

Report from the Executive Committee

Paul Mackenzie
mackenzie@fnal.gov

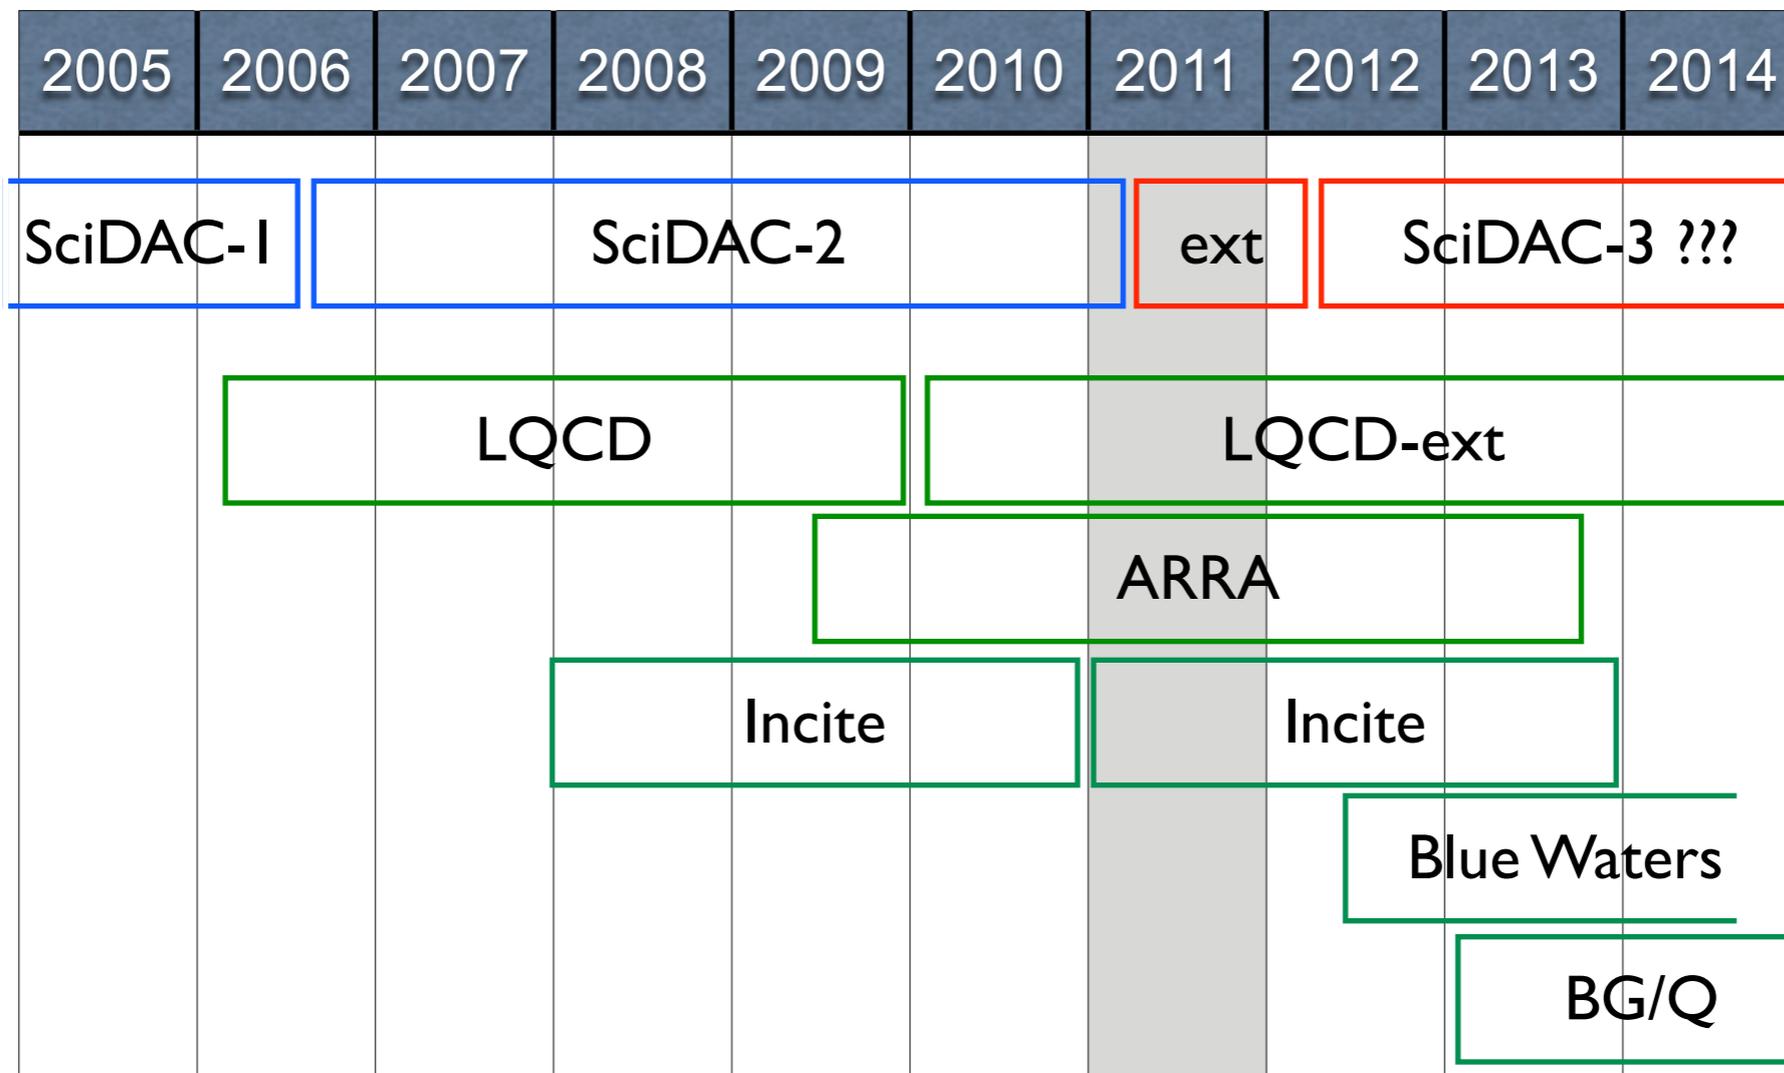
USQCD All Hands' Meeting
JLab
May 6-7, 2011

Outline

- LQCD-ext Project, 2010-2014
- LQCD-ARRA Project
- Incite Grant
- SciDAC-2 Grant, 2006-2011
- Surveys
- Travel Funds
- Coming Peta-scale resources



USQCD projects



The LQCD-ext Project, 2010-2014

- Continues to operate hardware from the LQCD project and before.
 - QCDOC (-2011), Kaon, 7n, and JPsi clusters acquired under LQCD.
- New hardware budget of \$18.15 M over five years.
- Areas of scientific emphasis
 - Fundamental parameters of the Standard Model, and precision tests of it.
 - The spectrum, internal structure and interactions of hadrons.
 - Strongly interacting matter under extreme conditions of temperature and density.
 - Theories for physics beyond the Standard Model.
- The proposal envisioned access to the DOE's leadership class computers as an essential component of the full program.



The LQCD-ext Project, 2010-2014

- First new hardware installation of LQCD-ext happening at Fermilab in FY10/11.
 - Ds1: 245-node, quad-socket, 8-core Infiniband cluster.
 - Ds2 being planned. Current plan: 176 more infiniband nodes+128 Fermi (scientific) GPUs. Will proceed when budget unfrozen. (This week?)
- We're working on metrics for several GPU-related quantities.
 - What fraction of GPU-enabled hardware should be contained in new purchases?
 - Moving target now as GPU use is just ramping up.
 - How should GPUs be related to CPUs in allocations?
 - Charge units could be based on current price of hardware.
 - How should we report the CPU power of a system including GPUs to the DoE?
 - Effective core-hours delivered by GPUs could be based on core-hours that would have been required to do the same calculation on CPUs.



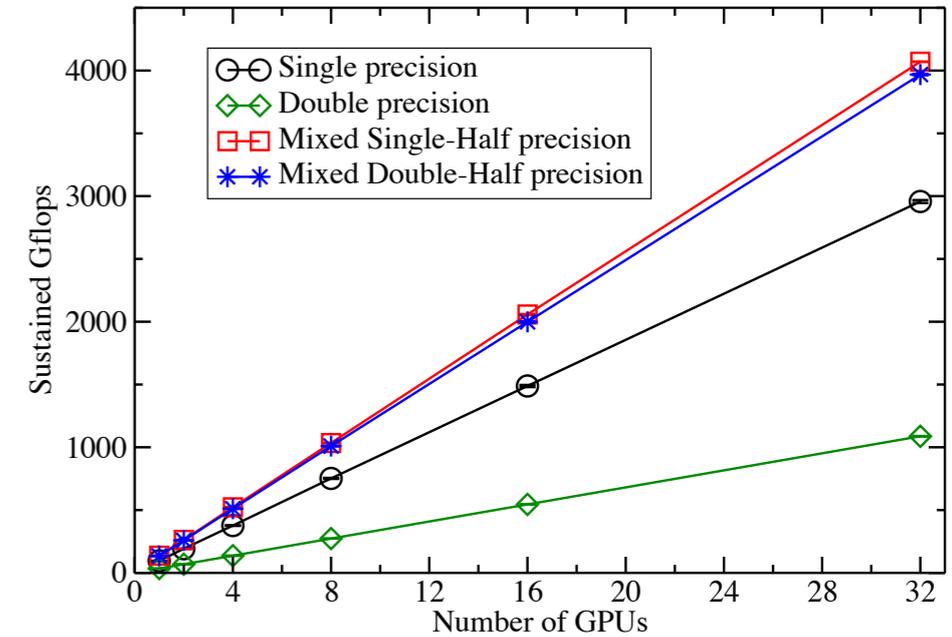
The LQCD-ARRA Project

- Separate project from LQCD-ext;
 - project management is separate and parallel to LQCD-ext.
 - Resources to be managed for science as a coherent whole.
- Sited at JLab, budget of \$4.96 M.
 - Combined budgets for the LQCD-ext and LQCD-ARRA projects around \$23 M, as we originally proposed. (Compared with ~\$9.2 M for LQCD Project.)
- Infiniband clusters 9q and 10q.
 - 512 nodes, dual quad core Infiniband cluster.
- GPUs
 - >500 GPUs of several types.
 - Both Tesla (scientific) and gaming cards



GPU progress

- Much progress with GPU codes this year.
 - Very good scaling with 1-D decomposition.
 - $64^{**}3^{*}128$ run with 4-D decomposition and so-so scaling.
- It's clear that GPUs can handle part of our capacity needs very well. How big is that part?
 - Current plan is for the FY11 Ds2 to be supplemented with a 128-GPU cluster.
 - The project expects to get permission to restart the Ds2 purchase this week.
 - FY12 purchase could include clusters, GPUs, or BG/Q. Information on expected GPU use by June would have maximum usefulness.



Babich, Clark, and Joo, arXiv:1011.0024v1

Japanese use during crisis

- USQCD has offered the Japanese lattice community the use of 10% of its cluster resources during the electricity crisis.
- Until more plants come on-line, supercomputer use is severely curtailed on the eastern grid, including Tokyo and Tsukuba.
- BNL and UK also planning help.
- Four projects will run at Fermilab and JLab.

USQCD Incite Award

- Time on the DOE's leadership class computers, the Cray XT5 at ORNL and the BlueGene/P at ANL, is allocated through the Incite Program.
- Last year, USQCD received a new three-year grant from Jan. 1, 2011 to Dec. 31, 2013.
 - Ours is one of the three largest allocations for 2011. It consists of:
 - 50 M core-hours on the ANL BlueGene/P,
 - 30 M core-hours on the ORNL Cray XT5.
- In 2010 the Cray is being used to generate anisotropic–Clover gauge configurations. The BG/P has been used to generate Asqtad and DWF gauge configurations and to do analysis on those configurations.



USQCD Incite Award

- At ALCF in 2008, USQCD was one of first projects ready to go, only one with three-year program mapped out.
- In one year we accomplished a three-year program of asqtad ensemble generation and the creation of DWF ensembles with a second, fine lattice spacing. We used 359 M core-hours in '08 (~1/3 of BG/P cycles), 279 M in '09, and 187 M in '10.
- Thanks Software Committee: James Osborn, Chulwoo Jung, Balint Joo ...



Allocations and Scientific Priorities

- The Scientific Program Committee (SPC) allocates all USQCD computing resources.
- It is the responsibility of the Executive Committee, in consultation with the SPC and the community, to put forward compelling physics programs in proposals.
- It is the responsibility of the SPC to accomplish the goals of a given proposal, bearing in mind the goals of the funders.
- E.g., charge number 1 to the May 10-11, 2011, LQCD annual review panel is to evaluate:
“The continued significance and relevance of the LQCD-ext project, with an emphasis on its impact on the experimental programs’ support by the DOE Offices of High Energy Physics and Nuclear Physics;”



Allocations and Scientific Priorities

- The Executive Committee will consult with the SPC and the community to create a compelling program of physics for the proposal.
- USQCD does not apply as a collaboration for resources at NERSC or on NSF supercomputers less powerful than Blue Waters. Of course, sub-groups within USQCD can and do apply for these resources.



Executive Committee

- Frithjof Karsch and Julius Kuti replaced Mike Creutz and Claudio Rebbi on the Executive Committee this year.
 - Thanks to Claudio and Mike for their years of service on the EC.
 - Thanks to Frithjof and Julius for being willing to serve.
- Current Executive Committee is Paul Mackenzie (chair), Rich Brower, Norman Christ, Frithjof Karsch, Julius Kuti, John Negele, David Richards, Steve Sharpe, and Bob Sugar.



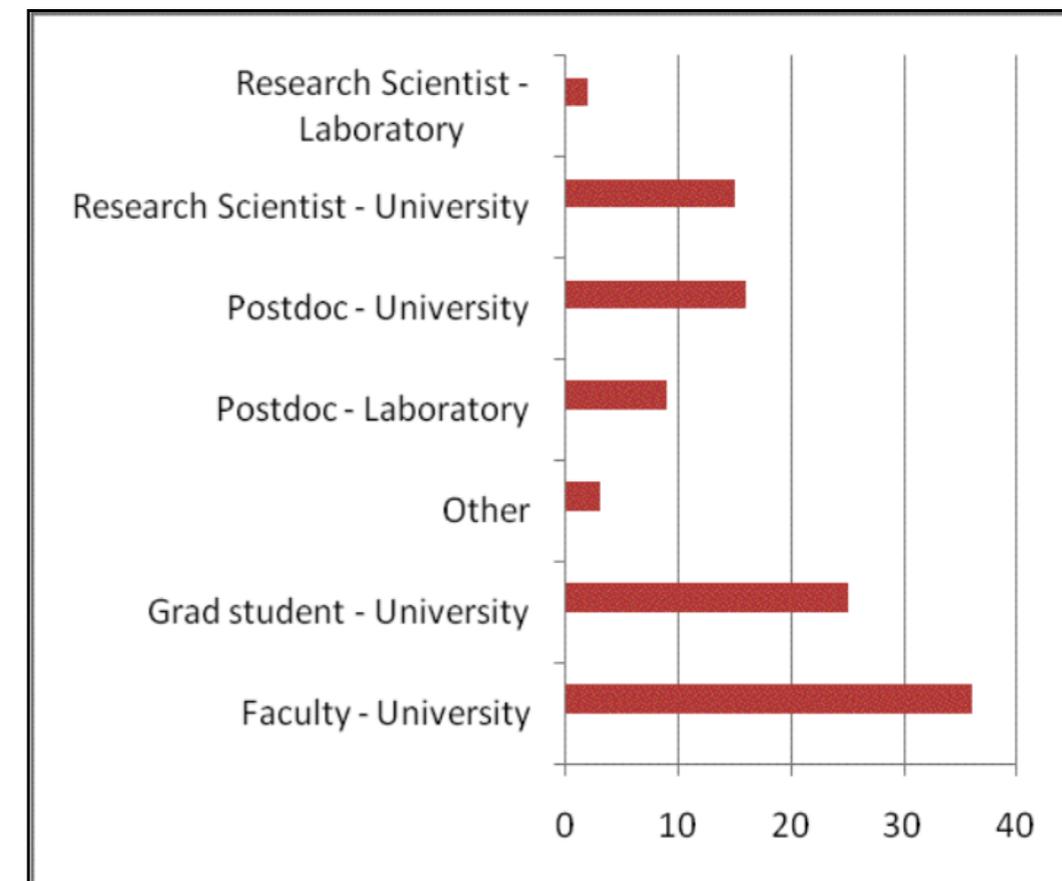
SciDAC-2 Grant

- Grant runs from 2006-2011. A one-year extension is being finalized now.
- We received \$2,359,000 last year.
- Recent efforts have focused on USQCD codes for the BlueGene/P and Cray XTs as well as new software tools for workflow, visualization and methods to meet the challenges of many-core hardware and multi-level algorithms. Rich Brower will give an overview of these activities for the Software Committee.
- One-year extension of SciDAC-2, 2011-2012 in the works.
- SciDAC-3 is being discussed to begin in 2012.
 - HEP and NP understand that SciDAC is essential for effective use of hardware resources and expect it to continue. Discussions now underway between HEP, NP, and ASCR.
 - Executive Committee and Software Committee members made a trip to Office of Science headquarters in Germantown in March to emphasize this. It seemed that our message was getting across.



Membership, demographic, and user surveys

- DoE asks the collaboration to take regular surveys on various topics.
 - More this year than usual.
 - We understand that this is a pain in the neck, but the information is useful for the DoE.
- DoE has asked the project to keep regularly updated **demographic information** on our field. New postdocs and students, new faculty members is a measure of the health of a field.



Membership, demographic, and user surveys

- New membership list and member email list.
 - Announcement will be sent out this week.
- Users survey.
 - DoE mandates that the project team take a user survey every year.
 - Only way for DoE to judge if users are happy with project management.
 - Logging in to a USQCD computer during the year constitutes an agreement to complete the survey.
 - Can be done rapidly.



Travel Funds

- As was indicated at last year's All-hands Meeting, limited travel funds are available for use by USQCD members.
 - Main priorities are USQCD Collaboration business, such as traveling to another USQCD institution to work on SciDAC software or USQCD hardware, or representing USQCD at an ILDG meeting.
- Those wishing to make use of these funds should send email to mackenzie@fnal.gov.
- Highest priority will be given to junior members of USQCD.



Coming peta-scale hardware

We expect to have access to several very large resources in the next few years.

- IBM Blue Waters at NCSA
- IBM BG/Q at Argonne
- Cray with GPU accelerators at Oak Ridge



Blue Waters, NCSA

- Expected mid-2012? 300,000 cores, eight-core POWER7 CPUs.
- Acceptance tests: *close to* 1 petaflop delivered on scientific applications including MILC asqtad configuration generation.
- Chroma and MILC are running on prototype hardware (Gotlieb, Joo, ...).
- Not much known yet as of now about how the NSF intends to allocate Blue Waters.
- As we learn more, we'll have to figure out how to apply in a way that maximizes our physics goals.



NSF PRAC Proposal for Blue Waters

- USQCD has submitted a proposal to Petascale Computing Resource Allocations (PRAC). We requested:
 - Travel funds to be used in the development and optimization of software for Blue Waters.
 - Early access to information regarding Blue Waters' architecture.
 - An early allocation of time on Blue Waters.
- The USQCD proposal has received a grant of \$40,000 for travel associated with code development.
- Nondisclosure agreements are still being negotiated between NCSA and the universities.



BG/Q at Argonne

- Early science starts early 2013. The ALCF's stated requirements for the 10 petaflops system include approximately 0.75 million cores with 16 cores per node.
 - <http://www.alcf.anl.gov/collaborations/early.php>.
 - USQCD through Columbia involved in design. (Peter Boyle dslash was the first realistic code running on simulator. Chulwoo Jung working on higher level code which could serve as basis for QLA, QDP, ... on the BG/Q.)
- Early science proposal.
 - Presented definite plan to do HISQ and DWF configuration generation, indicated that we would like do other projects such as QCD thermodynamics and BSM.
 - Argonne is aware that we couldn't be completely definite about what science will have highest priority two years in the future.
 - Awarded 150 M core-hours.
- Prototype BG/Q hardware at BNL late this year.



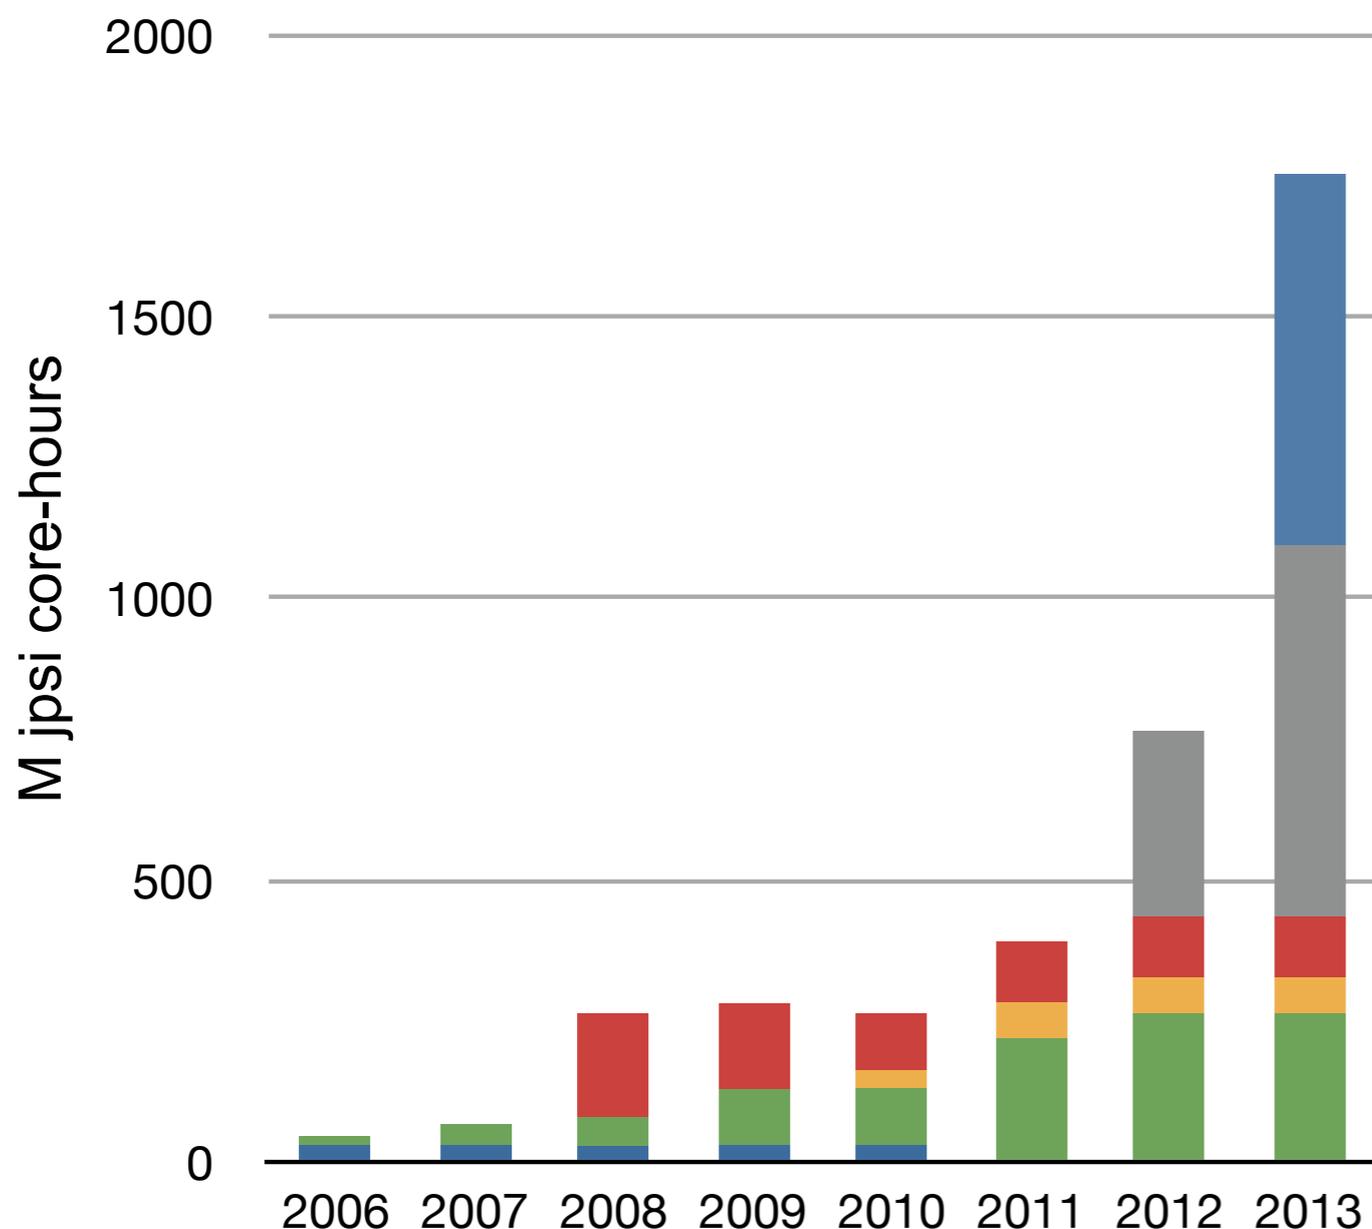
Oak Ridge 2012 machine, Titan

- Massively parallel NVidia Tesla GPUs.
 - Yikes.
- 20 PF peak.
- Possible collaboration with NVidia to prepare for it.
 - NVidia has decided that lattice QCD is an application they should support.



History of USQCD resources

USQCD Resources



Computing resources for calculations two or three years from now could be an order on magnitude larger than for current calculations.

USQCD ought to have a plan for spending 10% of expected US resources for 3 years.

It's possible that, as happened on the ALCF BG/P, we could get 30% of the resources for the first year (rather than 10%).

GPUs numbers are a lower bound and underestimate. For LQCD, includes no 2012/13 capacity hardware. For Incite, does not include Oak Ridge Titan. Assumes 10% of ALCF and OLCF; fraction could be much larger.

- QCDOC
- Clusters
- GPUs
- ALCF M jpsi c-h
- OLCF M jpsi c-h
- Blue Waters jpsi c-h
- BG/Q jpsi c-h



Extra or old slides



Hardware goals by fiscal year

Fiscal Year	Dedicated Hardware (Tflop–Years)	Leadership Class Computers (Tflop–Years)
2010	35	30
2011	60	50
2012	100	80
2013	160	130
2014	255	210
Total	610	500

1 Tflop-year = 3.5 M 6n node-hours

Computing resources from the use of dedicated hardware (column 2) and leadership class computers (column 3) needed to carry out our scientific program by fiscal year. Computing resources are given in Tflop–Years, where one Tflop–Year is the number of floating point operations produced in a year by a computer sustaining one teraflop/s.

Goals envisioned in the LQCD-ext proposal.



SciDAC-2 Grant

Table of Revised Budgets for SciDAC Proposal National Computational Infrastructure for Lattice Gauge Theory

Table 1 below shows the total budget for each participating institution in each of the five years of the proposed grant. Note that all funds for hardware research and development have been eliminated. The FY2006 budgets cover the six month period between September 15, 2006 and March 14, 2007. All other budgets are for twelve month periods.

Institution	FY06	FY07	FY08	FY09	FY10
BNL	218	362	378	390	406
FNAL	249	442	456	472	485
JLab	258	458	472	484	497
Boston U.	88	183	191	198	206
DePaul U.	32	66	68	70	72
IIT	15	30	30	30	30
Indiana U.	25	51	52	54	55
MIT	113	235	244	254	264
U. Arizona	25	51	53	54	55
U. North Carolina	55	113	116	119	122
UC Santa Barbara	15	30	30	30	30
U. Utah	27	55	56	58	60
Vanderbilt U.	37	75	76	76	77
Total	1,157	2,151	2,222	2,289	2,359

Table 1: Institution and Total Budgets in \$1,000

Travel Funds

- The Executive Committee believes that travel funds should be used for activities that directly address or report on USQCD activities. Some examples are:
 - Traveling to another USQCD institution to work on SciDAC software or USQCD hardware.
 - Representing USQCD at an ILDG meeting.
 - Attending a USQCD sponsored conference or summer school.
 - Attending a topical workshop to report on results obtained with USQCD computing resources.
- We cannot afford to support travel to Lattice Meetings, or to meetings of sub-groups within USQCD.

