

BG/Q

SNUG TechTalk

SciNet

www.scinet.utoronto.ca

University of Toronto
Toronto, Canada

November 5, 2012

Installed Hardware

- SOSICIP (Southern Ontario Smart Computing Innovation Platform)
- 2.5 Racks → 2048 (32,768 cores), 512 (8,192 cores)
- 1.6 GHz PPC 16 cores CPU (204.8 GFlops/CPU) 16GB RAM
- 419.43 TFlops Theoretical (345 TFlops HPL @ 160 kW)
- 40 I/O Drawers (6 GB/s R/W)
- 500TB GPFS dedicated (IB connected)
- 2x Power 730 Devel Nodes
- Redhat Linux 6.2 / CNK









How do I use this thing?

Workflow

- Cross compile on P7 devel nodes (bgxlc,bgxlf) statically linked
- Submit job using Loadleveler
- “runjob” instead of mpirun/mpiexec
- even serial/threaded apps launch with runjob
- minimum dedicated block is 64 nodes (1024 cores)
- sub-block jobs for smaller batches

Not the GPC

BG/Q Design

- Designed for extremely large scalable applications
- Hybrid computing model (MPI & Threading)

Not the GPC

BG/Q Design

- Designed for extremely large scalable applications
- Hybrid computing model (MPI & Threading)

Things to Note

- No shell on compute nodes
 - No direct access (ssh etc)
 - No local scripting
- No /dev/shm
- All I/O offloaded to I/O nodes (64-CN to 1-I/O)
- No TCP/IP on compute node (MPI internode)
- Statically linked binaries
- Big-endian
- 64 bit only

General BGQ Information

http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JUQUEEN/Documentation/Documention_node.html

SciNet Specific BGQ

<https://support.scient.utoronto.ca/wiki/BGQ>