A TeraGrid MATLAB Cluster - Exploring New Services for an XD Future

Cornell University
Center for Advanced Computing

PI David Lifka, Director, Cornell University Center for Advanced Computing
Co-PI: Gerhard Klimeck, Professor, Electrical and Computer Engineering,
Technical Director of the NSF Network for Computational Nanotechnology, Purdue University
Co-PI: Michael McLennan, Rosen Center for Advanced Computing, Purdue University

PURDUE UNIVERSITY

NSF Strategic Technologies in Cyberinfrastructure Award – 0844032
"A TeraGrid MATLAB Cluster – Exploring New Services for an XD Future"

Introduction

The increasing pervasiveness of distributed and parallel computing in a wide variety of fields has resulted in the need for computational power that scales seamlessly from the desktop to national resources without the steep learning curve required of traditional parallel programming techniques. Productivity in the HPC community is often measured by the performance gains seen from the parallelization of code. Instead, we propose that productivity can be measured by improved ease of access and utility. We have implemented a system that allows a researcher to scale from their desktop to a remote system transparently and with little knowledge of parallel computing. The Cornell team has deployed an experimental TeraGrid resource running MATLAB as an on demand computing utility. The broad usage of MATLAB in a wide spectrum of academic environments makes it an ideal application for this initial experiment. The MATLAB Distributed Computing Server running on this resource provides parallel and distributed computational services to interactive desktop users and Science Gateways. We have developed a client library that provides a safe and efficient means to ship computation from a user's workstation to our servers, where the computation can be run in a distributed or parallel style.

Project Homepage: http://www.cac.cornell.edu/matlab

Education & Training

Due to the pervasiveness of MATLAB in a very broad spectrum of academic environments, including small colleges and minority serving institutions, this experimental TeraGrid resource will be accessible to many new user groups beyond the traditional high-performance computing community. Parallel MATLAB offers support for both MPI parallelism and distributed computing parallelism. This will allow students that are new to parallel programming to focus on the opportunities for parallelism in their computational research without requiring them to start with extensive training in C, C++ or Fortran. The MATLAB on the TeraGrid project team will educate current and next-generation scientists on how to use the MATLAB on the TeraGrid cluster as an extension of their desktop and how to integrate it with Science Gateways. Classroom workshops will be held and Virtual Workshops will be available online. The interface, parallel computing learning modules, and MATLAB's user-friendly utilities will greatly soften the learning curve for parallel computing.

System Specifications

Microsoft Windows HPC Server 2008 Cluster
- Supports MATLAB Clients on Windows, Mac & Linux

64 Dell PowerEdge™ M600 blade servers
- 2 quad-core Intel® Xeon®/blade server
- 512 total cores
- 16GB RAM/blade server
- 1TB total RAM

8TB DataDirect Networks Storage
- RAID 6 file system with on-the-fly read/write error correction
- Accessible by all blade servers
- Accessible by GridFTP
- TeraGrid 10Gb connectivity

The MathWorks
- MATLAB Distributed Computing Server 4.1
- MATLAB Release 2009a
- MATLAB Parallel Computing Toolbox 4.1

System Architecture

The integration was accomplished by providing a Java-based integration with the MATLAB Parallel Computing Toolbox. The system hosted at the Cornell Center for Advanced Computing (CAC) is composed of a Dell blade cluster running Microsoft Windows HPC Server 2008 and an 8TB DataDirect Networks storage device which is mounted by a GridFTP server. Authentication is performed using x509 certificates retrieved from a MyProxy Server. A software package developed by CAC must be added to the client MATLAB classpath and provides access to the GridFTP server and to web services which perform job submission and maintenance tasks to the Microsoft Windows HPC Server 2008. With this integration a user can rapidly switch from running a multi-core process on their local workstation to multi-core and multi-node on our server.

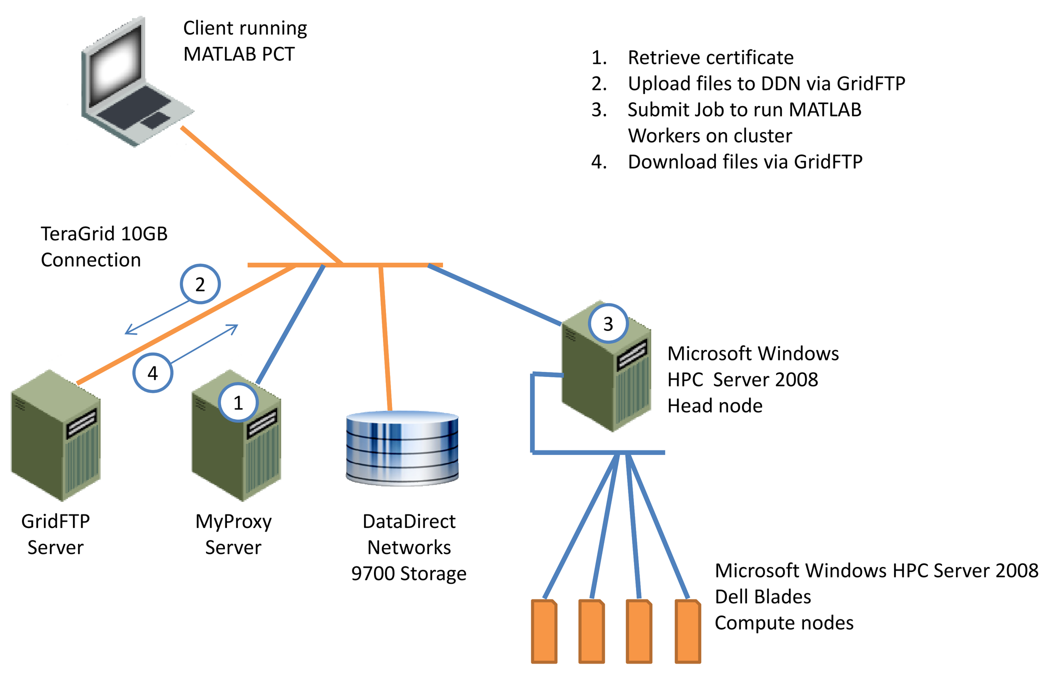

Client Architecture

The client is composed of a set of MATLAB functions callable by the user that integrate with the PCT through the generic scheduler interface and a Java component which handles the actual communication with MyProxy servers, GridFTP servers, and the Microsoft HPC Server 2008 Basic Profile service. The MATLAB platform allows native calls to Java objects, which makes the MATLAB functions little more than argument preparation for the Java component. We chose to use Apache CXF for web service operations and JGlobus for core MyProxy and GridFTP operations. This base functionality was then wrapped to provide for runtime-configurable transport mechanisms to allow connection reuse as well as advanced features like managing a connection pool, asynchronous transfers across the pool, reliable connections, pre-transfer compression, and dynamic selection of Extended Mode features like parallelism for connection pool members. Web service operations are secured via 2-way SSL and also allow WS-Security mechanisms for client authentication.

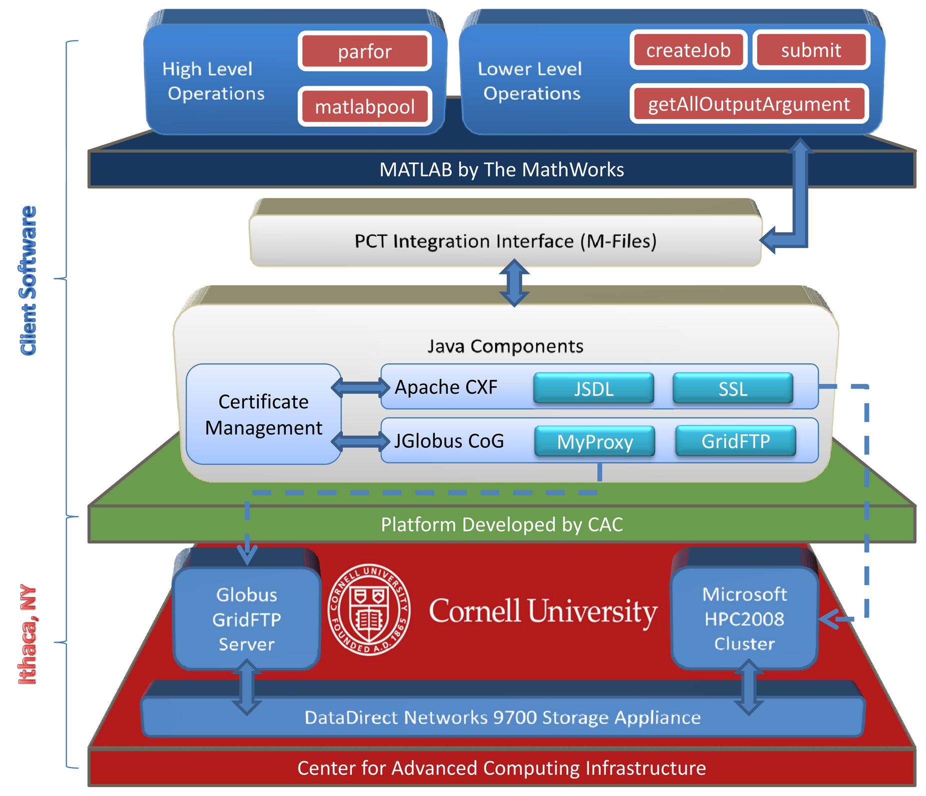

Build Your Own

This resource allows access to a parallel MATLAB resource in a secure manner nationally (and internationally) allowing a user to seamlessly scale from running a multi-core problem on their machine to running on multiple multi-core nodes. No special permissions are required for the client (i.e., there is no submission host), which is enabled by our SOA approach to this problem and by leveraging existing tools. The client works on 32 and 64 bit versions of all MATLAB-supported platforms (Mac, Linux, Windows, and Solaris) without needing to maintain separate releases and the server-side configuration is composed of commodity parts that can be easily duplicated. To this end, not only will the Cornell CAC support TeraGrid customers with this resource, CAC will be releasing documentation to aid other groups in setting up similar resources. Please visit the project homepage for more details: http://www.cac.cornell.edu/matlab

MATLAB Enables Broad Range of Researchers

The MathWorks produces software for technical computing and Model-Based Design for engineers, scientists, mathematicians, and researchers. Their two core products are MATLAB, used for performing mathematical calculations, analyzing and visualizing data, and writing new software programs; and, Simulink, used for modeling and simulating complex dynamic systems such as a vehicle's automatic transmission system. MATLAB is used by over 1,000,000 active users and 5,000 college/universities worldwide.

Focus on Science

MATLAB allows researchers to perform computationally intensive work without having to be experts in computer science. The MATLAB Parallel Computing Toolbox allows researchers to focus on more computationally intensive problems without having to become parallel computing experts. The MATLAB Distributed Computing Server allows applications developed with the Parallel Computing Toolbox to seamlessly scale to hundreds of cores on high performance clusters right from the researchers desktop.

Centers for Disease Control and Prevention (CDC) HCV Modeling

MATLAB® was used to study networks of coordinated amino acid variation in Hepatitis C virus (HCV), a major cause of liver disease worldwide. Mapping of coordinated variations in the viral polyprotein has revealed a small collection of amino acid sites that significantly impacts Hepatitis viral evolution. Knowledge of these sites and their interactions may help devise novel molecular strategies for disrupting Hepatitis viral functions and may also be used to find new therapeutic targets for HCV. Statistical verification of HCV coordinated mutation networks requires generation of thousands of random amino acid alignments, a computationally intensive process that greatly benefits from parallelization. This resource enabled a great reduction in HCV simulation times. The figure below is an HCV network (Campo et al, 2008). A genomic network of HCV coordinated substitutions is shown, where a vertex is an aa site and a link between two nodes is a significant correlation between two sites. The position of each vertex depends on its k-shell value and on the k-shell value of its neighbours. A color code allows for the identification of k-shell values, while the vertex's size is proportional to its degree. The k-shell decomposition and visualization was performed with LaNet-vi (Alvarez-hamelin et al, 2006).

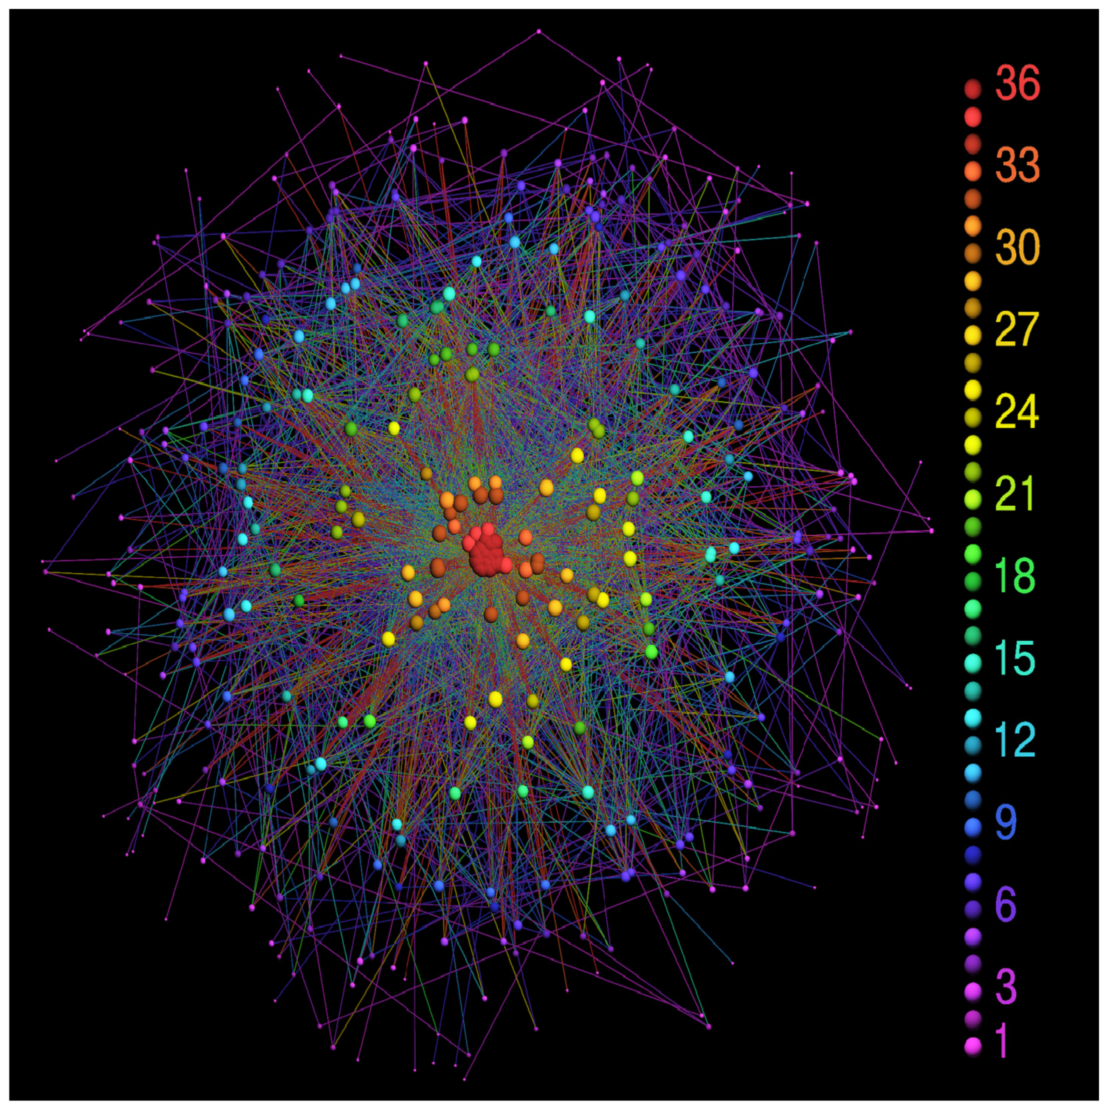

Science Gateways: Purdue University nanoHUB.org

TeraGrid Science Gateways such as nanoHUB are able to seamlessly run simulations over the network on the MATLAB on the TeraGrid experimental resource.

NanoNET is a tool to simulate the Nanobundle Network Thin Film Transistors (NB-TFTs). Random networks of carbon nanotubes with thousands of tubes and random orientation can be simulated using this tool. The final answer can be compactly formulated in the formula in the images below. Here ID is current and LC and LS is channel length and tube length of the transistor and m is the current exponent. For a normal Si MOSFET, m = 1 and the current is simply inversely proportional to channel length. But for these nanotube networks, m > 1 is also possible. Indeed, m = 1 for very high density networks but the value of m increases with decreasing tube density of the network. This abnormal behavior can be simply understood as follows: When the density of tubes is very high, most of the tubes take part in conduction and the current simply doubles on halving the channel length. But for a lower density network, there are some islands of pools of nanotubes that are not taking part in the conduction which start to connect as channel length is reduced. So not only the average path length reduces, but the number of paths also goes up with decreasing channel length which causes this super-linear increase in the current with channel length or m > 1. The smaller the density, the more pronounced is this effect and higher is the m. The following image depicts the default input screen. There are several parameters that the user must specify. The diagram defines the relationship between the various geometric input parameters. Additional inputs define the number of random samples to simulate and the computational resource requirement.

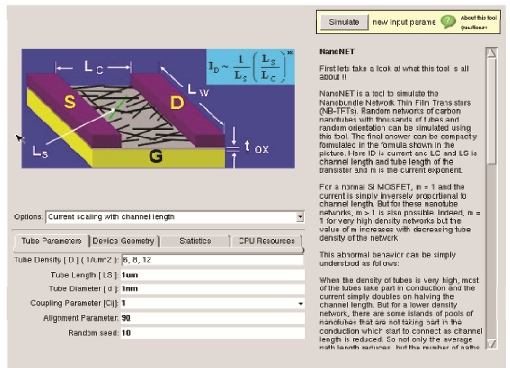

The next image depicts the results for a typical simulation.

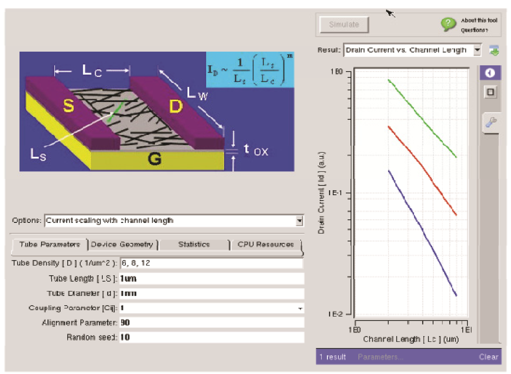

Multiple output images are available with the Drain Current vs. Channel Length being highlighted here. The following graph was obtained using the Rappture download feature.

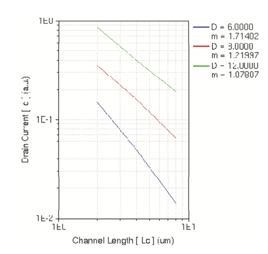

The following sequence of images shows the effect of tube alignment on current for the transistor. An alignment parameter of 0 indicates a completely aligned network, an alignment parameter of 90 indicates a completely random network. The simulation results show that the maximum current does not occur at the randomness extremes.

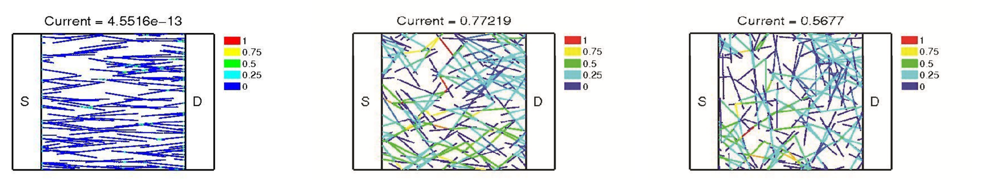

Earth and Atmospheric Science Modeling

This figure shows a snapshot from a numerical experiment studying diapirism in the mantle wedge of a subduction zone. The experiments are studying the physical and chemical processes that lead to arc volcanism above descending plates in tectonic settings like in the Cascades Volcanics (e.g. Mt. St. Helens, Lassen Peak) or Japan (Fuji, etc.). This experiment explores how metamorphic reactions in the descending slab release water that in turn leads to diapirism and partial melting in the overlying mantle wedge. This example was calculated on a 1e7 dof grid on a 32-processor cluster. Speed is color-contoured along an isosurface of a density contour, where the density reduction is due to the presence of extra water phases expelled from the slab. The use of parallel MATLAB made it possible to code a complete parallel-3D Stokes solver (incorporating multigrid-preconditioned and domain-decomposition-preconditioned conjugate gradient solvers), and a Semi-Lagrange transport solver including the effects of melting reactions in 1.5 years of man-work. (Joerg Hasenclever, from ongoing Ph.D. work at Cornell with Jason Phipps Morgan, 2009.)

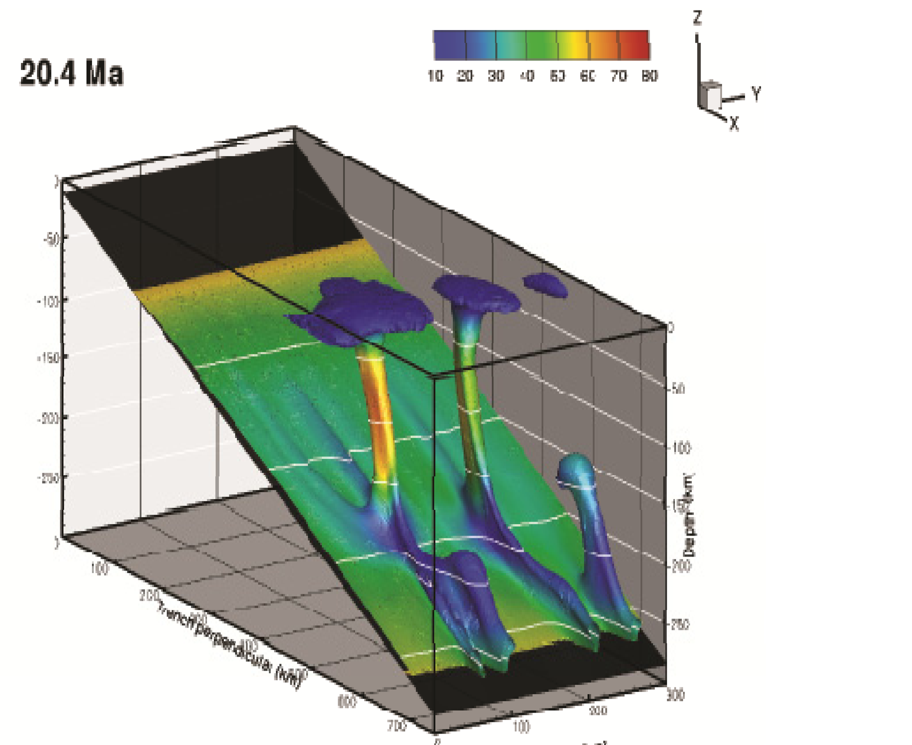

Acknowledgements
This research is sponsored by the National Science Foundation (Award #0844032), Dell, Intel, Microsoft, and The MathWorks. We would also like to thank our colleagues at the Centers for Disease Control for their efforts.

The MathWorks
intel
Microsoft
DELL
TeraGrid
Experimental Resource