

Northwest Indiana Computational Grid

Thomas Hacker

Computer and Information Technology, Purdue University

OSG All-Hands Consortium Meeting

Renaissance Computing Institute (RENCI)

University of North Carolina

March 3-6, 2008

What is NWICG?

- The Northwest Indiana Computational Grid (NWICG) is a partnership of researchers and educators at three university campuses in northwestern Indiana: Purdue University-Calumet, Notre Dame University, and Purdue University-West Lafayette.
- Under NWICG, resources at all three campuses are available to each other, and all campuses coordinate and plan future infrastructure with this partnership in mind.



Development of NWICG

- IP-Grid (Indiana-Purdue Grid) (2004)
 - NSF funded grid built on partnership for ETF proposal
 - Developed high performance networking infrastructure
- Teragrid partnership
 - Indiana and Purdue collaborated on a successful joint Teragrid proposal
- Large grid infrastructure projects were in development or underway
 - University of Michigan Grid (M-Grid) (2002)
 - UK e-Science, European Data Grid
- Motivated development of a computational grid focused on Northwest Indiana
- Purdue University (West Lafayette, Calumet), Notre Dame University
- Garnered Federal support to create Northwest Indiana Computational Grid
- Goals of NWICG
 - Leverage mutual interests amongst the three campuses with national science and research initiatives
 - Provide cyberinfrastructure to support science and educational efforts
 - Contribute to regional economic development

NWICG History

- 2004
 - NWICG first formed as a DOE-sponsored project
- 2005
 - Initial workshops
 - Condor pools at West Lafayette and Notre Dame deployed and expanded
- 2006
 - Adoption of the OSG software stack
 - Opteron cluster and storage array at Notre Dame brought online
 - Networking between West Lafayette and Notre Dame upgraded
- 2007
 - High-memory SGI Altix in West Lafayette installed
 - Condor pool at Calumet brought online
 - Networking to Notre Dame and Calumet upgraded

Networking Upgrades

- Initially, Notre Dame and West Lafayette added a dedicated 100Mb/s MREN (Metropolitan Research & Education Network) link between the two campuses
- Later Notre Dame obtained a 1 Gb/s “I-Light” link via Indiana's state-operated fiber network. Benchmarks were done to ensure this link is performing as expected between Notre Dame and West Lafayette.
- Calumet has brought online another DS3 circuit as a stop-gap measure pending their own I-Light connection.

NWICG HPC Resources

- **Calumet**

- Condor pool of over 200 Pentium4 Windows lab machines, accepting flocking from West Lafayette's Condor gatekeeper
- Small parallel computing cluster deployed at Calumet for use in training users and testing locally before involving NWICG grid partners



- **Notre Dame**

- 144-core Opteron Sun Fire X2100 system
- 36 TB NetApp FAS3050 storage array
- Condor pool containing more than 500 machines
 - Predominantly 32-bit and 64-bit Linux machines with some Sun machines

NWICG HPC Resources

- **West Lafayette**

- 128-processor Itanium SGI Altix 4700 with 512 GB of RAM (unique high-memory system to complement Notre Dame Opteron cluster)
- 30 TB SGI InfiniteStorage 4000 system
- 8-core FPGA SGI 450 RASC RC100 system (4 Xilinx Virtrex FPGAs)
- Tesla Quad-GPU system (CUDA) with dual 16-lane PCI buses
- Condor pool with over 7,000 processors of various type



- **Also smaller dedicated systems at each campus to support other resources**

Workflow Improvements

- Added Condor job submission hosts at each campus
 - Users won't need to configure local Globus/Condor-G clients.
- Dedicated interactive compilation node for the SGI Altix
 - Allows researchers to easily and quickly develop and test codes
 - Jobs submitted through Globus/Condor-G to SGI Altix
- Dedicated Serialize-Place-and-Route node being tested now at West Lafayette for grid users of the Altix 450 FPGAs
 - Allows researchers to create their own algorithms for the FPGAs and install/run them easily.

Organization / Communication

- Calumet and Notre Dame have dedicated AccessGrid facilities
 - AG used for regular coordination meetings
- NWICG website
 - Features a wiki for documenting and explaining how users can access NWICG resources.

Education / Outreach

- Condor workshop held in 2007 at Purdue
 - Introduced potential users and systems administrators to Condor
 - Demonstrations on using and configuring Condor
- Parallel computing workshop held at Calumet
 - Train students and faculty in parallel and distributed computing concepts and methods.
- Contributed funds to help build a teaching lab for high-performance computing at Purdue

Scientific Research using NWICG

- September 11 World Trade Center attack simulations using LS-DYNA

<http://www.youtube.com/watch?v=cddlgb1nGJ8>

- Purdue FPGAs confirmed to meet published BLAST benchmarks from SGI of 65.9x speedup over a normal CPU.





Image shows façade before impact. Reproduced courtesy of Mete A. Sozen, Ayhan Irfanoglu, Christoph M. Hoffmann.

- Ozone degradation research, including a very large 6-week job not able to be accommodated at national centers. This is the largest calculation ever completed for this problem.

Scientific Research using NWICG



*Image of virus structure.
Reproduced courtesy of Wen Jiang.*

- Purdue work to identify specific infectious virus structure using image reconstruction from electron cryomicroscopy. This work was published in Nature in February.

- Notre Dame research modeling of optimum automotive structures subjected to complex crash interactions using hybrid cellular automaton.

Scientific Research using NWICG

- Research on adsorption and reaction thermodynamics relevant to NO oxidation on Pt(111) at as a function of oxygen coverage (Notre Dame).
- Large prime number factorization research support (Purdue).
- Hurricane simulations from the 2000-2005 seasons as part of the Florida State University Mesoscale Superensemble Project (FMSP) (Notre Dame)




Image shows a 3-dimensional volume rendering of central pressure of Hurricane Katrina. The red vectors depict flow around and within the hurricane. Imagery produced by VAPOR (www.vapor.ucar.edu), a product of the National Center for Atmospheric Research.

NWICG Current Status

- Actively engaging users of major existing computational resources (Notre Dame cluster, West Lafayette high-memory Altix).
- In partnership with Teragrid, developing better processes for using FPGAs in a shared production setting, including issues related to bit-file generation, end executable compilation, and scheduling and execution on and of both host and FPGAs
- Setting up visualization facilities at Calumet now and actively recruiting a Calumet visualization center director.
- Awaiting major upgrade to Calumet networking via I-Light connection.
- Holding regular coordination and planning meetings for both infrastructure and steering committees, usually via AccessGrid on each campus.

Future Plans: Calumet

- Investigating curriculum development around simulation and modeling to be used in conjunction with existing programs.
- Support for much larger job/data transfer via upgraded network capacity.
- Development of visualization center and facilities for use by local users in conjunction with remote processing via NWICG partners.
- Investigating extension of Molecular Dynamics n-body techniques to Astrophysics problems using accelerator technologies.
- Exploring other research proposals as well.

Future Plans: Notre Dame

- Research into “Grid Heating”—the reuse of HPC waste heat as a heating-energy source.
- Work with FSU on satellite remote sensing techniques to improve and verify climate models.
- Deploying OSG Storage Element resource.
- Exploring other research proposals as well.

Future Plans: West Lafayette

- Open and refine access to FPGA resources already on the floor.
- Analyze power consumption of FPGA resources and determine power-time savings for certain jobs over normal CPU usage.
- Investigation of additional accelerator technologies for possible deployment, including GPUs and Cell processors.
- Development of image-processing service to use Tesla GPU (CUDA) system.
- Testing DVI over IP use as part of real-time remote visualization with NWICG partner institutions.
- Interoperate with TeraDRE to allow Condor-flocking of rendering jobs to other NWICG sites, and make TeraDRE resources available to NWICG users.
- Expanding cooperative Condor use to additional Indiana colleges via the BoilerGrid project.
- Exploring other research proposals as well.

Research Collaborators

- Ed Bensman
- Gary Bertoline
- David Braun
- Paul Brenner
- Jeffrey Evans
- Rachel Getman
- James Goldman
- Thomas Hacker
- Christoph Hoffmann
- Ayhan Irfanoglu
- Wen Jiang
- Bruce Loftis
- Mónica Martínez-Avilés
- Neal Patel
- Voicu Popescu
- William Schneider
- Shawn Slavin
- Carol X. Song
- Mete A. Sozen
- Doug Thain
- Samuel S. Wagstaff, Jr.

NWICG Personnel

- **Purdue-West Lafayette**
 - Gerry McCartney
 - Bruce Loftis
 - Chris Hoffmann
 - Preston Smith
 - Kevin Colby
 - David Braun
 - David Matthews
 - Birgitte Brydso
 - Carol X. Song
 - Randy Herban
 - Rick Irvine
 - Jeff Schwab
 - Mohamed Sayeed
 - Faisal Saied
 - Kay Hunt
 - Joma Roe
- **Notre Dame**
 - Peter Kogge
 - Olaf Wiest
 - Dewitt Latimer
 - Ed Bensman
 - In-Saeng Suh
 - James Rogers
 - Jean-Cristophe Ducom
 - Marcy Hull
 - Doug Thain
- **Purdue-Calumet**
 - Doug Sharp
 - Joe Morales
 - Shawn Slavin
 - John Alberts