LQCD (15 months in review)

Xeon Phi / KNL Cluster – "SciPhi XVI"

- 264 nodes (2^8 plus 8), single socket 64 core KNL CPU
- 16 GB high bandwidth memory on package (480 GB/s)
- 192 GB main memory / node
- Intel OmniPath 100 Gpbs fabric, 48 ports/switch, 32 nodes/switch
- LNET routers (4) to route file system traffic back to Infiniband
- IP routers (same 4) to NSF mount /home and to reach other services

Top 500! #397 in performance (HPL on 256 nodes) November 2016

Green 500! #10 in performance/watt

1 of 10 KNL clusters in the November 2016 list

1 of 28 OmniPath clusters in the November 2016 list

Bleeding edge!

- Had to change the firmware and O/S multiple time during commissioning
- Had to figure out LNET; got help from CNI for fault tolerant services
- During 2017 we are (still) replacing borderline early chips using a "screen" test (labor intensive, but free to Jefferson Lab)

Recent Storage Enhancements

- New /work file server (and /home for LQCD)
  - Physics quotas: 144 TB (with 2nd disk chassis, Jan 2018)
  - LQCD: 86 TB
  - SSD write accelerators, higher performance on small files
- LTO-7 integration
  - Active data on LTO-6: 12 drives (8 LTO-6, 4 LTO-7)
  - 1.4 GB/s for uncompressible new and recent data
  - Additional 0.4 GB/s for writing raw duplicates for ejection
  - GlueX wants to run at 800 MB/s (compressed) with 50% to 70% duty factor, so writing raw + duplicate requires up to 1.1 GB/s (more than half of current total library bandwidth)
- Upcoming: LTO-8 integration
  - Can write re-formatted LTO-7 media (called M8) with 50% more data (makes that media more cost effective than LTO-6 media)
  - I/O at 300 MB/s/drive; planning to install 4 to nearly double library bandwidth on recent + new data

Offsite Computing

We have started on support for local plus distributed computing models for FY18 and beyond.

- Peaks and valleys are becoming larger: LQCD no longer a large enough flywheel to smooth out load variations
- Provisioning to peaks is expensive (idle time wastes money)

Options: Send jobs to OSG, Supercomputer Centers, NERSC, Cloud, …

Jefferson Lab has submitted a request to NERSC and DOE for support of GlueX computing. The request is large, and so the decision process is a bit extended.

New Web Apps

The web pages for Physics computing and LQCD computing have now been fully converted to newer adaptive web apps, giving a better interactive experience and supporting screens of various sizes (including cell phones).

As time permits, new features are being added – requests welcome!

LQCD Software Support

The HPC / LQCD team within Scientific Computing continues the pattern of significant software deliverables every year. Recent successes:

- Optimization of graph code for contractions with significant reduction in memory footprint == last stage of the LQCD simulation pipeline that extracts the strength of quark − quark interactions (Jie Chen w/ Robert Edwards)
- Use of COCO template library to demonstrate architecture portable high performance code (Balint Joo)
- Investigations of sparse matrix methods for evaluating 4+ quark systems to avoid exponential cost growth (Frank Winter w/ Robert Edwards)
- ⋯ (many more examples possible)