LQCD Facilities at Jefferson Lab

Chip Watson

Apr 16, 2010

Jefferson Lab

Thomas Jefferson National Accelerator Facility

Infiniband Clusters

9q 2009 **Q**DR IB
2.4 GHz Nehalem
24 GB mem, 3 GB/core
320 nodes
Quad data rate IB

Segmented topology:
6 * 256 cores 1:1
1 * 1024 cores 2:1

7n 2007 **in**finiband
2.0 GHz Opteron
8 GB mem, 1 GB/core
396 nodes, 3168 cores
Double data rate IB

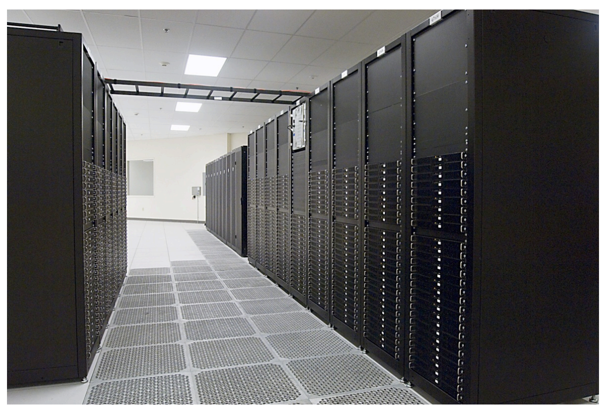

__7n__ has already changed since this photo was taken, shrinking to 11 racks to increase heat density to accommodate new clusters

2009 ARRA GPU Cluster

9g 2009 **G**PU Cluster
2.4 GHz Nehalem
48 GB memory / node
65 nodes, 200 GPUs

Original configuration:
40 nodes w/ 4 GTX-285 GPUs
16 nodes w/ 2 GTX-285 + QDR IB
2 nodes w/ 4 Tesla C1050 or S1070

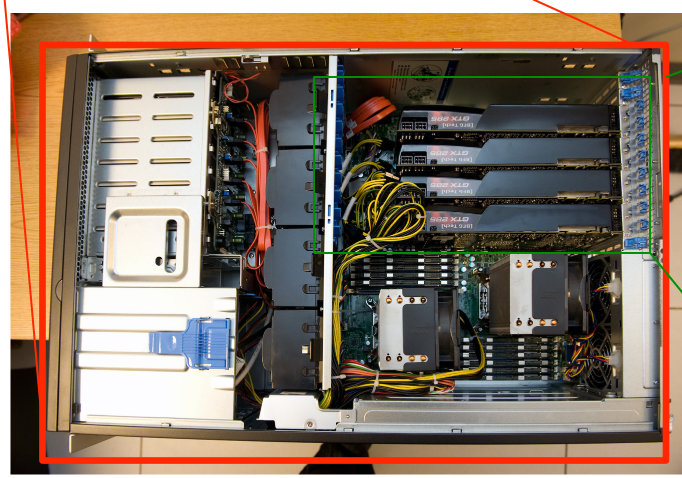

Operations

Fair share: (same as last year)

- Usage is controlled via Maui "fair share" based on allocations
- Fairshare is adjusted ~monthly, based upon remaining time
- Separate projects used for the GPUs, treating 1 GPU as the unit of scheduling, but still with node exclusive jobs

Disk Space:

- Added 200 Tbytes (ARRA funded)
  - Lustre based system, served via Infiniband
  - Will be expanded this summer (~200 TB more)
- 3 name spaces:

  /work (user managed, on SUN ZFS systems)

  /cache (write-through cache to tape, on ZFS, will move to Lustre)

  /volatile (daemon keeps it from filling up, project quotas, currently using all of Lustre's 200 TB)

9½ month Utilization

From: 2009-07-01 (yyyy-mm-dd) To: 2010-04-13 (yyyy-mm-dd) For: all

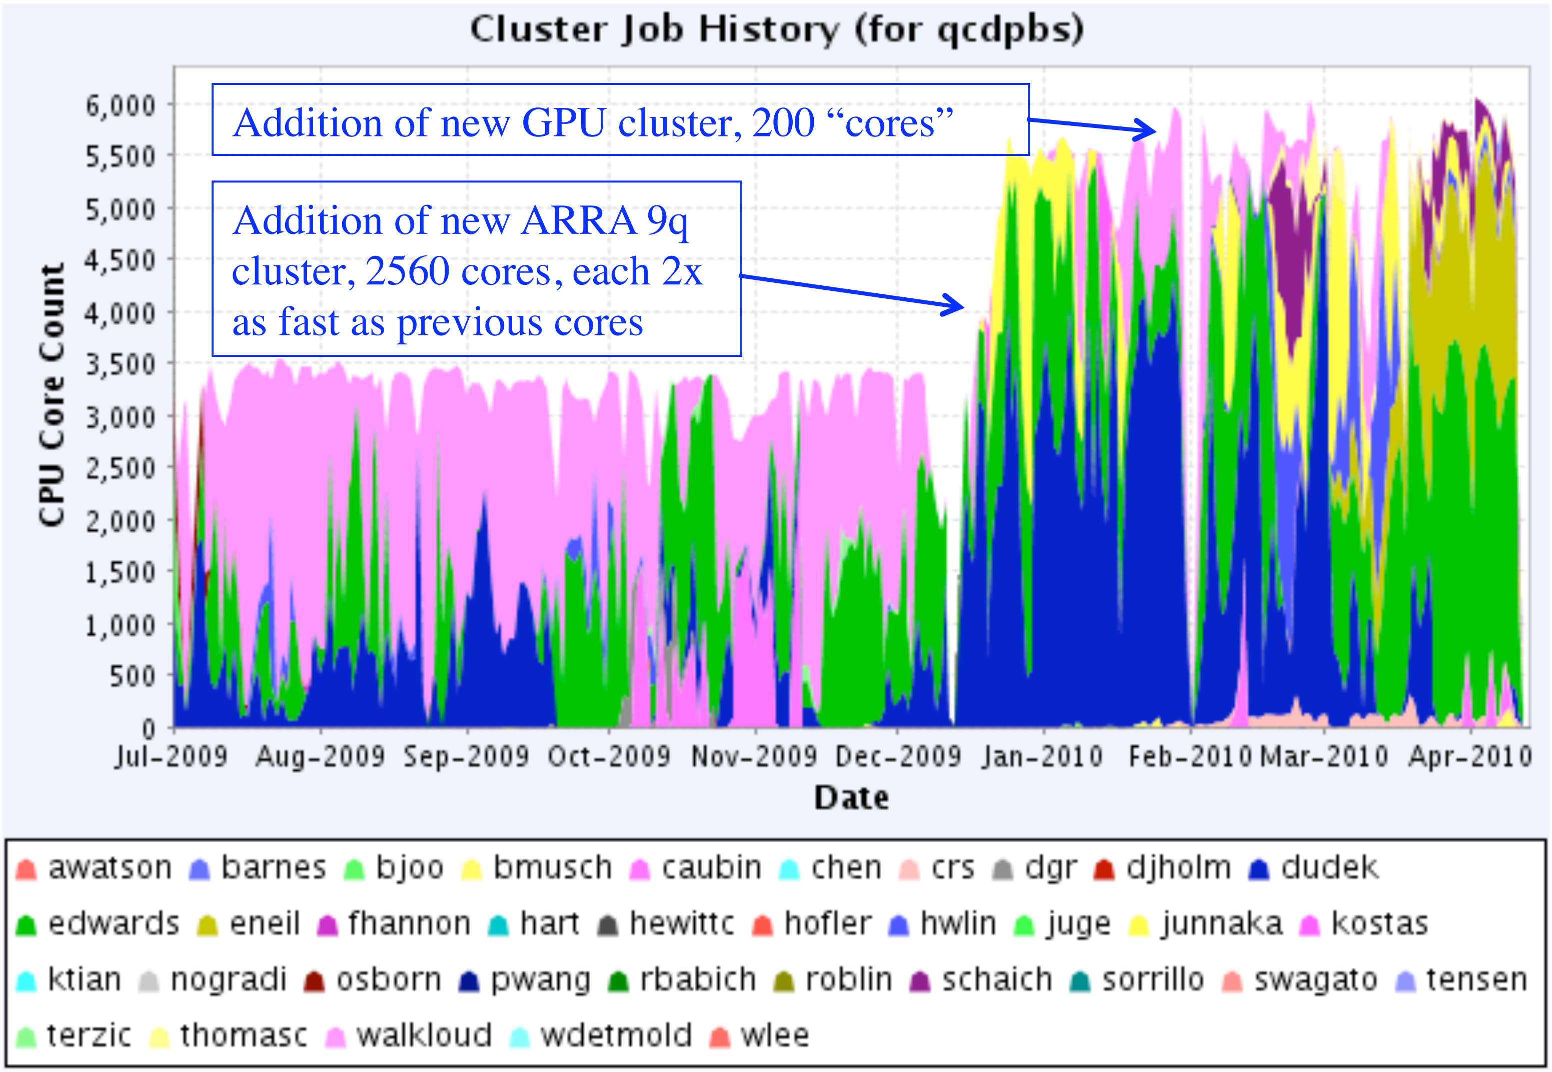

Note: multiple dips in 2010 are power related outages to prepare for installation of 2010 clusters plus O/S upgrade and Lustre file system deployment onto the 7n cluster.

Job Sizes

Last 12 Months

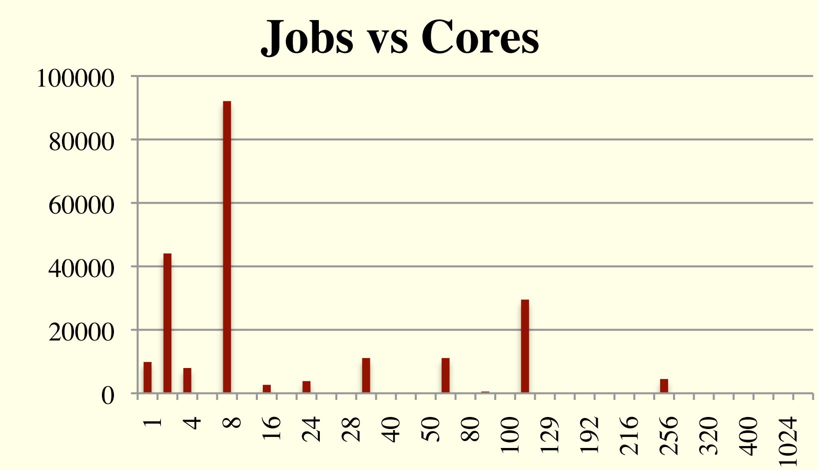

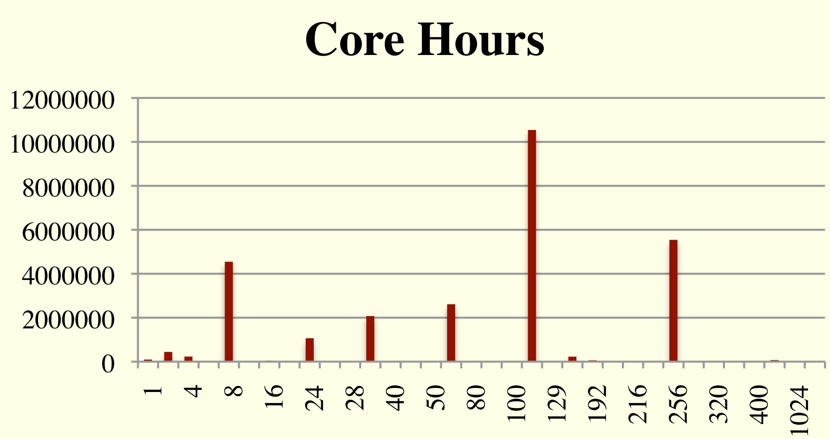

Last 3 Months

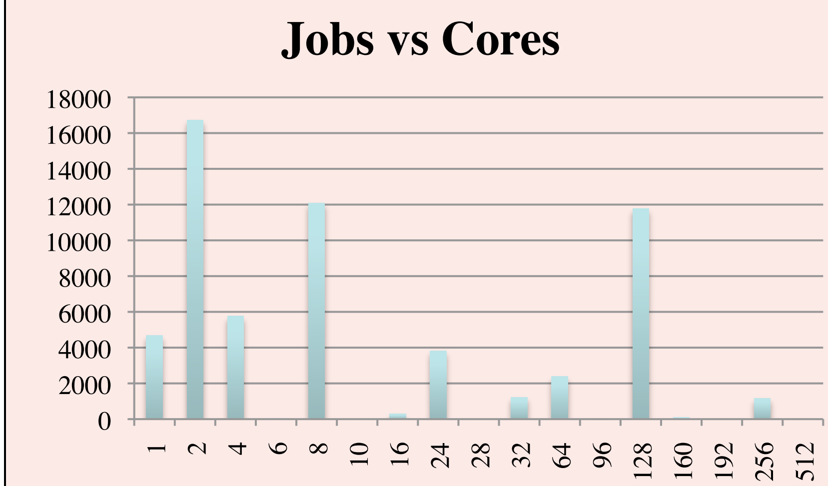

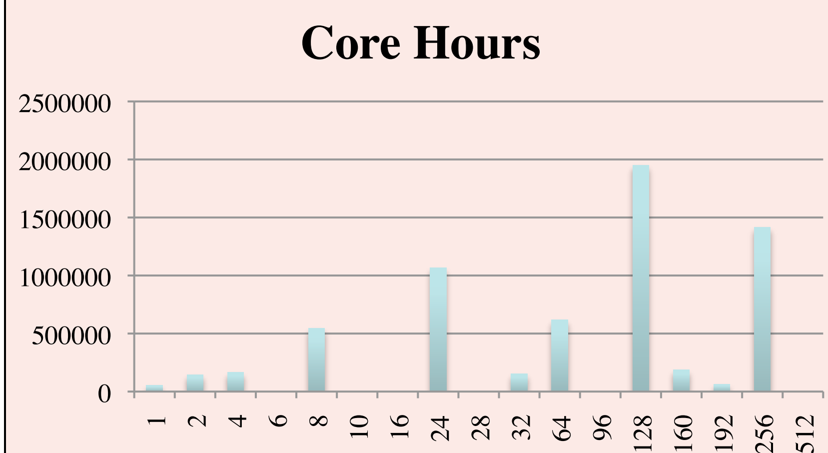

2010 Clusters

10q 2010 **Q**DR IB
2.53 GHz Westmere
8 core, 12 MB cache
24 GB memory,
224 nodes
Quad data rate IB
All node GPU capable

Segmented topology:
7 * 256 cores

Coming Soon!

10g 2010 **G**PU Cluster
2.53 GHz Westmere
48 GB memory / node
~50 nodes, ~300 GPUs

>100 Fermi Tesla GPUs
>100 GTX-480 gaming GPUs
16 nodes w/ QDR Infiniband
Some GPUs to go into 10q
Installation ~ June, 2010

Notes:

- Fermi Tesla GPU has 4x double precision of Fermi gaming cards, plus ECC memory, 2.6 GB with ECC on
- GTX-480 costs ¼ Tesla, $500 vs. $2000 per card, but has only 1.5 GB memory
- Fermi (both Tesla and GTX) have about 10% higher single precision performance of GTX-285 cards

Disruptive Technology -- GPGPUs

Spring 2009 slide

GPGPUs (general purpose graphics processing units) are reaching the state where one should consider allocating funds this Fall to this disruptive technology: hundreds of special purposes cores per GPU plus high memory bandwidth.

Integrated node+dual GPU might cost 25% - 75% more, but yield 4x performance gain on inverters yields 2.5x – 3x price/performance advantage

Challenges

- Amdahl's law: impact being watered down by fraction of time the GPGPU does nothing
- Software development: currently non-trivial
- Limited memory size per GPU

Using 25% of funds in this way could yield 50% overall gain.

Disruptive Technology -- Reality

Software status: *(further details in Ron Babich's talk)*

- 3 different code bases are in production use at Jlab
- Single precision is ~100 Gflops/GPU
- Mixed single / half precision is ~200 Gflops/GPU
- Multi-GPU software with message passing between GPUs is now production ready for Clover
- Many jobs can run as 4 jobs / node, 1 job per GPU, or with multi-GPU software 2 jobs of 2 GPUs which minimizes Amdahl's law's drag and yields 400-600 Gflops / node (rising as code matures)

Price Performance

- Real jobs are spending 50% - 90% in the inverter;
  at 80%, a 4 GPU (gaming card) node in mixed single/half precision yields >600 Gflops for $6K, thus 1 cents / megaflop
- Pure single precision and double precision are of course higher cost, and using the Fermi Tesla cards with ECC will double the cost per Mflops, but with the potential of reducing Amdahl's law's drag

Spring 2010

Weak Scaling: $V_s=24^3$

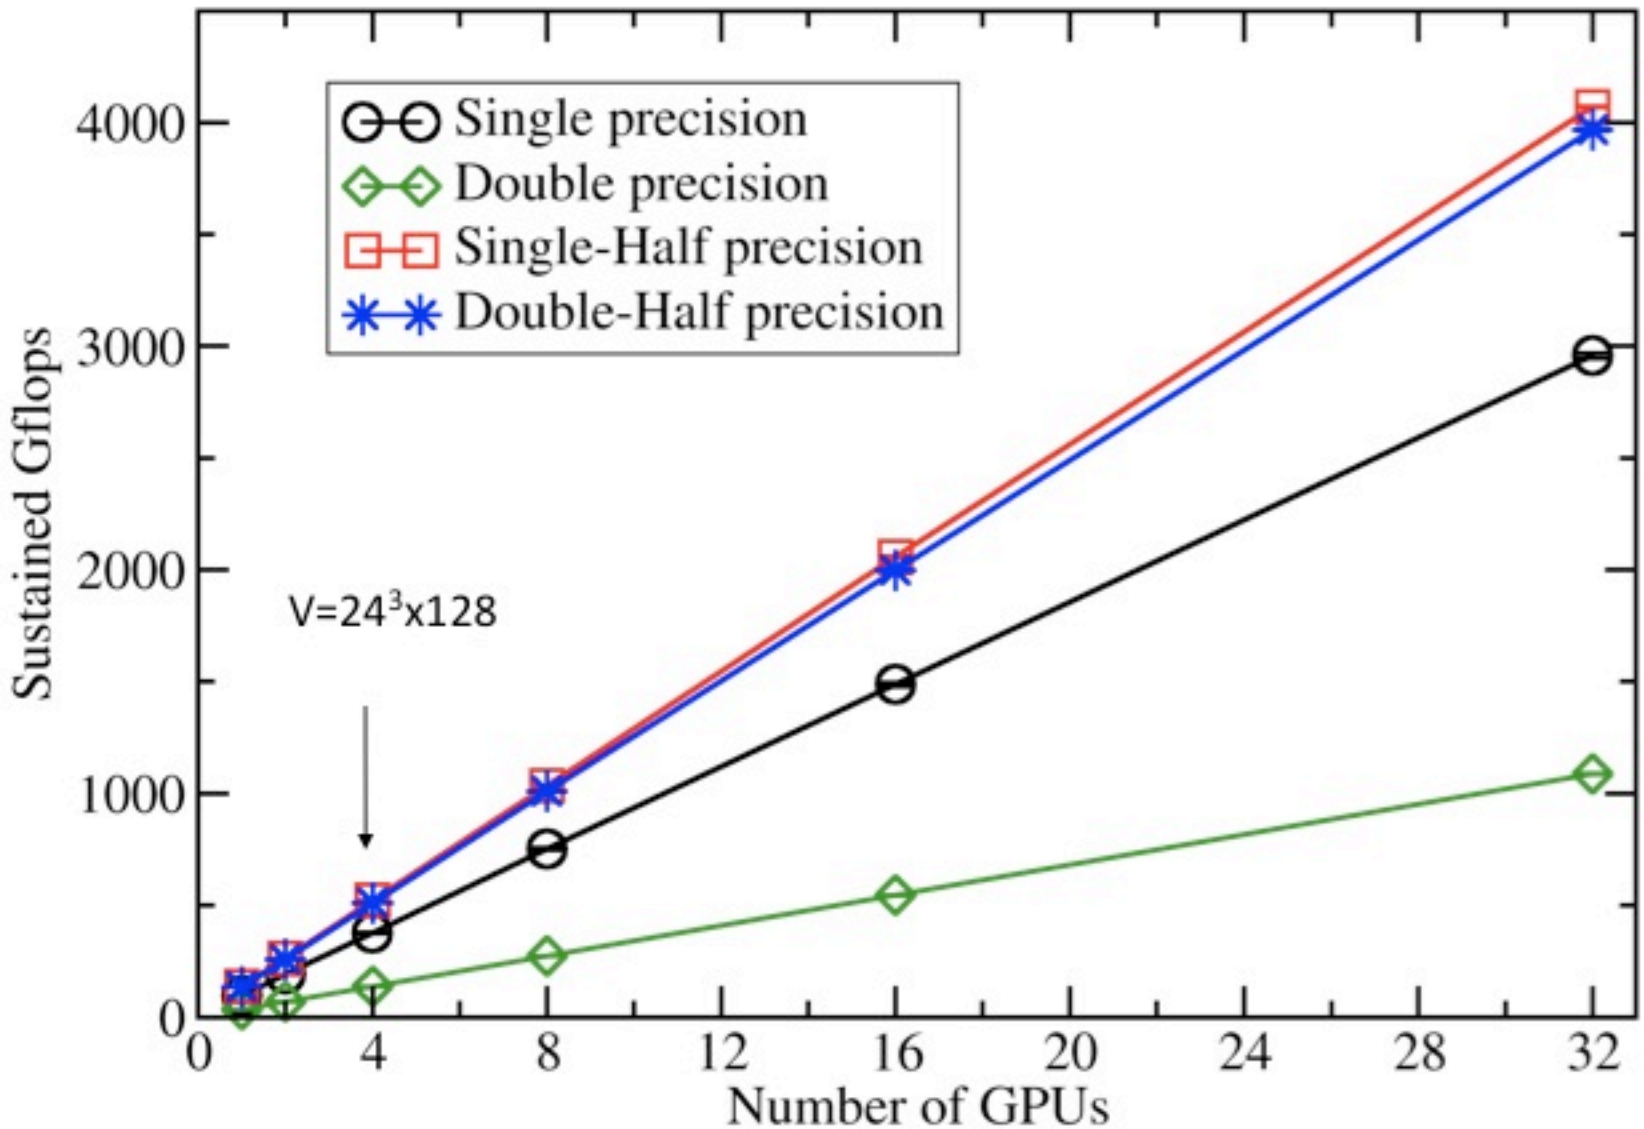

Weak Scaling: $V_s=32^3$

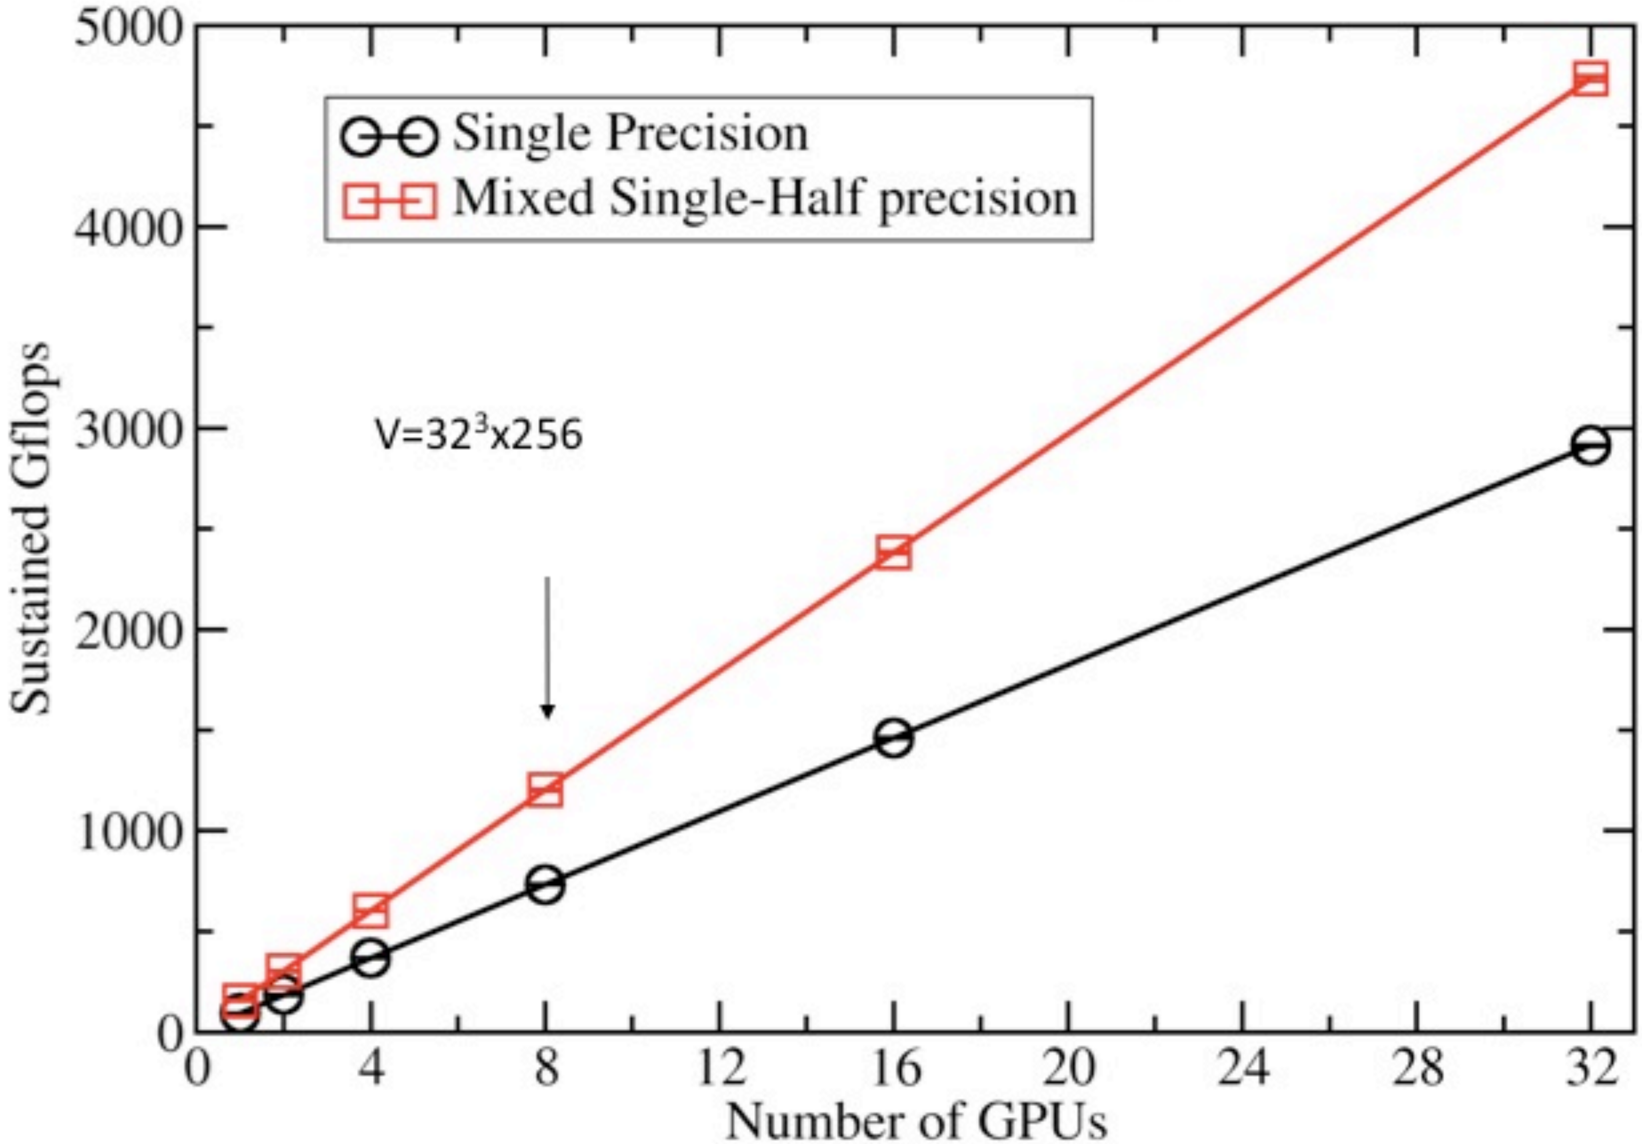

Strong Scaling: $V=32^3x256$

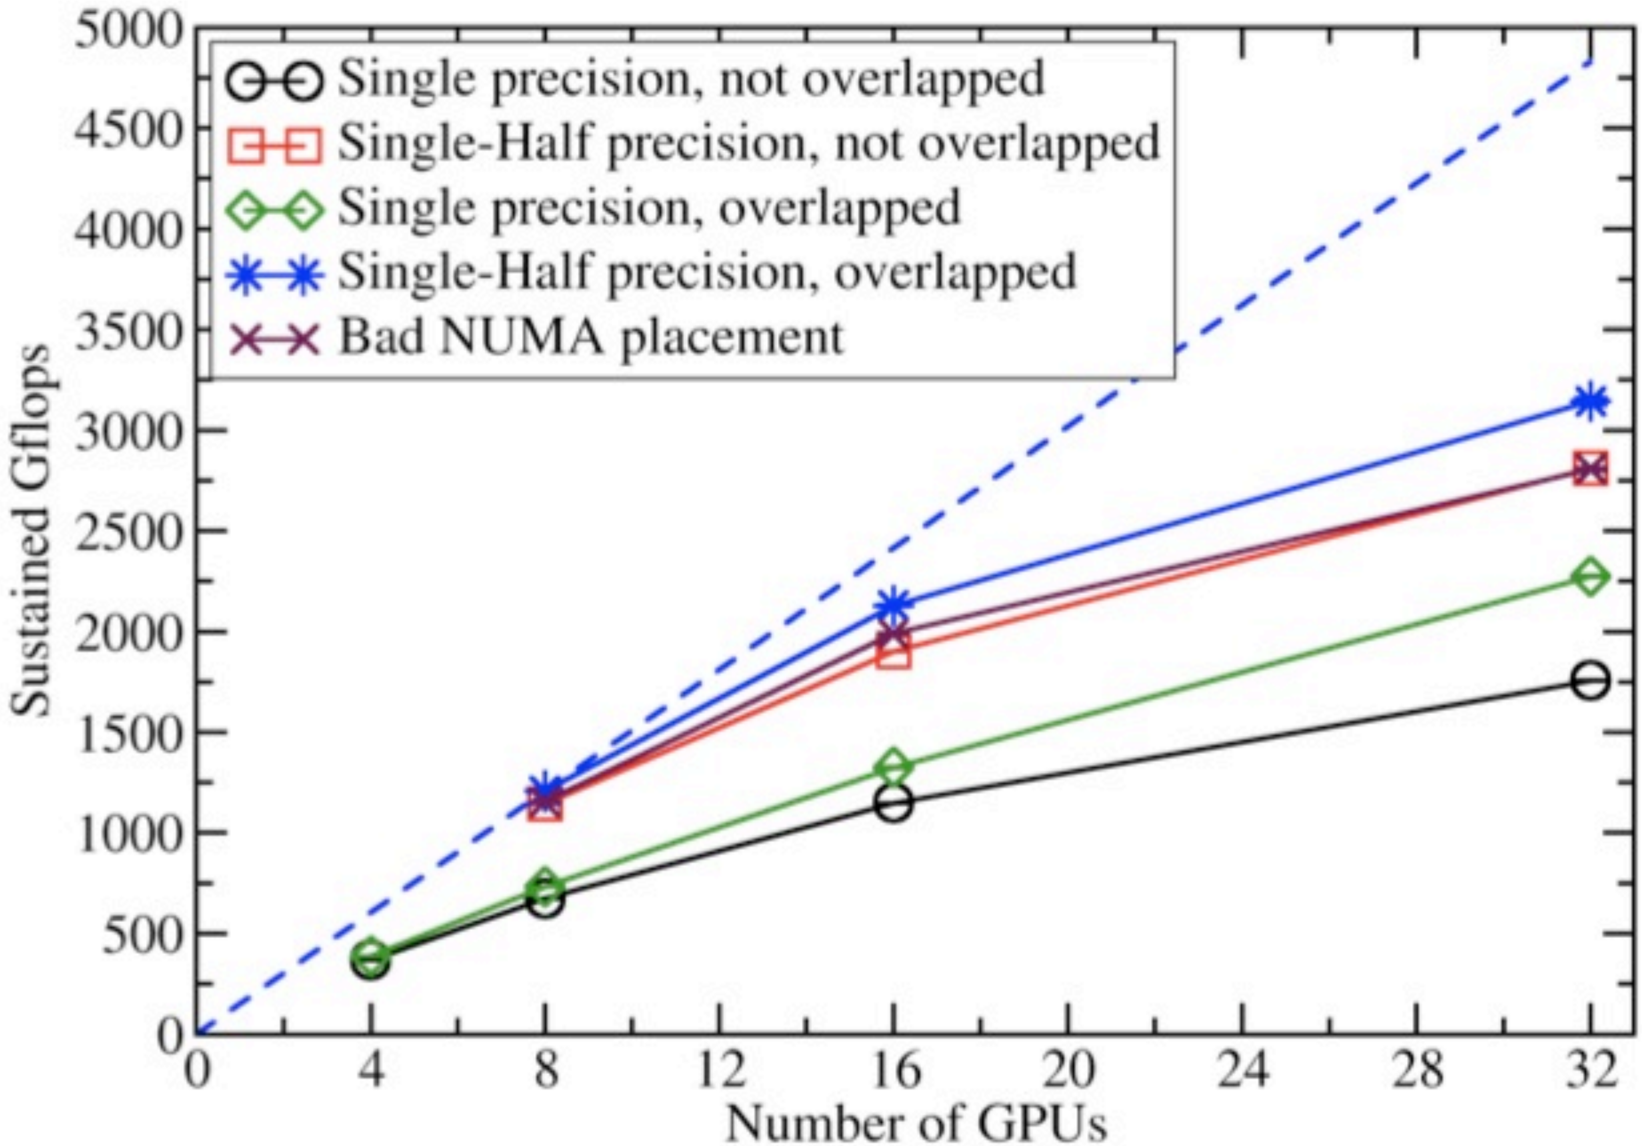

Strong Scale: V=24^3x128

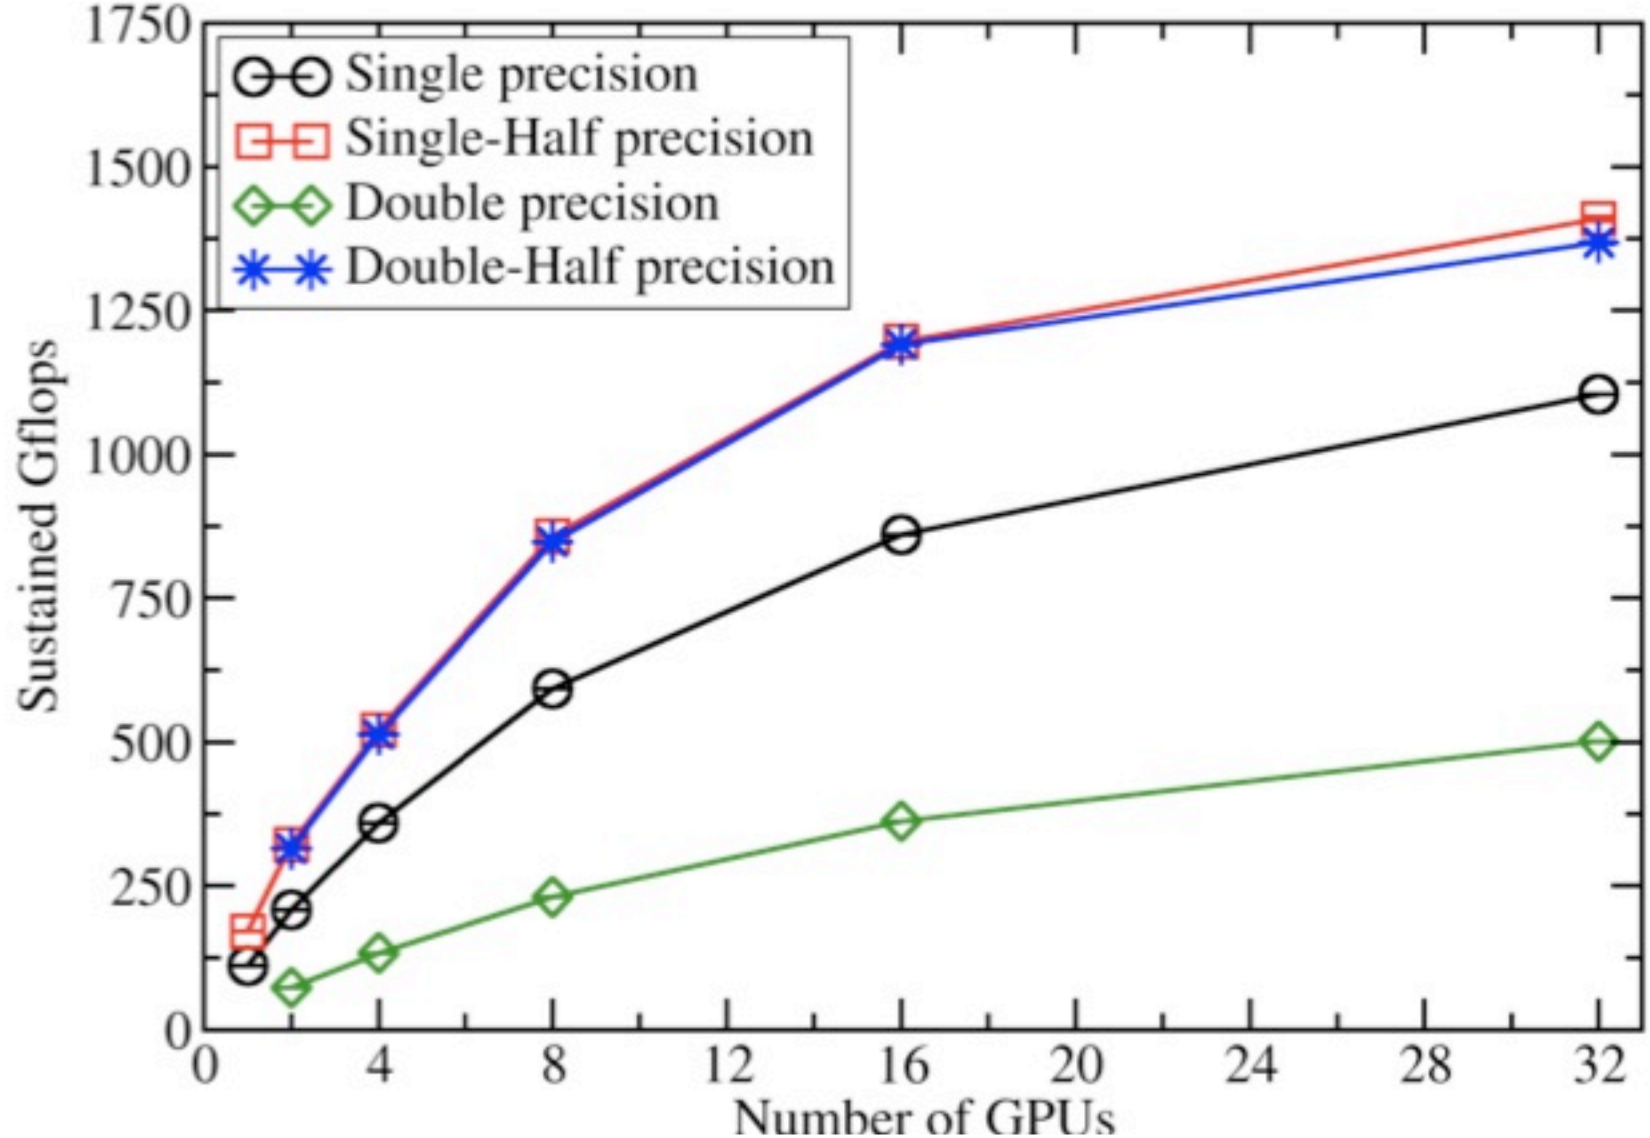

Price-Performance Tweaks: 24^3x128

3 ways of performing the V=24^3x128 calculation

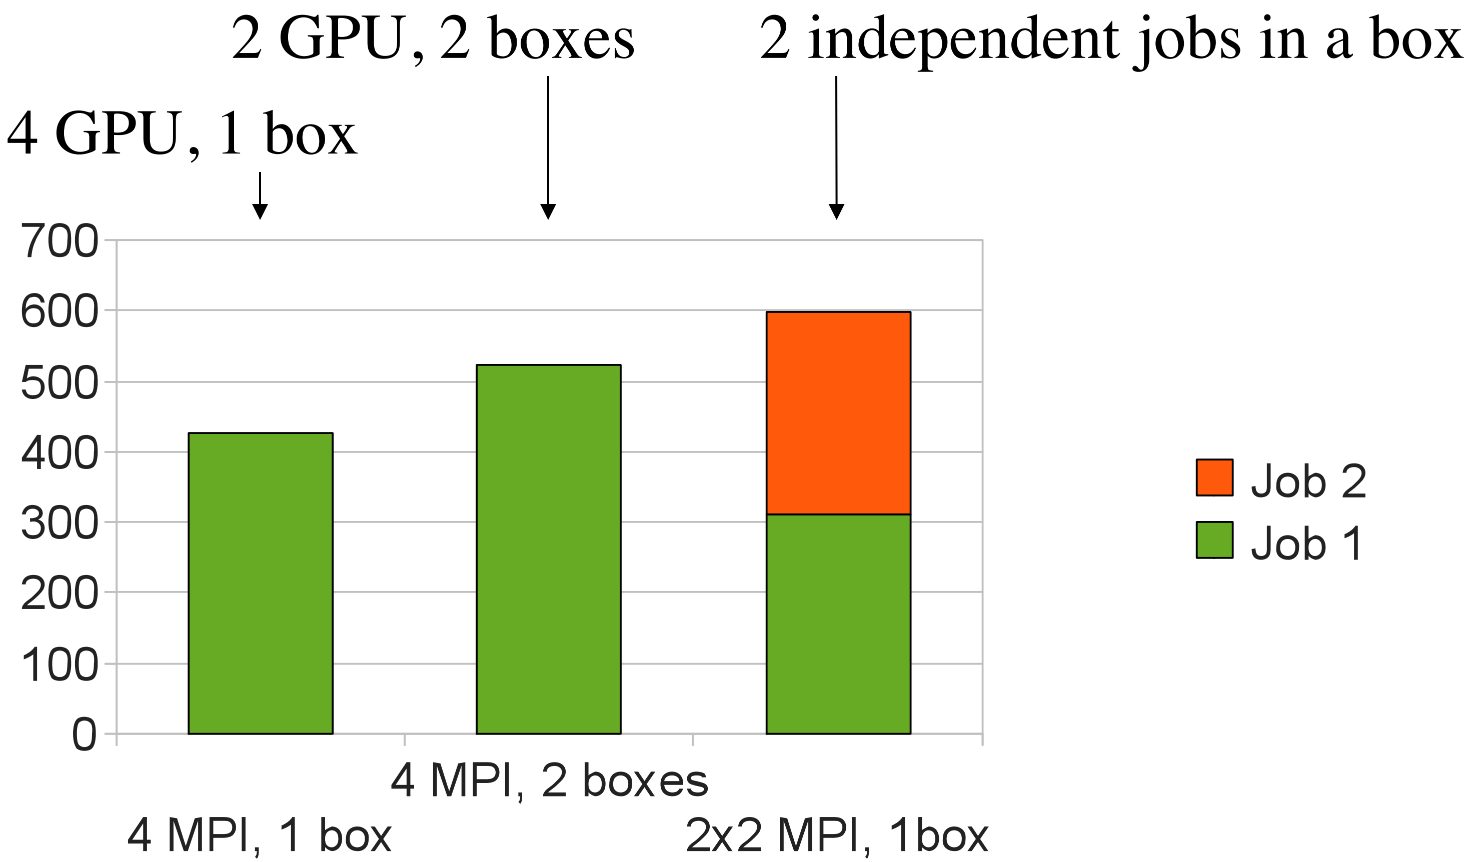

$/Mflops	$0.015	$0.024	$0.01

Price-Performance Tweaks: 32^3x256

4 ways of performing the V=32^3x256 calculation

(all using 8 GPUs, the minimum to hold the problem)

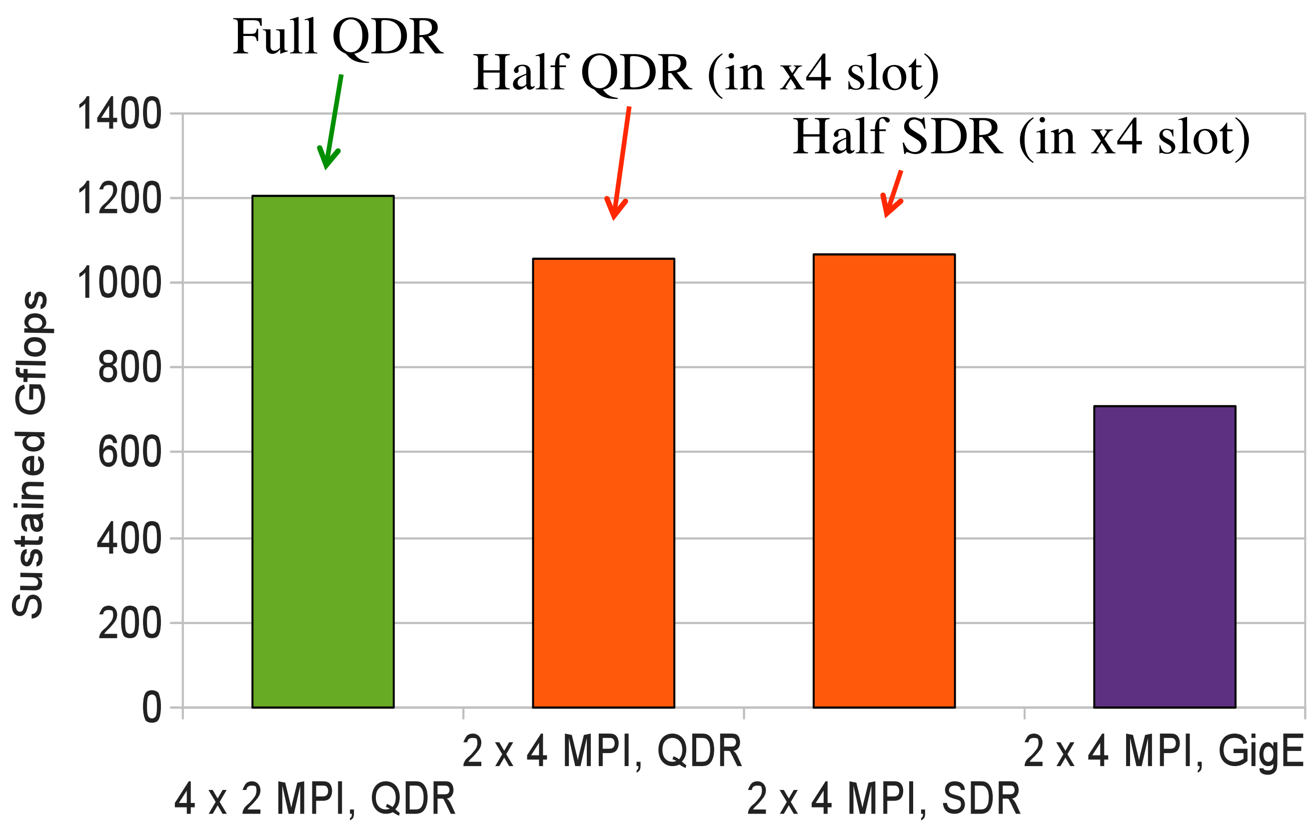

$/Mflops	$0.02	$0.014	$0.012	$0.017

Notes:

Since SDR is as good as QDR for 2 nodes, additional scaling to 5-10 TFlops is feasible using QDR.

All non-QDR GPU nodes will be upgraded with SDR recycled from 6n

A Very Large Resource

500 GPUs at Jefferson Lab

- ★ 190 K cores (1,500 million core hours / year)
- ★ 500 Tflops peak single precision
- ★ 100 Tflops aggregate sustained in the inverter, mixed half / single precision
- ★ 2x as much GPU resource as all cluster resources combined (considering only inverter performance)

All this for only $1M with hosts, networking, etc.

Disclaimer: to exploit this performance, code has to be run on the GPUs, not the CPU (Amdahl's Law problem). This is both a software development problem (see next session), and a workflow problem.

Potential Impact on Workflow

Old Model

2 classes of software

- Configuration generation, using ~50% of all flops

  3-6 job streams nationwide at the highest flops level (capability)

  a few additional job streams on clusters at 10% of capability

- Analysis of many flavors, using ~50% of all flops

  500-way job parallelism, so each job running at <1% capability

New Model

3 classes of software

- Configuration generation, using < 50% of all flops
- Inverter intensive analysis jobs on GPU clusters using ???% of all flops
- Inverter light analysis jobs on conventional clusters using ???% of flops

Summary

USQCD resources at JLab

- 14 Tflops in conventional cluster resources (7n, 9q, 10q)
- 20 Tflops, soon to be 50 Tflops, of GPU resources

  (and as much as 100 Tflops using split precision)

Challenges Ahead

- Continuing to re-factor work to put heavy inverter usage onto GPUs
- Finishing production asqtad and dwf inverters
- Beginning to explore using Fermi Tesla cards with ECC for more than just inverters
- Figuring out by how much to expand GPU resources at FNAL in FY2011

QUESTIONS ?