Simulation of sample and sample environment.
- E. Farhi (largest NSTG cluster user) has just completed a preliminary study to computationally compute scattering kernels using VASP (private communication)

http://neutron.risoe.dk/

McStas: Sample kernel simulations
(private communication, E. Farhi, ILL)

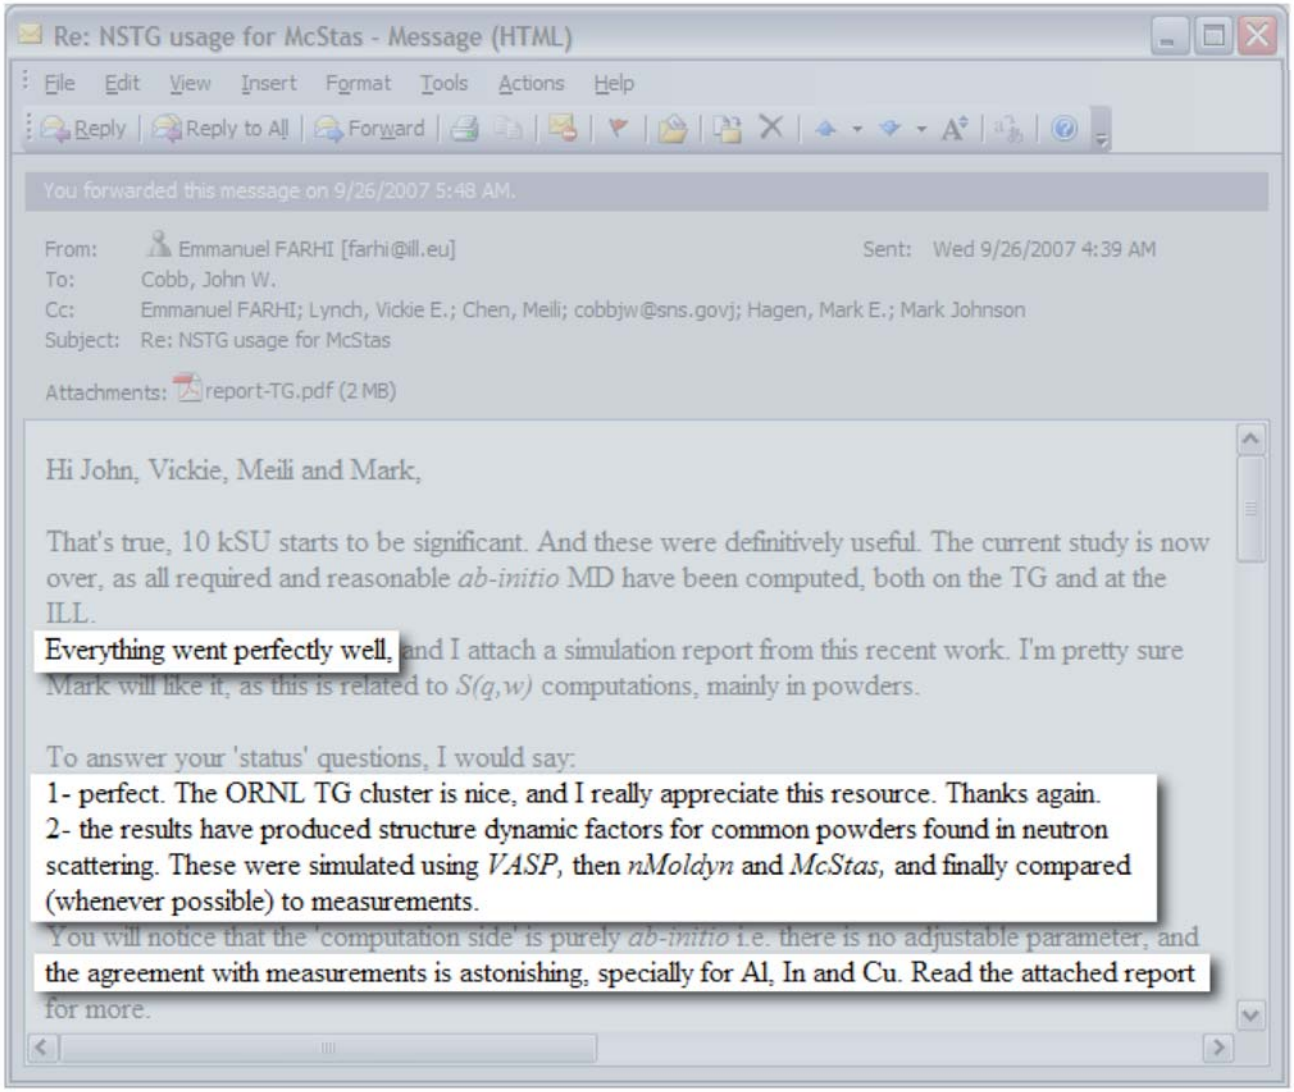

Re: NSTG usage for McStas - Message (HTML)

File Edit View Insert Format Tools Actions Help

Reply | Reply to All | Forward

You forwarded this message on 9/26/2007 5:48 AM.

From: Emmanuel FARHI [farhi@ill.eu] Sent: Wed 9/26/2007 4:39 AM
To: Cobb, John W.
Cc: Emmanuel FARHI; Lynch, Vickie E.; Chen, Meili; cobbjw@sns.govj; Hagen, Mark E.; Mark Johnson
Subject: Re: NSTG usage for McStas
Attachments: report-TG.pdf (2 MB)

Hi John, Vickie, Meili and Mark,

That's true, 10 kSU starts to be significant. And these were definitively useful. The current study is now over, as all required and reasonable *ab-initio* MD have been computed, both on the TG and at the ILL.

Everything went perfectly well, and I attach a simulation report from this recent work. I'm pretty sure Mark will like it, as this is related to *S(q,w)* computations, mainly in powders.

To answer your 'status' questions, I would say:

1- perfect. The ORNL TG cluster is nice, and I really appreciate this resource. Thanks again.
2- the results have produced structure dynamic factors for common powders found in neutron scattering. These were simulated using *VASP*, then *nMoldyn* and *McStas*, and finally compared (whenever possible) to measurements.

You will notice that the 'computation side' is purely *ab-initio* i.e. there is no adjustable parameter, and the agreement with measurements is astonishing, specially for Al, In and Cu. Read the attached report for more.

GumTree

http://gumtree.codehaus.org/

- An integrated scientific workbench
- Eclipse-based rich client
- Neutron science data acquisition and analysis software
- Under development at Australian Nuclear Science and Technology Organisation's (ANSTO) Bragg Institute for instruments on the OPAL reactor source
- ANSTO Team: Nick Hauser, Tony Lam
- Mark Green (TechX, Buffalo) looking to integrate into SNS portal and application efforts
- See OGCE workshop
  - November 11–12, 2007

REDDnet
Research and Education Data Depot Network

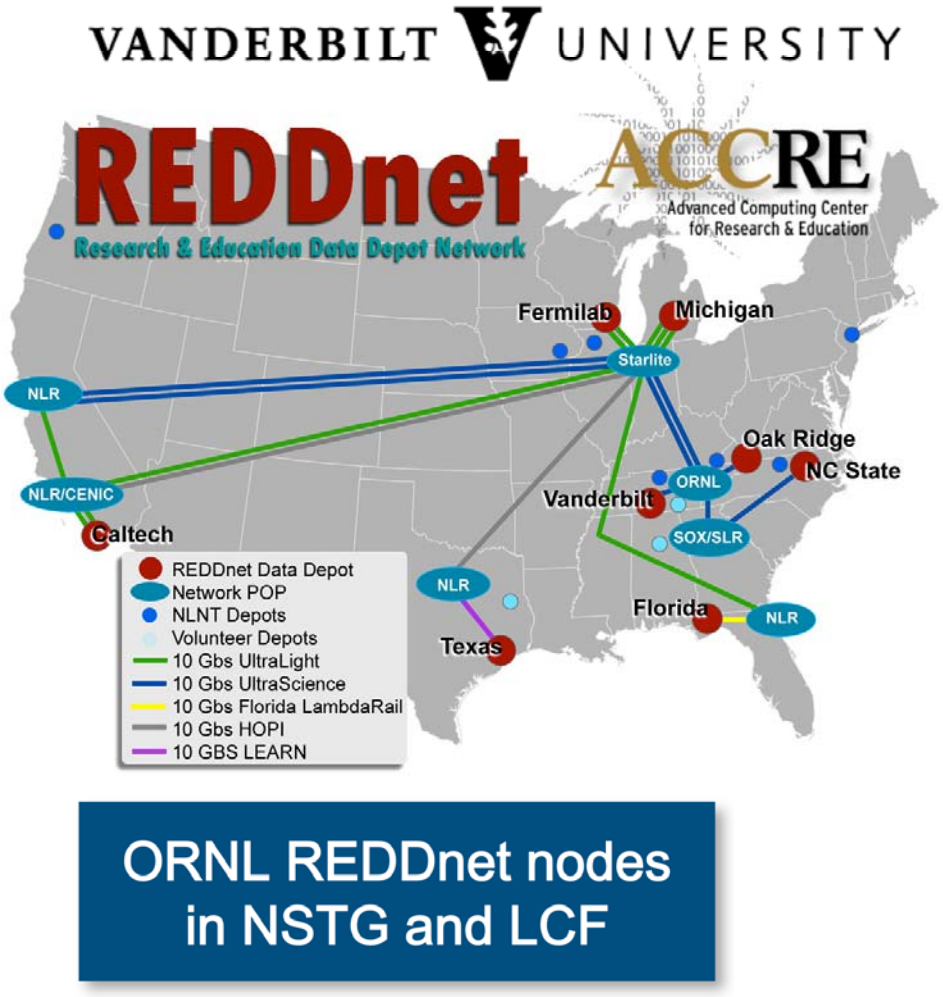

ORNL REDDnet nodes in NSTG and LCF

- National-scale logistical storage network
- Distributed (but acts local!)
- Funding from NSF and Library of Congress
- >700 TB (500 disk, 200 tape)
- Multiple application domains
  - Satellite imagery (AmericaView)
  - High energy physics: LHC
  - Terascale Supernova Initiative
  - Structural biology
  - Vanderbilt TV news archive
  - National Geospatial Data Archive
- Data everywhere (under the cover tools to manage movement, replication, …)

See Vanderbilt research booth for SC'07 demo and more details

http://www.reddnet.org/

Open Science Grid (OSG)

- Similar to TeraGrid in many ways:
  - Distributed resources
  - Includes Globus toolkit in software stack
  - Science focus
- Significant operational differences and culture from TeraGrid
- NSTG cluster is “dual deployed”: Both a TeraGrid and an OSG resource
  - Supports production jobs on both grids
  - Reports usage accordingly
  - Serves TG roaming accounts and most OSG Vos
  - Pick TG or OSG environment using “@” macros
- Pursued at request of neutron science community
- Significant overlap with REDDnet partners
- Purdue TeraGrid partner: Another TG/OSG integration case

http://www.opensciencegrid.org/

OAK RIDGE National Laboratory

Earth Systems Grid (ESG)

- TeraGrid transport for backend transport (ORNL ⇔ NCAR)
- ESG host in NSTG TeraGrid is at ORNL
- ESG host talks to HPSS at ORNL and launches GridFTP transfers for ESG
- See other TeraGrid partners: NCAR

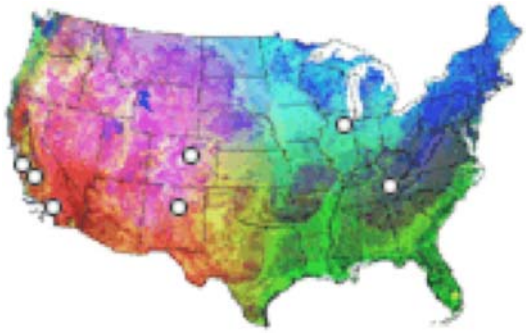

www.earthsystemgrid.org/

Acknowledgements and thanks

Staff and Effort

- NSTG: M. Chen[1], J. Cobb[1], V. Hazlewood[1], S. Hicks[1], G. Hinkel[1], V. Lynch[1], P. Newman[1], J. Nichols[1], D. Pack[1], G. Pike[1], Jim Rome[1], Bill Wing[1]
- SNS: J. Bilheux[1], G. Granroth[1], M. Hagen[1], J. Kohl[1], D. Mikkelson[9], R. Mikkelson[9], S. Miller[1], P. Peterson[1], S. Ren[1], M. Reuter[1], J. Schwidder[1], B. Smith[8], T. Swain[8], J. Trater[1], S. Vazhkudai[1]
- McStas group: E. Farhi[3], K. Lefmann[5], P. Willendrup[5]
- TechX Corp.: M. Green[6]
- OSG: R. Pourdes[2]
- ESG: D. Bernholdt[1], D. Middleton[4]
- UK eScience: L. Sastry[7]
- TeraGrid partners and others

Support

- National Science Foundation
- Department of Energy
- ORNL Laboratory Director's research and development fund

Organization

[1]ORNL, [2]Fermilab, [3]ILL, [4]NCAR, [5]RISO, [6]TechX Corp., [7]UK eScience, [8]UT, [9]U Wisc., Stout

OAK RIDGE National Laboratory

Contact

John W. Cobb, Ph.D.

Computer Science and Mathematics
(865) 576-5439
cobbjw@ornl.gov