


Introduction to CCV

John Huffman
Mark Howison

User Services & Support

Staff

- ▶ Jan Hesthaven, Faculty Director
- ▶ John Spadaro, Interim Executive Director
- ▶ Sharon King, Manager



User Services & Support

- ▶ John Huffman
- ▶ Mark Howison
- ▶ Andy Loomis
- ▶ Lingsheng Dong (Biomed)

Systems

- ▶ Sam Fulcomer
- ▶ Dave Johnson
- ▶ Nate Wood

Mission Statement

- ▶ Support of educational and research activities in advanced computing and visualization for classic users and nontraditional users of such technology.
- ▶ We maintain publicly available parallel computing clusters , large scale storage facilities for data manipulation, backup facilities, network infrastructure, and visualization systems.
- ▶ We are a service provider for services related to HPC, data backup, etc, and offer technical advice with the purchase of advanced HPC and visualization equipment.
- ▶ These services are provided to support research and educational efforts in HPC only. Storage and back-up facilities are solely used for large-scale research related data.

Contact Info

▶ For more information:

- (401) 863-7557
- info@ccv.brown.edu

Hours

8:30 – 5:00 (Academic Year)
8:00 – 4:00 (Summer)

▶ For support:

- support@ccv.brown.edu

▶ Keeping in touch:

- System updates from Twitter: [ccv_brown](#)
- User mailing list: users@ccv.brown.edu
- Website: ccv.brown.edu

Locations

94 Waterman Street

- ▶ 2nd floor, above EF Bishop
- ▶ Offices and conference room



180 George Street

- ▶ Corner of Brook and George
- ▶ Machine room
- ▶ Visualization facilities and the CAVE

Data Services

▶ RData (Research Data Storage)

- All faculty receive a 256GB free allocation
- Additional storage can be purchased on an annual or long term basis
- High-performance, parallel file system
- Accessible via **CIFS mounts** from any PC or Mac

▶ Backups

- Tivoli Storage Manager (TSM) backup and archival server with tape library

▶ Dropbox (dropbox.brown.edu)

- Send files that are too large for email, up to 2GB

HPC Services

- ▶ **Oscar** cluster (IBM iDataPlex)
 - 204 nodes for running batch jobs
 - Each node has:
 - 2 x quad-core Intel Xeon 5540 (2.53 GHz) processor
 - 24GB DDR3 memory
- ▶ **smp** nodes (smp000 .. smp007)
 - Individual access for interactive and large-memory jobs – accessible by VNC
 - Each node has:
 - 4 x quad-core AMD Opteron (2.6 GHz) processor
 - 64 or 128 GB memory
- ▶ Both systems share a QDR Infiniband interconnect and a high-performance GPFS file system




Visualization

▶ The CAVE

- An immersive 3D display system in an 8' cube
- CrystalEyes LCD-shutter glasses for stereo depth-perception
- Intersense and Polhemus tracking devices
- ***Undergoing a complete redesign for Fall 2011***

- ▶ Custom solutions for individuals and labs (talk to us!)




▶ Tiled Display Wall

- high-resolution stereo display for data exploration, educational, and lab use
- testbed for new CAVE technologies

The Future of Brown Visualization

▶ The New CAVE

- Much larger space
- 160M pixels, 70 projectors
- Wireless Tracking



▶ Tiled Display Wall

- Single computer, easier to use (desktop applications)
- Higher resolution, larger screen

Condo Membership

- ▶ PI's can contribute grant funding for the purchase of additional hardware
- ▶ Owner has top priority access to that portion of the system for 5 years



Available Libraries & Packages

- ▶ MATLAB
- ▶ Mathematica
- ▶ Maple
- ▶ Python, numpy, scipy
- ▶ Perl
- ▶ R
- ▶ Java
- ▶ BLAST, mpiBLAST
- ▶ CPLEX
- ▶ ATLAS, GotoBLAS
- ▶ ScaLAPACK
- ▶ UMFPACK
- ▶ Boost
- ▶ GNU Scientific Library
- ▶ FFTW
- ▶ HDF5, netCDF
- ▶ ...

Current Users of CCV Services

- ▶ Creative Writing
- ▶ Geological Sciences
- ▶ Biomed (CCMB, MCB, EEB, Neural/BSP)
- ▶ Cognitive Science
- ▶ Classics
- ▶ Economics
- ▶ Mathematics
- ▶ Chemistry
- ▶ URI, MBL, others

Highlighted Users & Projects:



CaveWriting



- ▶ Taught each semester by John Cayley



Highlighted Users & Projects:

▶ ADVISER (Jim Head PI)

- Large image processing
- Multiple images
- High end storage
- High end visualization



Highlighted Users & Projects:

▶ Illumina Genome Analyzer

- Sequencer located at the Center for Genomics and Proteomics at Ship Street
- Data transferred to CCV (~1TB)
- Quality must be established while reagents are still viable (time sensitive)
- Data can be further analyzed at a later time
- Primary contact: Lingsheng Dong

<https://sites.google.com/a/brown.edu/bioinformatics-in-biomed/>

Staff Support Services

- ▶ Software installation and maintenance
- ▶ Application troubleshooting and performance tuning
- ▶ Hourly consulting rates for smaller projects
- ▶ Partial grant funded FTEs for larger projects
- ▶ Range of expertise including HPC, IO and data tuning, immersive and desktop visualization

(BREAK)

Tutorials

- ▶ Logging in
- ▶ Virtual Network Computing
- ▶ Mounting RData
- ▶ Software modules
- ▶ Running interactive and debug jobs
- ▶ Running production jobs

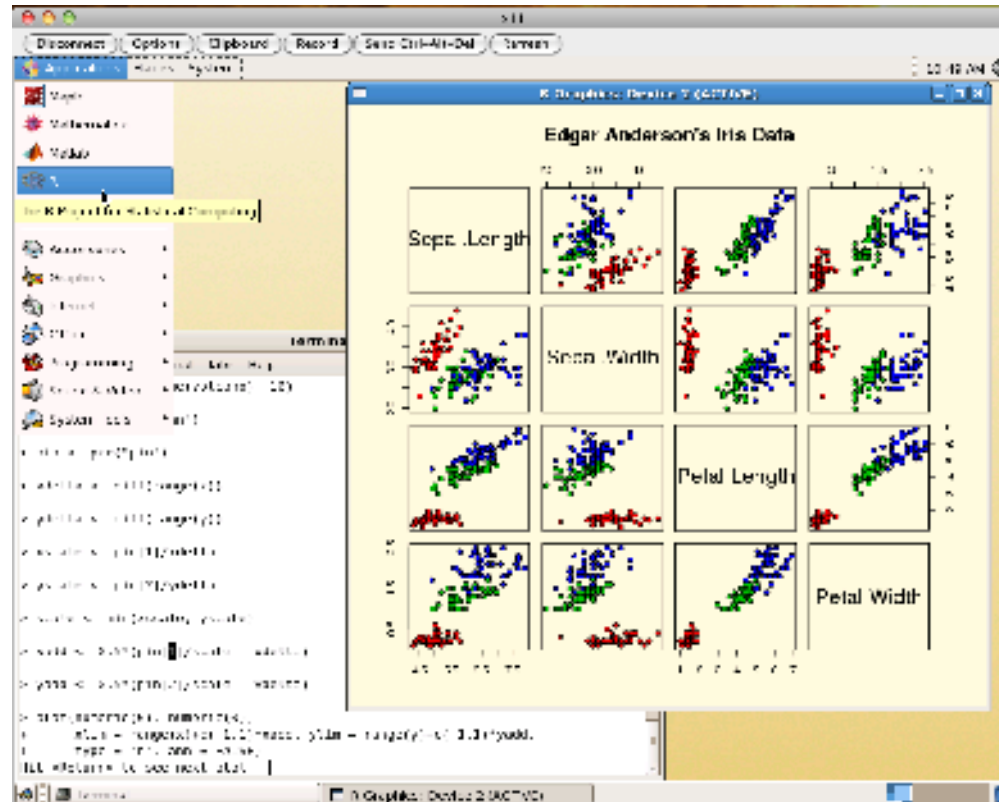
(See more at <http://www.brown.edu/CCV/doc>)

Logging in

- ▶ CCV uses the Secure Shell (SSH) protocol
- ▶ You will need an SSH client
 - Linux / OS X comes with a command-line client
 - Windows: try PuTTY or Cygwin
- ▶ Connect to the “ssh” server
 - Linux / OS X / Cygwin:
`ssh username@ssh.ccv.brown.edu`
 - PuTTY: enter `ssh.ccv.brown.edu` in “Host Name”

Virtual Network Computing

- ▶ OR connect through CCV's Virtual Network Computing client, available here:
 - <http://www.brown.edu/Departments/CCV/vnc>



Mounting RData with Samba/CIFS

- ▶ Allows you to mount remote file systems (i.e. **RData**) as if they were local drives
- ▶ Request a Samba account from **support@ccv**
- ▶ Create a Samba password by logging into Oscar and using **smbpasswd**
- ▶ On Windows, “Map Network Drive” with address **\\oscarcifs.ccv.brown.edu\username\data**
- ▶ On Mac, “Connect to Server” with address **http://oscarcifs.ccv.brown.edu/**
 - Select the “data” volume when prompted

Software modules

- ▶ Software on CCV systems is organized using the Modules system
- ▶ You can load and unload modules to control your software environment:

```
$ module load matlab/2008b
```

```
$ matlab
```

- ▶ View available modules with `module avail`
- ▶ View loaded modules with `module list`
- ▶ Module commands in your `~/.modules` file will automatically execute when you login

Running and debugging interactive jobs

- ▶ Log into one of the smp nodes (smp001–smp009)
 - From Oscar, `ssh smp001`
 - Execute serial/threaded program interactively
- ▶ smp001-smp006 have 64GB memory
- ▶ smp007 has 128GB
- ▶ smp008-009 have GPUs and 32 cores
- ▶ Resources (CPU, memory) are shared with other users who are logged onto the same node

Running production jobs

- ▶ Submitting a production job guarantees resources
- ▶ Specify requirements with a “batch script”

```
#!/bin/bash
#PBS -N jobname
#PBS -r n
#PBS -j oe
#PBS -l nodes=2:ppn8,walltime=00:30:00
```

```
# execute an MPI program
NPROCS=`wc -l < $PBS_NODEFILE`
mpirun -np $NPROCS mpirun <args>
```

```
# execute a binary in your path
program <args>
```

```
# change to the directory where the script was submitted
# and execute a local program
cd $PBS_O_WORKDIR
./myprogram <args>
```

- ▶ Submit the script with `qsub scriptname`

Managing jobs

- ▶ See the status of the entire queue with `showq`

JOBID	USERNAME	STATE	PROCS	REMAINING		STARTTIME
1170187	jfranck	Running	64	11:48:19	Mon Oct 18	10:23:40
1168909	mdeng	Running	8	18:47:20	Sun Oct 17	05:22:41
1165154	leyson	Running	8	19:05:14	Thu Oct 14	05:40:35
1165155	leyson	Running	8	19:08:12	Thu Oct 14	05:43:33
1168910	mdeng	Running	8	19:16:54	Sun Oct 17	05:52:15
1168917	mdeng	Running	8	20:30:33	Sun Oct 17	07:05:54

...

- ▶ See a summary of your jobs with `myq`
- ▶ See another users jobs with `myq username`
- ▶ Alter a job with `qalter <options> jobid`
- ▶ Delete a job with `qdel jobid`

(BREAK)

(HANDS-ON TUTORIAL)