
A First Look at Blue Gene/L

**Martin W. Margo, Christopher Jordan,
Patricia Kovatch, Phil Andrews**
San Diego Supercomputer Center

2005 LCI: HPC Revolution conference



SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



Overview

- What is BG/L
- BG/L software
- Why does SDSC have one?
- BG/L journey from Rochester to San Diego
- How we set it up at SDSC
- Performance numbers
- Lessons learned
- Community, collaboration and other resources
- Questions



SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



What is BG/L?

- Short for Blue Gene/Light
- Massively Parallel system with high density slower CPUs (700 MHz at SDSC)
- 1 rack:
 - 2048 CPU
 - ~ 5.7 Tflops
 - 3 different network for MPI communication
 - 3D Torus for point-to-point communication
 - Fat tree for collective I/O
 - Global interrupt for low latency MPI barrier communication
- DB2 database controls the machine (monitoring, admin, and logging capability)
- Compute node running CNK (Compute Node Kernel), barebones Linux based kernel
- I/O nodes provides permanent storage (disk), runs more advanced Linux kernel

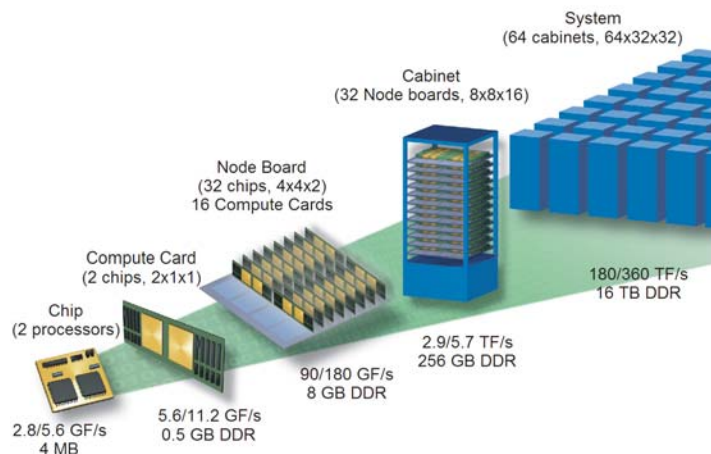


SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



What is BG/L? (cont)



Reference: IBM BG/L Redbook



at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



BG/L software

- **Numerous HPC software are available & more coming soon**
- **IBM provided:**
 - XLC, XLF compiler
 - DB2 client and servers
 - Load Leveler (coming soon)
 - GPFS v.2.3 (coming soon)
 - Mpirun
- **Benchmarks run at SDSC: Linpack, NAS, CG, MP, and Enzo**



SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



Why does SDSC have one?

- **High End Computing + data site**
- **Data intensive applications**
 - ENZO (Astronomy), SCEC (earthquake engineering)
- **Large multi-level storage infrastructure (SAN, SAM-QFS, HPSS)**
- **GPFS over WAN capability**

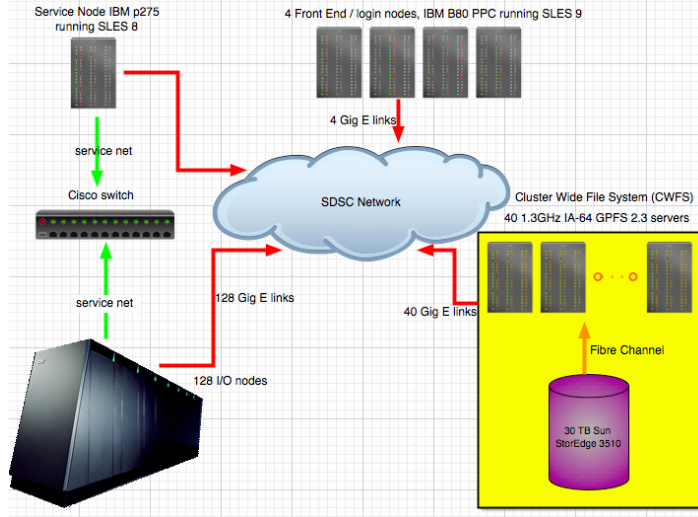


SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



SDSC's setup

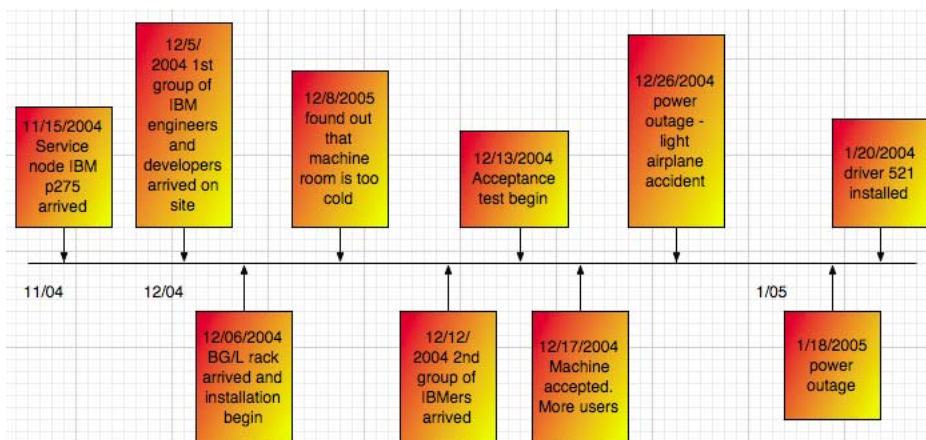


SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



From Rochester to SDSC



SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



BG/L setup at SDSC

- First BG/L machine to have 128 I/O nodes equipped with 1 copper G/E each
- 8:1 compute node: I/O node ratio (max)
- GPFS 2.3 with Multiclustering served over NFS to the I/O nodes
- GPFS served over 40 1.3GHz IA64 servers
- GPFS storage: Sun StorEdge 3510 RAID Array with multi-path
- HPSS access (6 PB + mass archival storage)



SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



SDSC's BG/L

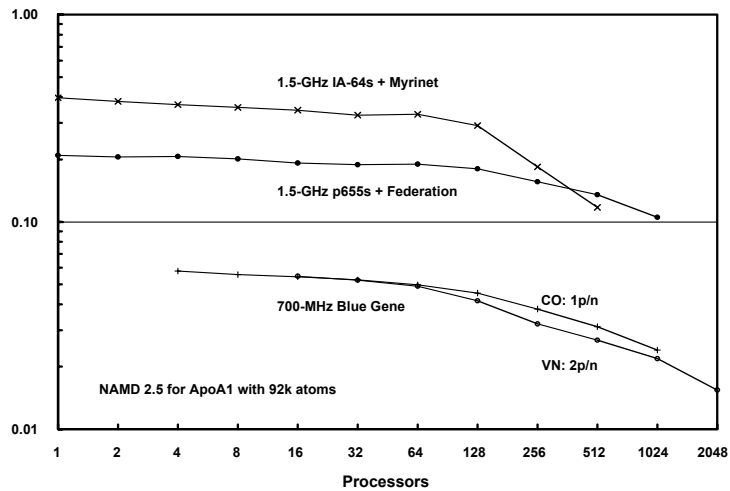


SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



Performance with NAMD



Reference: Wayne Pfeiffer, SDSC



SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



Lessons Learned

- Site/physical data center requirements are moving targets (e.g. needs 3500 cfm / jet engine)
- Obscure requirement (cable must be shielded else it will crash an F-16)
- No authoritative documentation (change by the minute)
- Only the developer can support it. (e.g. Constantly engaging IBM's Mike Woolwood)
- Need weekly calls with IBM BG/L group to get bug fixes, software patches, report bugs, and ask questions.



SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



Collaboration

- **BG/L consortium**
- **Members are from industries, universities, and national laboratories.**
- **A community to foster collaboration between BG/L owner sites.**
- **Provide technical expertise, tools, and porting help for BG/L**
- **Visit <http://www-fp.mcs.anl.gov/bgconsortium/> to learn more**



SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



Thanks to

- **Don Thorp, SDSC**
- **Wayne Pfeiffer, SDSC**
- **Giri Chukkapalli, SDSC**
- **Mark Duffield, IBM**



SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO



Questions ?

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

SDSC

SAN DIEGO SUPERCOMPUTER CENTER

at the UNIVERSITY OF CALIFORNIA, SAN DIEGO

