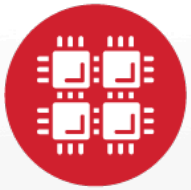

Ohio Supercomputer Center

An OH·TECH Consortium Member

Feb 2014 HPC Tech Talk

Agenda

- Mission
- WebEX tips
- Overview of service changes
- Attendee-driven discussion
- "Tech Notes" (30 minutes) - GPU utilization at OSC for computational chemistry
- Slides are available at
  - http://www.osc.edu/tech_talks

Mission

- We want to better engage the daily users of the system
- Provide another avenue for the community to raise issues and talk about unmet needs.
- This event is for you! Please ask questions, make comments, and provide feedback as to how these can be improved to better serve you.

WebEX tips

- You can use the “Raise hand” icon to ask a question or make a comment (this will notify us so that we can acknowledge you)
- You may also use the Q&A section to ask questions.
- Please mute your microphones when not talking, to avoid feedback and interference noises.

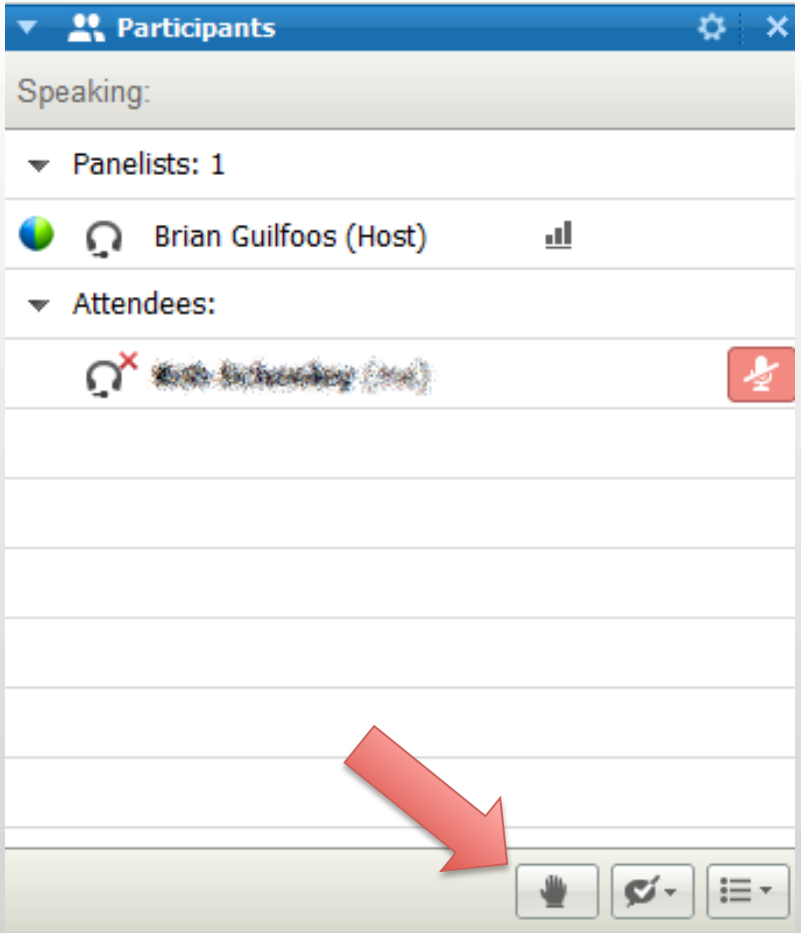

Downtime Summary

- Updated Red Hat Enterprise Linux (RHEL) operating system on all clusters to newer version
  - The Glenn cluster is now running RHEL 5.10, and Oakley is running RHEL 6.5
  - If something no longer works as before, this may be the culprit. Feel free to email oschelp@osc.edu
- GPFS upgrades
  - Much faster I/O speeds, for those of you who have allocations on /nfs/gpfs.

Software Updates

- LS-DYNA 971-R7.0.0 now available on Glenn
- Gaussian license has been renewed
- Intel Compilers and the mvapich2 MPI library are updated regularly on Oakley
  - most recent Intel compiler module is intel/14.0.0.080
  - latest stable mvapich2 release is mvapich2/1.9
- Hybrid MPI/SHMEM or MPI/UPC programming for friendly users with mvapich2-x on Oakley, contact ktomko@osc.edu for more information

Module changes for ANSYS/Abaqus

- Main modules updated to direct users to the correct license server
- Old modules have been deprecated
  - Will be removed in the future
- Effective on both clusters
- I.e., use *`module load abaqus/6.12`* instead of *`module load abaqus/6.12-nimbis`* or similar (Oakley syntax)

Updated software pages

- Now capable of filtering by cluster or by application area.
  - http://www.osc.edu/supercomputing/software

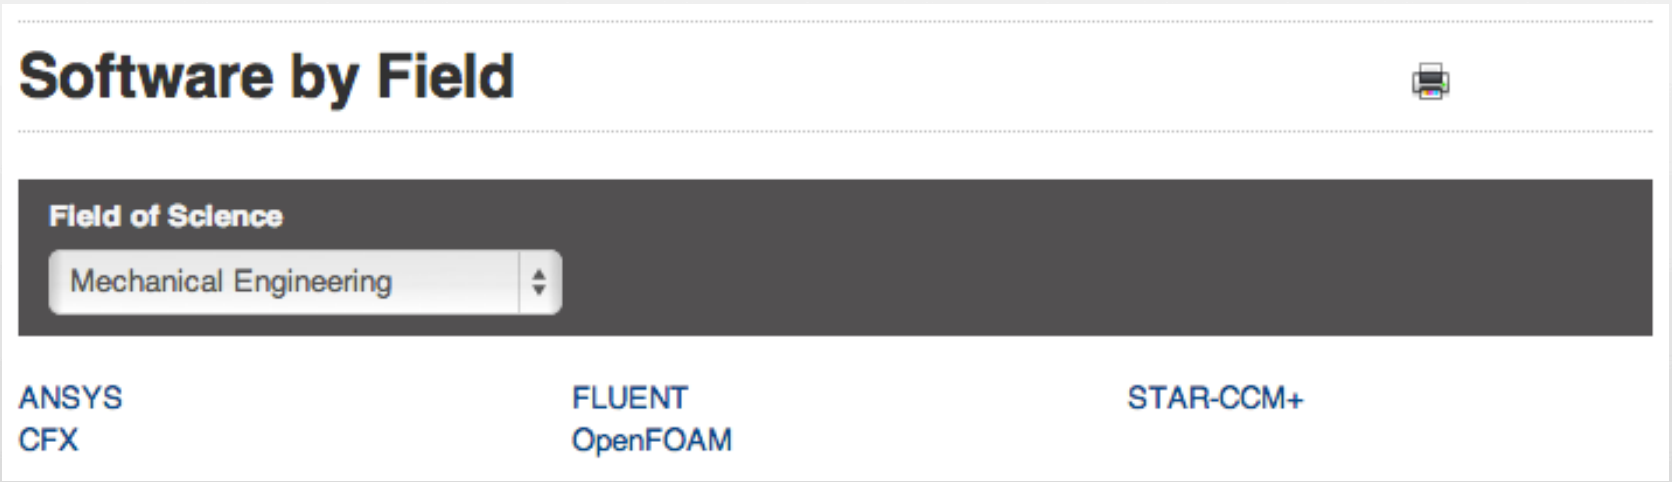

SUG Software Focus Groups

- We'd like your input on decisions regarding software installed for general use (via modules) on our systems
- An invitation to participate will be sent out to users in the next few weeks, or send email to Karen Tomko, ktomko@osc.edu
- Establishing several initial focus groups:
  - Bio informatics/Bio Sciences
  - Fluid Dynamics
  - Structural mechanics
  - Quantum Chemistry/ Materials
  - Physics
  - Atmospheric and Environmental Modeling

New Identity Management Site

- Password and shell changes are no longer performed on ARMSTRONG.
- Accounts can be recovered via email if you have forgotten your password.
- You can change your email address
- If your name is not correct when you log in, contact OSC Help!
- https://my.osc.edu/

Glenn is still in production!

- Long waits on Oakley, but often no wait at all on Glenn
- We can provide assistance in migrating your jobs to Glenn, or helping you decide if you will benefit from switching.
- If there are software packages missing on Glenn that you require, contact OSC Help to have them added.

Known Issues

- You cannot check your quota on GPFS. No workaround; the quota numbers reported when you log in are calculated once per day.
- Look at the bottom of http://www.osc.edu/supercomputing for up to date issue reporting.

Discussion period

- Floor open to questions of presented material, or for the community to raise other issues to discuss.

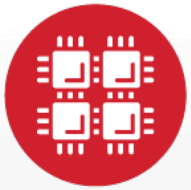

Ohio Supercomputer Center

An OH·TECH Consortium Member

Tech Notes

Computational Chemistry on NVIDIA GPUs at OSC

Dr. Scott Brozell

Goals

- Increase awareness and facilitate usage of underutilized hardware resources and significant software capabilities
- Obtain community feedback

Outline

- Introduction
- GPU Hardware
- GPU Batch Usage
- GPU Policy
- GPU Enabled Software Applications
  - Molecular Dynamics
  - Electronic Structure
  - Examples

Introduction

- Graphics Processing Unit (GPU)
  - Specialized coprocessor for 3D computer graphics
  - Massively data parallel processing
- General Purpose computing on GPUs (GPGPU)
  - Evolving floating point arithmetic support
  - Single Instruction Multiple Data (SIMD) algorithms
- Compute Unified Device Architecture (CUDA)
  - NVIDIA parallel programming model and toolset
  - Partition computation across hybrid computing platform

Csuri Advanced GPU Environment

HP Intel Xeon
Oakley Cluster

IBM 1350 AMD Opteron
Glenn Cluster

Batch Usage

- PBS directives and environment variables
  - Oakley
    - #PBS –l nodes=n:ppn=p:gpus=g
      - +1 <= n, 1 <= p <= 12, and 1 <= g <= 2
    - #PBS –l nodes=n:ppn=12:gpus=2:vis
      - Starts an X server
    - $PBS_GPUFILE
      - Specifies a file containing the allocated GPU identifiers
  - Glenn
    - #PBS –l nodes=n:ppn=8:gpu
      - +1 <= n
      - Automatically allocated a whole node of 2 GPUs

Batch Policy

- Scheduling
  - GPU nodes reserved for GPU jobs
    - Except for backfill – very short jobs
- Charging
  - No charge for GPUs per se
  - Based on total CPU time

- GPU resources are underutilized

GPU Enabled Software Applications

NVIDIA Computational Chemistry Site
  - Complete listing
- Visualization
  - VMD
    - Oakley
    - Automatically detects computational resources
      - qsub –I –X –l nodes=1:ppn=12:gpus=2
      - module load vmd; vmd
    - Automatically recognizes display settings
      - If no display then defaults to text mode
      - glxinfo command reports OpenGL status

Molecular Dynamics Software

- Amber
  - Explicit and implicit solvent simulations
    - Significant speedups
    - Many features
    - Effective GPU only programming model
      - For I/O, one CPU core needed per GPU
      - Oakley nodes=1:ppn=2:gpus=2
        - » pmemd.cuda &; pmemd.cuda &; wait
        - » mpiexec pmemd.cuda.MPI
      - Modest scalability
      - Size limitations
    - Amber 14 planned release in April
      - Improved performance, scaling; more features

Molecular Dynamics Software

- Gromacs
  - 4.6 has native GPU acceleration
    - Non-bonded forces calculated on GPUs
    - Explicit and implicit solvent simulations
    - Significant speedups
    - Full featured
    - Hybrid CPU and GPU programming model
      - Use multiple CPU cores per GPU
      - Oakley nodes=n:ppn=12:gpus=2
        - » cutoff-scheme=verlet
      - Modest scalability

Molecular Dynamics Software

- NAMD
  - Native GPU acceleration
    - Non-bonded forces calculated on GPUs
    - Explicit and implicit solvent simulations
    - Significant speedups
    - Full featured
    - Hybrid CPU and GPU programming model
      - Use multiple CPU cores per GPU
      - Oakley nodes=n:ppn=12:gpus=2
        - » namd2 +idlepoll
      - Modest scalability

Molecular Dynamics Software

- LAMMPS
  - GPU standard package
    - Long-range Coulombics calculated on GPUs
    - Significant speedups
    - Hybrid CPU and GPU programming model
      - Use multiple CPU cores per GPU
      - Oakley nodes=n:ppn=12:gpus=2

Molecular Dynamics Software

- LAMMPS
  - USER-CUDA package
    - Long-range Coulombics calculated on GPUs
    - Significant speedups
    - Mostly GPU only programming model
      - For I/O, one CPU core needed per GPU
      - Oakley nodes=1:ppn=2:gpus=2
  - GPU versus USER-CUDA
    - Generally USER-CUDA has more features
    - Benchmarking recommended

Electronic Structure Software

- Emerging functionality on GPUs
  - GAMESS
    - RHF, RMP2, RCCSD(T)
      - Modest speedups
  - NWChem
    - Various CCSD subtasks
      - Modest overall speedups
    - Installation of 6.3 is planned
  - Q-Chem
    - RI-MP2
      - Significant speedups
      - Installation and testing in progress

Electronic Structure Software

- GPU functionality coming
  - Gaussian
    - Under development
  - Molpro
    - Density fitting methods: HF, DFT, MP2, etc.
      - Modest overall speedups
    - Purchase request was approved

Example Batch Scripts

- Input files in /nfs/10/srb/workshops/compchem

- Amber Oakley dual GPU:
  - amber/ jac9999dualcuda.oakley.pbs
- Amber Oakley parallel GPU:
  - amber/ jac1000cudampi.oakley.pbs

Summary

- Powerful GPU resources are underutilized
- Significant HPC GPGPU performance improvements exist in software applications especially in the area of molecular dynamics
- GPU batch policies are very favorable
- Community feedback is encouraged

Q&A

- Any questions?
- oschelp@osc.edu