## Research Computing resources at Georgia State University

Information Systems & Technology (IS&T) provides systems and applications support for research computing for Georgia State University research. This document is provided for researchers interested in an overview of research computing resources available to them, or needing "boiler plate" descriptions of these resources for preparing funding grants. We will be glad to help you with additional information tailored to your needs or grant.

IS&T's advanced computer services team supports the deployment of Georgia State University cyberinfrastructure for research computing, including high performance computing, grid computing, data storage and research networking.

Georgia State's core research computing resources include:

- IBM System p5 575 with POWER5+ processors (ursa.gsu.edu)
- IBM Cluster 1350 with Intel® Xeon® Quad-Core processors (octans.gsu.edu)
- SURAgrid

Georgia State is building a diverse and powerful supercomputing grid to support a growing body of academic and scientific research. Researchers include more than 100 faculty and Ph.D. students from many disciplines as well as from other universities such as Georgia Southern University and Southeastern Universities Research Association (SURA) sites such as Virginia Commonwealth University, Louisiana State University and University of Delaware.

Experiencing increasing demand for supercomputer resources has led to a university-wide strategy to build computational resources and services that support a wide variety of disciplines, with an especial focus the University's strategic areas of focus, such as Molecular Basis of Disease (http://www.biology.gsu.edu/mbd/) and Brains & Behavior (http://www.biology.gsu.edu/brains&behavior/).

As researchers in chemistry, biology, computer science, finance and other disciplines gain access to more capacity and capability, they naturally are expanding their research goals and thereby driving demand. The following provides an overview of research computing resources available to Georgia State University researchers.

*For more information on research computing resources, visit http://www.gsu.edu/ist/acs/ or contact:*

***Art Vandenberg***
*Information Systems & Technology*
*avandenberg@gsu.edu*
*404-413-4743*

**IBM System p5 (*ursa.gsu.edu*)** is a supercomputer for research computing. Acquired June 30, 2006 as an enterprise resource, it ran its first applications by September 2006. Initial applications include **Amber** for molecular dynamics simulation on bio-molecules and **Gaussian**, based on quantum mechanics, for prediction of energies, molecular structures, properties, and vibration frequencies of molecular systems. **SAS** is installed for business intelligence and statistical analysis.

**IBM System p5 P575 with Power5+ Processors**
576 CPU. 1408GB RAM. 18TB SAN storage

**IBM System p5 575 with Power5+ Processors**

- **Symmetric multi-processor supercomputer**
  - *36 nodes with 16 x 1.9 GHz Power5+ CPU*
  - *576 CPU total, 1408GB Total RAM*
    - *34 Nodes have 32G RAM (2GB per CPU)*
    - *1 Node has 256GB RAM (16GB per CPU)*
    - *1 Node has 64GB RAM (4GB per CPU)*
    - *(On a node one CPU may access up to all node RAM)*
  - *64bit configuration*
  - *High-speed federated switch node interconnect*
- ***1.3 TeraFLOPS estimated for 12 nodes***
  - *Linpack benchmark on a single p575 node (16 x 1.9 GHz Power5+ CPU) as published by IBM, is 108.79 GFLOPs.*
  - *8 nodes Linpack = 866.1 GFLOPs (99.51% theoretical max)*
- **18 TeraBytes of high speed storage**
  - *IBM DS4800 SAN Storage delivers up to 1,600 MBps throughput at up to 550,000 input/output per second*
  - *Two 4-Gb-per-second IBM 16-port SAN Switches*
  - *224 Fibre Channel (FC) or 224 Serial ATA (SATA) drives*

The IBM 575 runs **AIX 5.3** (IBM UNIX) operating system and IBM **GPFS** (General Parallel File System) for high performance shared file access and processing. **LoadLeveler** manages serial and parallel job scheduling; **XL/Fortran**, **XLC** compilers and **IBM Engineering and Scientific Subroutine Library** (ESSL) are available. **Data Backup/Recovery System** is comprised of IBM's Tivoli Storage Manager Server and IBM Ultrium 3 LTO Tape Library. Tivoli Storage Manager (TSM) automates data backup/restore**.**
More info: http://www.gsu.edu/ist/acs/ursa.html

**SURAgrid** – Georgia State's IBM System p5 is part of SURAgrid-IBM System p5 initiative, a partnership of Georgia State, Southeastern Universities Research Association and IBM. SURAgrid's IBM System p5s will boost SURAgrid compute capacity and Georgia State researchers may request access to these AIX compute cycles on SURAgrid. SURAgrid and IBM will link these systems via optical networking (cf. Southern Light Rail) creating aggregated System p5 cycles for researchers. More info: http://www.gsu.edu/ist/acs/24678.html

**Representative Applications**: Amber; Gaussian; CHARMM; AMMP; SAS.
**Representative Researchers**: Chemistry; Physics & Astronomy; Biology (molecular modeling; quantum mechanics; proteomics); Risk Management & Insurance (financial econometrics); Political Science (Monte Carlo simulation).

**IBM Cluster 1350 (*octans.gsu.edu*)** is the newest addition to research computing. It was acquired September 30, 2007, delivered and verified by November 30, 2007 and is being prepared for user access early 2008.

**IBM Cluster 1350 with Intel® Xeon® Quad-Core processors**

- **IBM BladeCenter® H technology**
  - ***HS21XM Blades** with **Intel® Xeon® Quad-Core 2.67GHz processors***
  - ***320 CPU total, 640GB Total RAM***
    - *2GB RAM per CPU*
  - ***64bit configuration***
  - ***3.0 TeraFLOPS***
- ***10 Gigabit Ethernet switch for compute and data network***
- **Management & User nodes**
  - *Each an IBM x3650*
  - *Total 1.8 usable Serial Attached SCSI hard disk drive*
- **Access to 18 TeraBytes of storage**

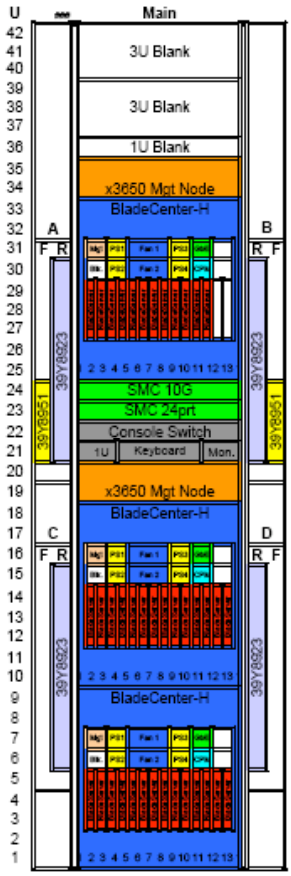


**IBM Cluster 1350 with Intel® Xeon® Quad-Core processors -** 320 CPU, 640 GB RAM, 3TeraFlop

The IBM Cluster 1350 runs **RedHat Enterprise Linux ES 4.0** operating system. IBM's xCAT provides cluster management tools and Platform Computing's LSF scheduler will be used for job submission and monitoring. IBM **GPFS** (General Parallel File System) will be used for high performance shared file access and processing. The 1350 will also have access to existing Tivoli Storage Manager and 18 Terabytes of storage that Georgia State has acquired in 2006 with the IBM System p5. This will allow researchers to easily share results from the IBM System p5 with the new IBM System Cluster 1350 – and vice versa.

Georgia State University's high performance computing users included than 40 faculty and Ph.D. students from many disciplines as well as from other Southeastern University Research Association sites such as Virginia Commonwealth University, Louisiana State University and University of Delaware. Adding the Cluster 1350 provides support for research that requires commodity-level Linux computing.

Researchers in chemistry, biology, computer science and political science will be among the new machine's first users. The Cluster 1350 along with Georgia State's high-end IBM System p5 575 supercomputer offer scientists a range of in-house computing capabilities.

**Example Applications**: Gaussian; MatLab; Dock; Auto-Dock
**Researchers**: Chemistry; Physics & Astronomy; Biology (molecular modeling; quantum mechanics; proteomics); Finance (financial simulation); Political Science (Monte Carlo simulation).

More info: http://www.gsu.edu/ist/acs/26220.html

**SURAgrid** is an outcome of an NSF Middleware Initiative (NMI) Integration Testbed Program in which Georgia State University took part. SURAGrid is progressing in three key areas:

- Grid community building - promoting grid technologies for research and education.
- Grid building – middleware infrastructure, portals, and inter-institutional authentication.
- Grid-enabling applications – identifying researcher applications that benefit from grids.

Georgia State's ACSrocks and IBM System p5 are part of the SURAgrid resources that currently include 16 clusters from 12 Southeastern University Research Association institutions, computing resources with 2,041 CPUs, 12.6 TeraFlops of computation and over 56 Terabytes of disk.

Georgia State's IBM System p5 (**ursa.gsu.edu)** is accessible to SURAgrid by way of the Globus gatekeeper **janus.gsu.edu**. Ursa.gsu.edu is one of many computing resources available to SURAgrid members, including Georgia State researchers, through the SURAgrid Portal.

For more information see: http://www.gsu.edu/ist/acs/24678.html

**United Devices, Grid MP** is a solution for harvesting idle CPU cycles from PCs and workstations in Georgia State University student labs and classrooms. Georgia State has licenses for up to 1,000 concurrent PCs. Grid MP offers excellent security for distributing work (fully encrypted data and program packages) with a web interface and portal for job submission, monitoring and retrieving results. United Devices is used by 5 of the 6 top pharmaceutical companies for research and development, clinical simulations (via R modules, SPSS modules, SAS modules – with caveats!), and data mining and processing.

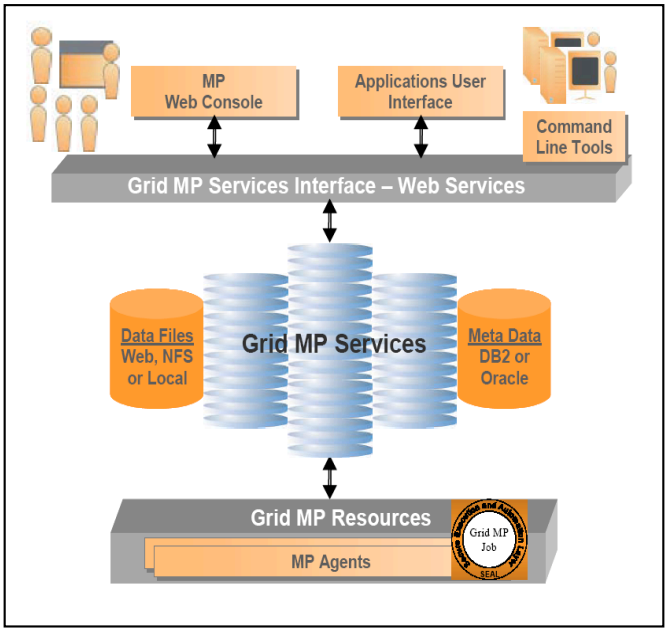


**James Reid**, PhD student **Dr. Ying Zhu**'s lab uses United Devices to support Blender (an open source software for **3D modeling, animation**, rendering, post-production, interactive creation, playback) with tremendously improved graphic rendering times since, rather than one PC rendering 250 frames, 250 PCs simultaneously render 1 frame each.

**David Cofer**, PhD student **Dr. Don Edward**'s lab is working on **AnimatLab** AnimatLab is an attempt at enabling reconstruction of relevant parts of the crayfish nervous system, and the world where the animal lives in an interactive computer simulation. Reconstruction of neural circuits and simulation of their activity is now a mainstay of computational neuroscience through such open source programs as NEURON and GENESIS. However, those simulations do not include the body and the world, and do not account for their dynamic interaction with the nervous system that occurs in real animals. AnimatLab provides a generic solution to this problem with a toolkit that allows the investigator to build computational models of the relevant neural circuits, the animal's body, and the world. AnimatLab is designed to be easily usable by investigators studying the neural control of dynamic behavior in a wide variety of organisms.

**Dr. Bob Calin-Jageman**, post-doc with **Dr. Paul Katz**'s lab, is working with **Bill Xie**, PhD student of **Dr. Yi Pan**, to implement a neuron application to model the interactions of the brain.

Citation: Robert J. Calin-Jageman, Chao Xie, Yi Pan, Art Vandenberg, Paul S. Katz, "**NEURONgrid: A Toolkit for Generating Parameter-Space Maps using NEURON in a Grid Environment**," Proceedings of the 2007 International Symposium on Bioinformatics Research and Applications, Lecture Notes in Computer Science, Springer, Atlanta, Georgia, May 7-10, 2007.

**Abstract**. Neuroscience research increasingly involves the exploration of computational models of neurons and neural networks. To ensure systematic model exploration, it is often desirable to conduct a parameter-space analysis in which the behavior of the model is catalogued over a very large range of parameter permutations. Here we report the development and testing of a tool-kit called NEURONgrid for conducting this type of analysis in a grid environment using NEURON (Hines & Carnevale, 1997, 2001), a popular and powerful simulation platform for the neurosciences. NEURONgrid provides helper classes within NEURON for manipulating parameters, a package of NEURON for running in a grid environment, and a management client that enables neuroscientists to submit a parameter-space analysis, monitor progress, and download results. NEURONgrid provides a user-friendly means for conducting intensive model exploration within the neurosciences. It is available for download at http://neurongrid.homeip.net.

**NOTE: Georgia State is Partner in IBM's World Community Grid** (WCG) (http://www.worldcommunitygrid.org/) has over 500,000 devices using United Devices agents (or BOINC http://boinc.berkeley.edu/) engaged in projects such as the Human Proteome Folding and FightAIDS@Home – all "projects that hold tremendous potential to benefit humanity." Georgia State is a World Community Grid Partner.

From more info, see: http://www.gsu.edu/ist/acs/24666.html