

MIGRATING TO THE SHARED COMPUTING CLUSTER (SCC)

SCV Staff

Boston University

Scientific Computing and Visualization



Topics

- Glenn Bresnahan – Director, SCV
 - MGHPCC
 - Buy-in Program
- Kadin Tseng – HPC Programmer/Consultant
 - Code migration
- Katia Oleinik – Graphics Programmer/Consultant
 - GPU computing on Katana and SCC
- Don Johnson – Scientific Applications/Consultant
 - Software packages for E&E

Highlights of Shared Computing Cluster (SCC)

- At Massachusetts Green High Performance Computing Center, Holyoke
- Hardware
 - 1024 cores (64 Intel Xeon E5-2670 nodes with 16 cores each) -- more later !
 - 594 TB of usable disk space
 - No GPUs at present. Katana GPUs will be moved to SCC by July 1, 2013
- Software
 - Most system and application software packages have been installed
 - Webpage available to check status of packages (later slide)
- Operating system is Linux Centos 6.4
- Login nodes: *scc1.bu.edu*, *scc2.bu.edu*, *scc3.bu.edu*, *scc4.bu.edu*
- Network speed between main campus and MGHPCC is 10Gbit/s
- Who has access to the SCC? Users with Katana/Blue Gene access
- Home directory disk storage for each user is 10GB
- **All projects** get 50 GB each on /project and /projectnb to start.
- /project is backed up daily (just like on Katana)
- In addition, /project and /projectnb are backed up by *snapshots* — deleted files are *retrievable* by owner
- PI may ask for up to 1 TB storage for free; purchase more through buy-in

Where are my files on SCC and Katana ?

- For majority of users, */project* and */projectnb* migration is done !
 - If your */project* files not migrated yet, you will be contacted soon
- Migrated */project* and */projectnb* are shared by SCC and Katana
- Home directories (\$HOME) on SCC and Katana are distinct (no migration)
- In general, **don't copy dotfiles** (e.g., *.cshrc*, *.bash*) among systems, otherwise things may break. Adding contents to them is OK
- If you use *.aliases* on Katana, check before copying over to SCC
- On SCC, to access Katana home directory, prepend file path with **/katana**

```
scc1:~ % cd /katana/usr3/graduate/user
scc1:~ % cp myfile /katana/usr3/graduate/...../ .
scc1:~ % cp -r /katana/usr3/graduate/...../myfolder .
```

How to Keep \$HOME Files In Sync

To maintain a single copy of a file on both SCC and Katana home directories, you may accomplish it with:

- makefile

```
# source is on Katana, executable is local on SCC
```

```
BASE = /katana/usr3/graduate/username
```

```
mm_mpi_gnu: ${BASE}/matrix_multiply_mpi.f90
```

```
    mpif90 -o $@ ${BASE}/matrix_multiply_mpi.f90 -O3
```

```
# source on Katana, build exec for AMD & Intel & move back to Katana
```

```
mm_mpi_katana: ${BASE}matrix_multiply_acc.f90
```

```
    mpif90 -o $@ ${BASE}matrix_multiply_acc.f90 -O3 -tp x64
```

```
    mv $@ ${BASE}/
```

- Symbolic link abc.c in folder Misc on SCC points to a real file on Katana
scc1:~ % **ln -s** /katana/usr4/ugrad/name/abc.c Misc/abc.c

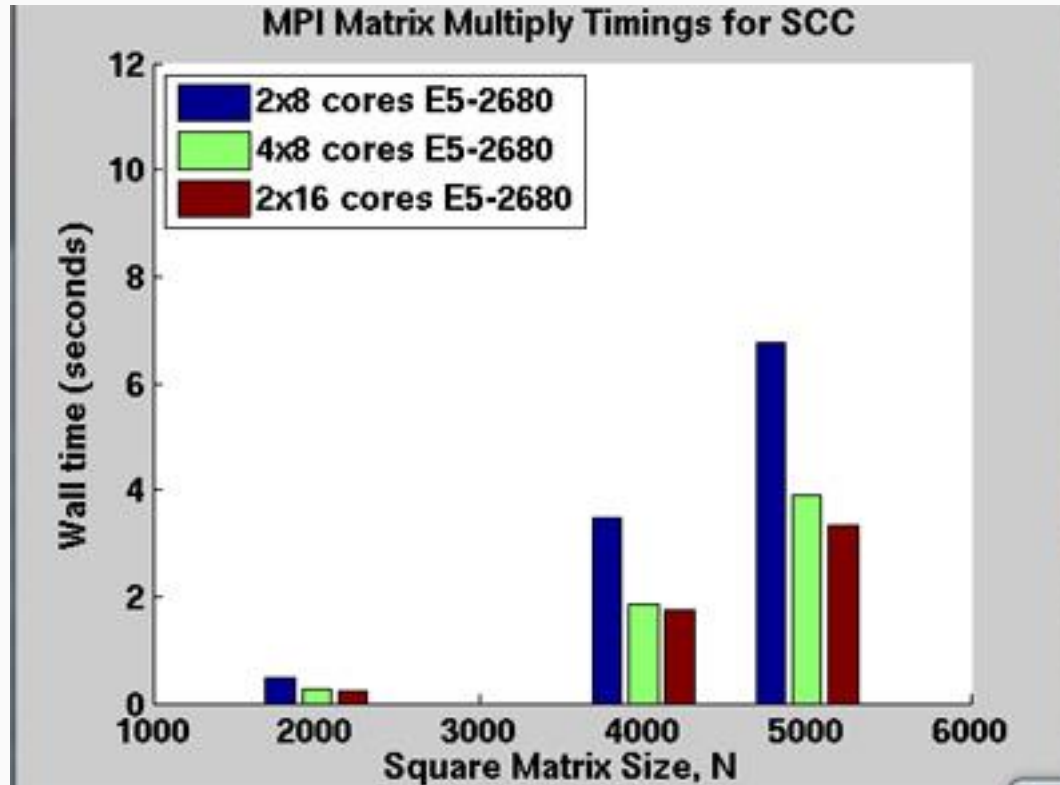
Compilers . . .

- Compilers
 - GNU: *gcc, c++, gfortran* -- optimize with *-O3*
 - PGI: *pgcc, pgCC, pgfortran* -- optimize with *-fast*
- MPI
 - **MPI wrapper:** *mpicc, mpiCC, mpif90* – built for GNU compilers (PGI later)
 - **MPI implementation:** OpenMPI
 - *mpif90 -show* shows what compiler family, linked libs, etc.
- OpenMP – shared-memory directive-based parallel computing for CPUs; turn on with *-mp* (PGI); *-fopenmp* (GNU)
- OpenACC – directive-based parallel computing for GPUs; turn on with: *-acc* (PGI)
- Cuda-based GPU computing (Katia will talk about this)

Batch Usages

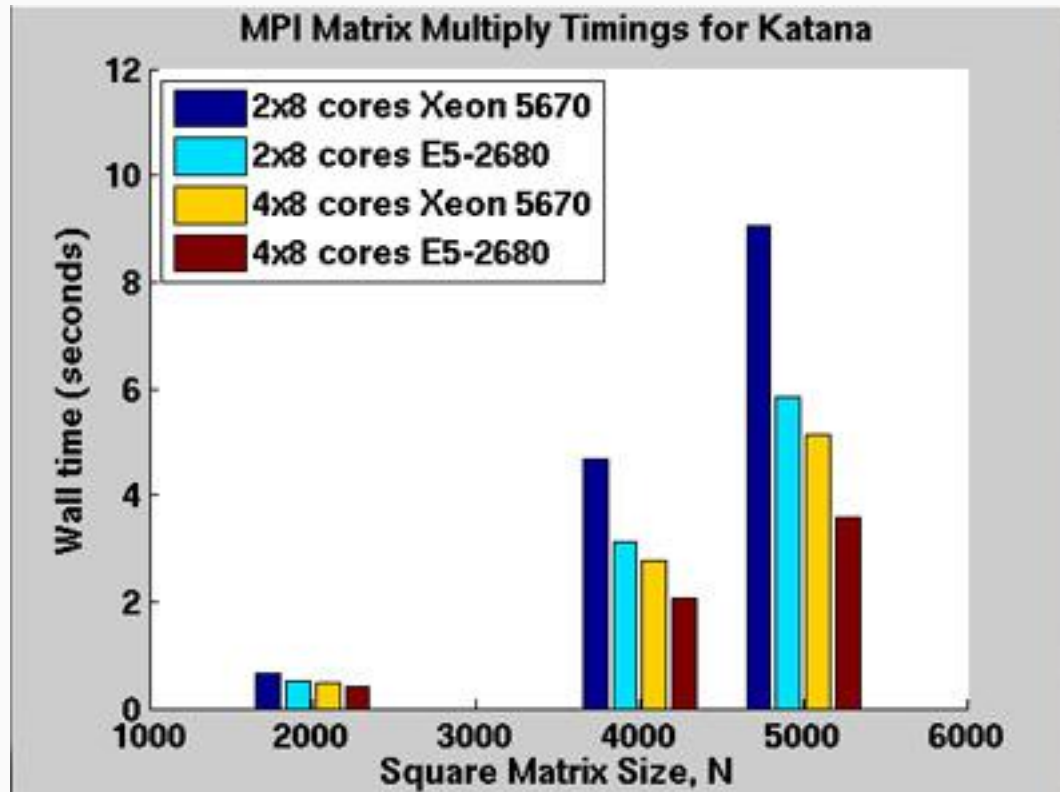
- **Open Grid Scheduler (OGS)** is open source version of Katana's SGE Batch scripts for both are essentially interchangeable
 - Batch commands: *qsub, qstat, qdel, qhost, . . .*
 - *Files needed by a compute node must reside on SCC, e.g., executables and I/O data files (/katana not supported on nodes)*
 - OpenMP or thread-based jobs: *-pe omp N (1 ≤ N ≤ 16)*
 - MPI jobs: *-pe mpi_X_tasks_per_node N (X=4, 8, 16; N=i*X; i=1, 2, 3, . . .)*
 - *Example: X=4, N=16 => 4 nodes with 4 cores each (rest of each node shared)*
 - *Example: X=16, N=16 => exclusive use of an entire node; 2 nodes if N=32*
 - Serial jobs: omit *-pe*
 - Cuda-based GPU jobs: available July 1; use Katana now
 - OpenACC-based GPU jobs: available July 1; use Katana now
 - Default walltime limit: *12 hours* (2 hours on Katana)
 - Maximum walltime limit: *no limit for now* (24/72 hours on Katana)
 - Core use limit: *≤ 256 cores* (fewer later) concurrently per user (128 on Katana)

MPI Square Matrix Multiply on SCC

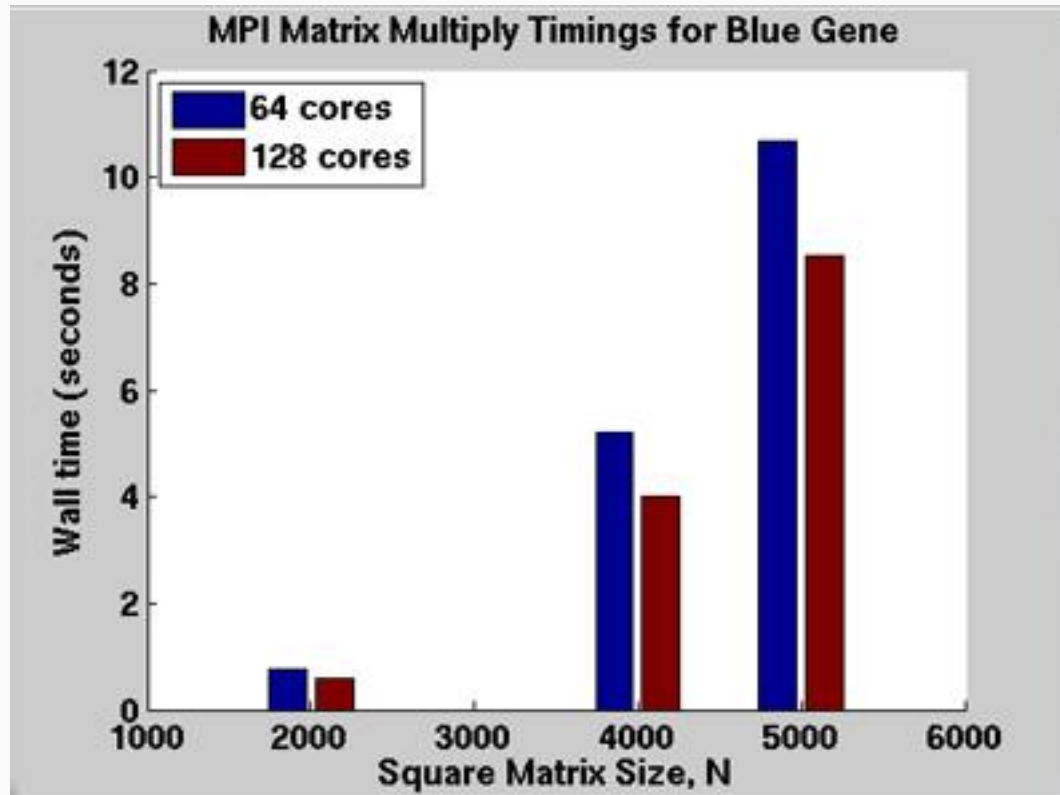


- 2x8 cores case uses `-pe mpi_8_tasks_per_node 16`
- 4x8 cores case uses `-pe mpi_8_tasks_per_node 32`
- 2x16 cores case uses `-pe mpi_16_tasks_per_node 32`

MPI Square Matrix Multiply on Katana



MPI Square Matrix Multiply on the Blue Gene



For this application, wall time is shorter with 16 cores on the SCC than 128 cores on the Blue Gene. With significantly more memory and application software, the SCC is more general-purpose and versatile.

Scratch Disks

- **Primarily to store temporary files during batch jobs**
 - */scratch* on login and compute nodes
 - Not backed up. File life is 10 days
 - Nodes are assigned at runtime. Node info accessible from \$PE_HOSTFILE
 - I/O addressing */scratch* points to local compute node's scratch
 - After batch job completes, assigned nodes are reported in *your-batch-script.poXXXX*. To access your data . . .

```
scc1:~ % cd /net/scc-ad8/scratch
```

Phasing Out of Katana & Blue Gene ?

- Before the end of 2013

Relevant SCC Webpages

- SCC home page
<http://www.bu.edu/tech/about/research/computation/scc/>
- Status of Software packages:
<http://www.bu.edu/tech/about/research/computation/scc/software/>
- Software Updates
<http://www.bu.edu/tech/about/research/computation/scc/updates/>
- Snapshots / IS&T Archive / Tape Archive
<http://www.bu.edu/tech/about/research/computation/file-storage/>
- Technical Summary
<http://www.bu.edu/tech/about/research/computation/tech-summary/>

Useful SCV Info

- **Research Computing Home Page**
(www.bu.edu/tech/about/research)
- **Resource Applications**
www.bu.edu/tech/accounts/special/research/accounts
- **Help**
 - **System**
 - help@scc.bu.edu
 - Web-based tutorials
(www.bu.edu/tech/about/research/training/online-tutorials)
(MPI, OpenMP, MATLAB, IDL, Graphics tools)
 - Consultations by appointment
 - Katia Oleinik (koleinik@bu.edu) – Graphics, R, Mathematica
 - Yann Tambouret (yannpaul@bu.edu) – HPC, Python
 - Robert Putnam (putnam@bu.edu) – VR, C programming
 - Kadin Tseng (kadin@bu.edu) – HPC, MATLAB, Fortran