J3.5 A REPORT ON PLANS FOR A TERAGRID

Robert Wilhelmson* and Dan Reed
National Center for Supercomputing Applications (NCSA)
University of Illinois at Urbana-Champaign

1. INTRODUCTION

The TeraGrid[1] is a plan to advance the cyberinfrastructure for 21st century science and engineering. It is a response to the pressing need for greater computational power to enable experimentation, modeling, data analysis, and visualization activities that often involve large volumes of data and use of an expanding national computational grid. It is a vision to develop and deploy a comprehensive computational, data management, and networking infrastructure of unprecedented scale and capability that couples distributed scientific instruments, terascale and petascale computing facilities, multiple petabyte data archives, and gigabit and beyond networks, all widely accessible by researchers including those in the atmospheric sciences, oceanography, and hydrology.

The TeraGrid is an outgrowth of the NSF PACI[2] program that has two leading edge sites, NPACI in San Diego, CA and NCSA[3] in Champaign/Urbana, IL. Facilities and resources from these and affiliate sites have been used by a growing number of geoscientists during the last decade. Over the past several years there has been increased emphasis within PACI on use of the national grid[4] in support of distributed computation, distributed data access and processing, distributed datamining, and distributed collaboration. Infrastructure developments in support of this emphasis have included Globus[5] middleware for authentication and security and the Access Grid[6] for collaborative group meetings and distributed presentations. The Teragrid is a natural extension of these and other developments. It will provide the geoscience community (including atmospheric science, oceanography, and hydrology) an infrastructure for actively and daily pursuing new research problems where location of resources being used and of people participating does not inhibit progress.

*Corresponding author address: Robert B. Wilhelmson, National Center for Supercomputing Applications, U. of Illinois, 605 E. Springfield Ave., Champaign, IL 61820

[1] http://archive.ncsa.uiuc.edu/MEDIA/vidlib/new.html (streaming video introduction)

[2] http://www.interact.nsf.gov/cise/descriptions.nsf/pd/paci?openDocument

[3] http://www.ncsa.uiuc.edu, http://www.ncsa.uiuc.edu/Science/Stories/, http://www.npaci.edu/online/ , http://www.paci.org

[4] http://www.gridforum.org/

[5] http://www.globus.org, http://www.ncsa.uiuc.edu/News/Access/Releases/010924.Middleware.html

[6] http://www.accessgrid.org

2. THE TERAGRID

The Teragrid plans call for provision to the national research community of multiple teraflop computing capabilities in the form of Linux clusters that support multi-teraflop model simulation, petabyte data analysis capabilities, high-performance visualization on remote platforms, and high resolution wall displays. Further, the software infrastructure necessary to support research productivity in this environment will be enhanced. The foundational support for the TeraGrid comes from a recently funded NSF proposal entitled “The TeraGrid: Cyberinfrastructure of 21st Century Science and Engineering.”

The TeraGrid facility will be deployed among four major institutions, NCSA, SDSC (San Diego Supercomputing Center), Argonne National Labs and Caltech. In partnership with IBM, Intel, Qwest, Oracle and SUN, the facility will consist of terascale Linux clusters that use Intel’s next-generation McKinley microprocessor, large-scale storage archives and data management software and a 40 Gb/s optical network to interconnect the components. This facility will increase PACI computing, storage and communication capability by an order of magnitude. The aggregate computing power will be 13.5 teraflops with roughly 600 terabytes of disk storage. Tertiary storage will be in the multiple petabytes. Operating as a distributed resource that can be co-allocated and scheduled, the TeraGrid will support traditional scientific computing and emerging disciplinary scientific groups. A unified TeraGrid Operations Center will coordinate management, user support, and access.

Fig. 1 shows the proposed TeraGrid layout. It will be augmented by facilities at the various sites. For example, at NCSA there will be more than 10 teraflops of computing capability including the current teraflop IA32 (Platinum) and IA 64 Linux clusters (Titan) along with NUMA SMP capability. By the end of 2002 there will be 240 terabytes of secondary disk storage and nearly 1.5 petabytes of archive storage. A large display wall (see 12.8 in this preprint volume) and Access Grid capabilities are already present.

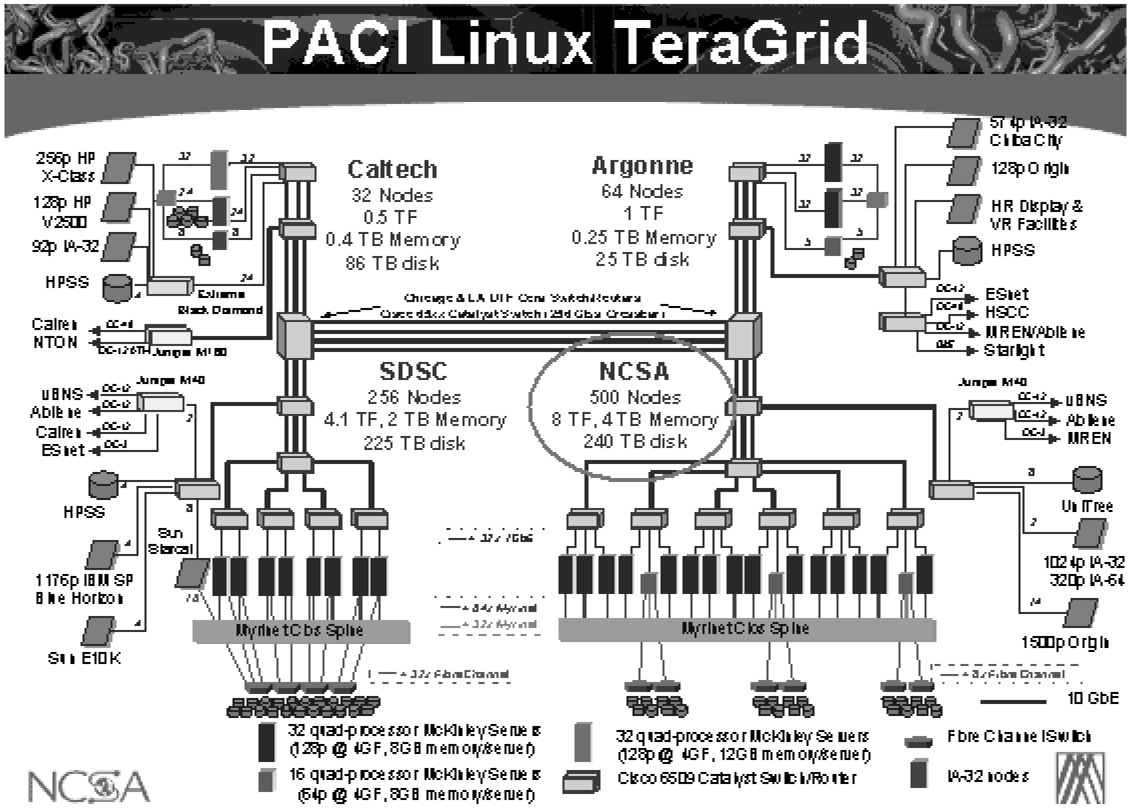

Fig. 1. This figure shows the various components of the TeraGrid as currently envisioned. The network between the sites will be 40 Gbits per second, over 10 times the bandwidth of the Abilene network (OC-48, 2.5 Gbits per second). The aggregate computing power will be 13.5 teraflops with roughly 600 terabytes of disk storage. Tertiary storage will be in the multiple petabytes.

3. 21st CENTURY TERASCALE APPLICATIONS

There is an unending need for computing power in the geosiences directed at modeling efforts that include data assimilation, model simulation, data analysis, and visualization. Equally important are the power needs for data collection, processing, and visualization. With the increase in power and data handling capabilities available across national and international networks that is achieved through use of distributed architectures (e.g. Linux clusters, distributed data repositories), researchers in all fields including the geosciences are becoming more dependent on the development of community codes and data handling tools that can fully tap into this infrastructure. For example, the developers and users of emerging terascale application codes face many challenges including the efficient usage (both computing and I/O) of parallel computing systems (and collections of these systems). Other users are creating or mining multi-terabyte data sets data from multiple sources to extract scientific insights. Still others are coupling scientific instruments (e.g., Doppler radar, radio telescopes, electron microscopes, beam lines) for near real-time data processing and analysis. Some applications embody all three aspects – these Grid-enabled applications will couple geographically separated computing, storage systems, and instruments to achieve breakthroughs. For example,the LHC (Large Hadron Collider) experiments will require filtering of petabytes of data to detect "one in a trillion " events. Such codes embody common usage models for high-end applications and exemplify "community toolkits" used by many more researchers than just the developers. The PACI partnerships plan to work with the developers of these community codes and domain toolkits to optimize their performance for the distributed terascale facility and its TeraGrid, immediately broadening the research base for terascale infrastructure.

The Weather Research and Forecasting Code (WRF)[7] is an example of one geoscience code being ported to the TeraGrid including enhancements for use with emerging portal, visualization, and

[7] http://www.wrf-model.org

datamining software. The WRF model is currently being ported to the new NCSA Itanium cluster using IA 64 technology from Intel and currently is running on Pentium III Linux cluster at NCSA. HDF5 parallel I/O functionality will be implemented in the future. At the November 2001 Supercomputing Conference in Denver, TeraGrid functionality will be demonstrated with the help of visitors to the NCSA booth. These visitors will collectively launch up to two hundred WRF simulations on grid-based computers. These simulations will be used to help study how two merging storms influenced tornado development as observed on April 19th, 1996 in Illinois. These idealized simulations will involve the initialization of two storms within a numerical model with the location, intensity, and appearance of the second storm determined by the visitor. Visitor input will be handled through a web browser and XCAT software will be used to launch the simulation on one of several grid computers. Results will be collected, visualized, and analyzed for incorporation into a data mining learning environment (D2K)[8]. The latter will be used to aid the visitor in choice of parameters and updated after every simulation. The results will be simultaneously displayed on a tiled wall[9] along with those initiated by others. Tiled wall display of WRF ensemble simulations is being implemented by using each wall panel (there will be 20 in all arranged 4 across and 5 down) containing a different view of a simulation (e.g., horizontal slice, three dimensional display). Animations of these different views will be available in lock step mode.

Other geoscience efforts that will be utilizing the TeraGrid are focusing on regional ocean and biological behaviors, coupling of storm scale watershed models (ARPS and EMGIS, WRF and JANSWERS)[10], and earthquake engineering through “teleobservation” and “teleoperation.” For example, the ROMS model being used in the LEO-15 observational/modeling project[11] is being ported to the Linux TeraGrid environment. It will be linked to a Spatial Modeling Environment (SME) for incorporation of biological model components. In the future (a few components are already operational), one can imagine an expanded set of resources including remote, biodegradable sensors in the ocean that monitor temperature, biological materials, and key chemical concentrations, transmitting the measurements via wireless technology to digital libraries, mining and visualizing this data directly to derive new insights, using the refined data in large scale predictive models, redeploying the sensors to refine the system as a result of the predictions, and finally, triggering nanoactuators to remove inappropriate concentrations of effluent or other non-native materials. All resources of this system, from the sensors to the networks, blur the distinctions among computation and data, hardware and software, machine and environment.

Another example is an earthquake engineering system that integrates “teleobservation ”and “teleoperation,” enabling researchers to control experimental tools —seismographs, cameras, or robots at remote sites —along with real-time remote access to data generated by those tools. Coupled with video and audio feeds, large-scale computing facilities for integrated simulation, data archives, high-performance networks, and structural models, researchers will be able to improve the seismic design of buildings, bridges, utilities, and other infrastructure in the United States. Like all multidisciplinary research and infrastructure, this earthquake engineering system combines research from multiple disciplines, including computing research in distributed systems and collaboration, structural engineering analysis of seismic activity on buildings and bridges, tsunami research on the formation and effects of tsunamis; and geotechnical engineering of the interactions of seismic activity with subsurface soil and foundations. The recently funded NEESGrid[12] is an early prototype of this environment.

The deployment of these application infrastructures and their associated grid components is encompassed in the X-In-a-Box strategy[13]. The initial In-a-Box initiative involves four interrelated software packages: commodity computational cluster software, grid toolkits, Access Grid collaboration systems, scalable tiled displays. Applications-In-A-Box will utilize these packages.

4. ACKNOWLEDGEMENTS

Text from the TeraGrid proposal (PI: Dan Reed) is edited and used in this preprint. Thanks to all those who contributed. This preprint can be found at http://redrock.ncsa.uiuc.edu/AOS/home.html under publications.

---

[8] D2K: Data To Knowledge is a next generation workspace architecture for the creation of data analysis applications in any domain area. It is a visual programming environment that allows users to easily connect software modules together in a unique data flow environment to form an application.

[9] The NCSA tiled wall and its use is described in preprint 12.8 in this volume.

[10] http://www.ncsa.uiuc.edu/Science/EnHydro/

[11] http://www.imcs.rutgers.edu/mrs/, http://marine.rutgers.edu/cool)

[12] http://www.ncsa.uiuc.edu/News/Access/Releases/010827.NEESgrid.html, http://www.eng.nsf.gov/nees/default.htm

[13] http://www.ncsa.uiuc.edu/News/Access/Stories/IAB/