LQCD-ext FY2012
Hardware Acquisition Plans & Status

Chip Watson

JLab Scientific Computing Group

Outline

- Procurement Planning for 2012
  - Two clusters, one accelerated and one not
- Performance Goals for 2012
- GPU Performance Data
- Un-accelerated Infiniband Cluster Procurement
- Accelerated Cluster Procurement Plans
- File Services Expansion

Procurement Planning

Each year the project goes through an alternatives analysis to decide upon the best strategy for the coming year's hardware procurements.

This process, executed last summer, resulted in the decision to deploy two clusters at Jefferson Lab in FY2012:

1. A standard x86 Infiniband cluster
2. An Infiniband cluster with accelerated nodes

The combination of these resources would provide the highest performance increment while supporting a mix of applications (some able to exploit accelerators and some not).

With the planned de-commissioning of the 2007 LQCD-ext cluster at JLab, plus a planned air-conditioning increase, adequate facilities would be available at JLab.

Performance Goals

The original project goal for 2012 was a system sustaining an aggregate of 24 Tflops average inverter performance. This goal, set in 2009, was based upon the Moore's Law trend for x86 clusters, with some contingency built in.

With the good experience in exploiting GPUs in the ARRA LQCD project, and the growth in the demand for GPU cycles, it was decided that additional accelerated-node capacity could be profitably absorbed. This was true even with the anticipated release to production of the GPU cluster at Fermilab.

A decision was taken to defer the exact dollar split between the standard x86 cluster and the accelerated cluster, but data available indicated that the optimal split would be somewhere around 40% - 60% for each cluster, allowing healthy growth for each type of cluster.

Amdahl's Law Impacts

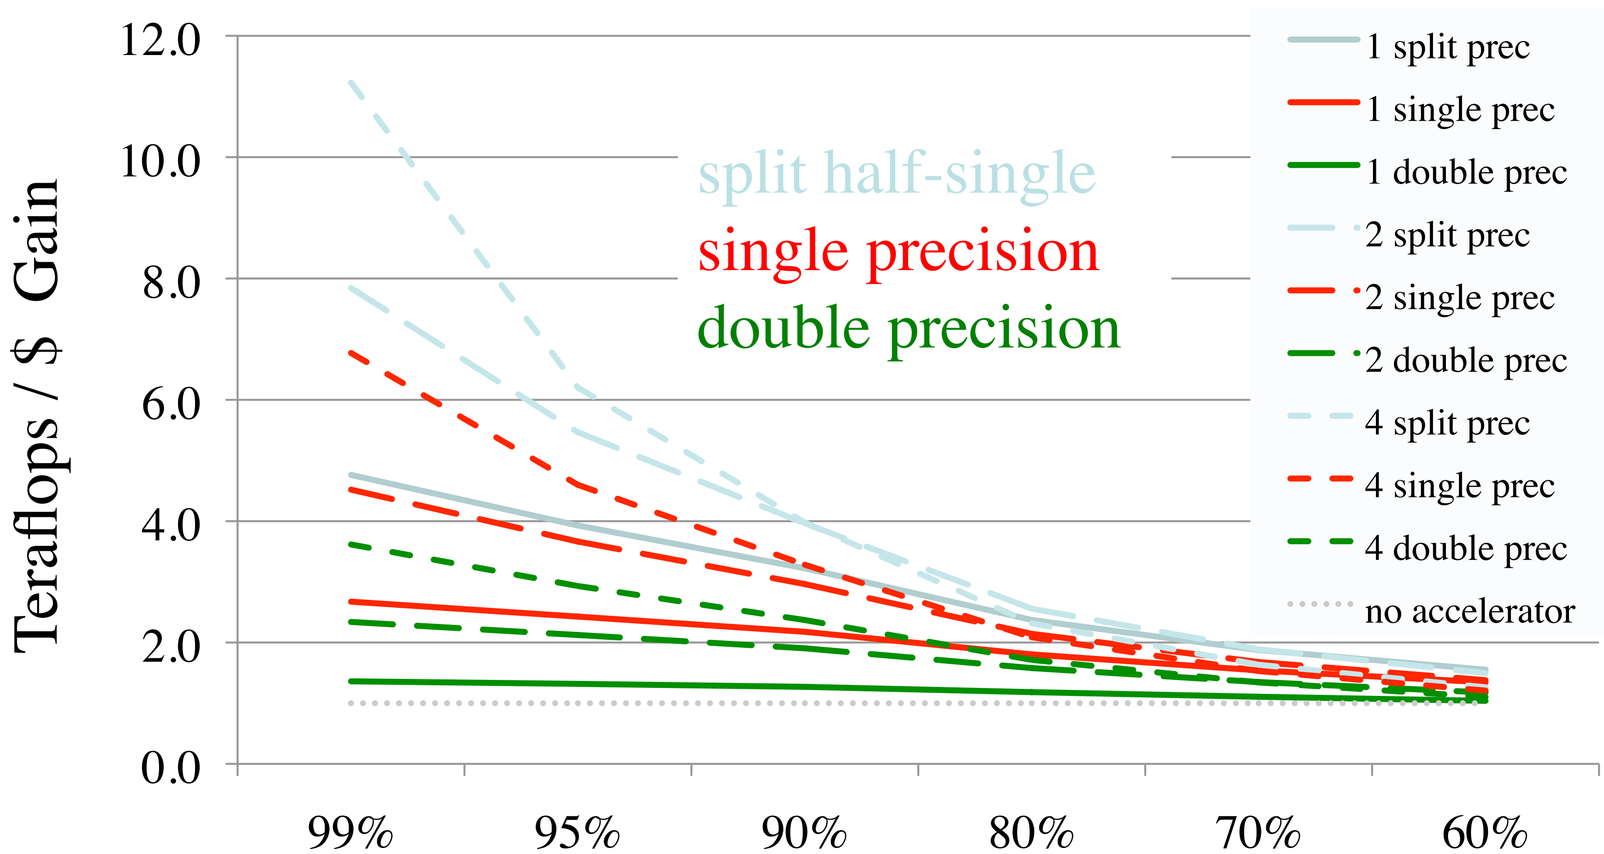

Data above is for the split half-single precision clover inverter, for 1, 2 or 4 GPUs, and for different floating point precision, for Westmere Intel and Fermi NVIDIA generations. Gains are clearly highly application dependent.

Project Goals for 2012 (2)

The ARRA project deployed considerable resources with ~70% of the GPUs being gaming cards targeting the "low hanging fruit" of heavy (~99%) inverter usage.

The LQCD-ext project recently deployed 152 GPUs (76 duals) with Fermi Tesla GPUs, more than doubling the ECC enabled resource for accelerating more than inverters.

For 2012 we anticipated continuing that trend towards ECC enabled systems. Double precision is also expected to become more important. Even though the single precision performance per dollar is lower than gaming cards, the applicability of the Tesla line is wider, and double precision will be better.

As a conservative estimate of the applications' need for ECC, and likely lower acceleration of more complex software than the inverters, we currently estimate an average 3x improvement for Teraflops/dollar over conventional clusters.

Performance Goals for 2012 (3)

The Alternatives Analysis done last summer concluded that the likely x86 cluster would have performance of 14 Tflops/$M, and the accelerated cluster a gain of ~3.5x. Assuming $1M for each cluster, this would yield 14 + 50 Tflops = 64 Tflops.

The first 40% installment of the x86 cluster has now been procured at 12.5 Tflops/$M (details to follow), and the more conservative 3x gain now leads to a total performance range estimate of 45 – 55 Tflops, approximately twice the original 2009 goal.

2012 Procurement Plan

High Level Synopsis:

1. Buy the x86 cluster first, as Interlagos and Sandy Bridge would both be available by early 2012.
2. Buy the accelerated cluster as late as possible, hopefully to include as an option the next generation NVIDIA Kepler GPU (Tesla version with ECC).
3. Under a continuing resolution, the x86 procurement could be started in January, and the second procurement in June.
4. Buy disk as needed, pacing the growing compute capacity.
5. Follow a schedule that would allow all hardware to be delivered by the end of the fiscal year.

x86 Cluster Best Value Procurement

We aim our procurements at aggregate performance for the available budget. We use single node benchmarking, but take into account the network impacts. We further shape this so that the operational costs are not excessive.

1. Infiniband fabric must be non-over-subscribed QDR for nodes with an average performance for 4 benchmarks (asqtad, clover, DWF, new dslash) of 50-100 Gflops.
2. Nodes of performance below 52 Gflops would be assessed an additional 5% scaling penalty to reach > 1 Tflops.
3. Single node benchmark average must be at least 20 Gflops.
4. Nodes must have a power of 2 number of cores.
5. The cost of power was assessed as $2.5K / KW, reflecting the cost of power at Jefferson Lab over a 3 year lifetime.

x86 Cluster Procurement

Timeline

Jan 4	RFI issued
Jan 24	RFP issued
	date was set to allow 40% of funds to accumulate under a continuing resolution by an anticipated award date of late February or early March; due date later changed to March 2 due to lag in getting benchmark codes available (especially Interlagos)
Mar 30	award made for 42% of the hardware budget
May 30	anticipated delivery (vendor anticipates 1 week early)
June	early use period
July 1	first phase in production, in time for new allocations

x86 Cluster Award

The winning system is based upon dual socket 8 core 2.0 GHz Sandy Bridge Xeon, with the required 32 GB memory (4 GB DDR3-1600 per memory channel). Higher clock speeds were not cost effective as the applications are memory bandwidth bound.

42% of funds awarded:

- 212 nodes, configured as 53 2U quad servers in 6 racks
  - 18 nodes per 36-port leaf switch, 12 leaf switches
    - 6 QDR switches as a spine switch, 3 uplinks/leaf/spine
    - 2 fiber optic QDR uplinks to the file system
      - (which is why there are only 212 nodes instead of 216)

The nodes are rated at 50 Gflops (asqtad, DWF) assuming a 10% network overhead, so this is a 10.6 Tflops system (to be confirmed on multi-node runs). This is very close to the conservative 2009 estimate of 12 TF/$M, indicating that we have consumed most of our contingency on unaccelerated x86 systems to a slowing Moore's Law.

Setting the % Split

In the 2012 call for proposals for allocations, the x86 and GPU resources were similarly oversubscribed in the responses, each by a factor of about 1.7.

Since the cost of the deployed x86 clusters is much larger than the GPU clusters, it would require setting the split to favor as much as possible the x86 cluster to scale them both up to be equally over subsubscribed. Using 60% for x86 would leave the GPUs more oversubscribed until that system was put into production in the fall, and then less oversubscribed afterwards, on average balancing well over the year.

Additional input gathered at the USQCD collaboration meeting May 4-5 at Fermilab reinforced this choice, and a procurement of an additional 2 racks is planned.

Accelerated Cluster Procurement

Timeline

Jun 1	RFI issued (also x86 expansion awarded)
Jun 15	RFQ issued, responses due July 18
July 30	award made for the remaining compute budget
Sept 30	system delivery
Oct-Nov	early use
Dec 1	production running

Risk Factors:

- Kepler Tesla GPU might not be deliverable by Sept 30
- Performance might be lower than planned
- Exploitation of GPUs by an increasing fraction of the collaboration might stall

Possible Problems & Responses

- Kepler late
  - could buy a mix of C2070 Fermi Tesla and GTX-580 to achieve performance goals, but with less ECC and double precision (will mine proposal data to determine feasibility, likely impacts)
  - could accept lower performance to favor ECC
- Kepler under performs
  - could buy a mix of Kepler Tesla and GTX-580 or -680 to achieve performance goals, but with less ECC
  - could accept lower performance to favor ECC
- GPU exploitation stalls
  - if known by late July, could order some nodes without accelerators and use for x86 jobs, effectively increasing the x86 fraction to above 60%, with lower total performance, leaving the door open to later acceleration

Disk Capacity Growth

New workflows are becoming more and more disk intensive for intermediate files (propagators, perambulators). Instead of the 5% of hardware the LQCD-ext project had previously used, we are increasing this to 8% of hardware, with some going to the non-deployment site.

2012: procure additional Object Storage Servers for the existing (ARRA) Lustre file system:

- 34 of 2TB disks: 3 * RAID-6 8+2, plus mirrored journal file system and mirrored system disk; w/ QDR Infiniband; yields 40 TB usable / server
- Two servers were already bought this year, and additional systems of this configuration are planned for this summer.
- Next systems might use 3 TB disks, for 60 TB usable.

Summary

The LQCD-ext project continues to grow its GPU resources as software developments permit, while continuing to deploy conventional x86 resources for applications that are not appropriate for GPUs, and for software developments where the focus is upon time-to-develop more than performance.

GPU systems more than compensate for a slowing Moore's Law for LQCD on x86, and have allowed us to tackle relevant science problems in a more timely fashion.

Disk demands are growing, and the split between compute and disk has been correspondingly adjusted.

A just-in-time process for determining the split among the various resources optimizes our science capabilities each year.