

Report from the Executive Committee

Paul Mackenzie
mackenzie@fnal.gov

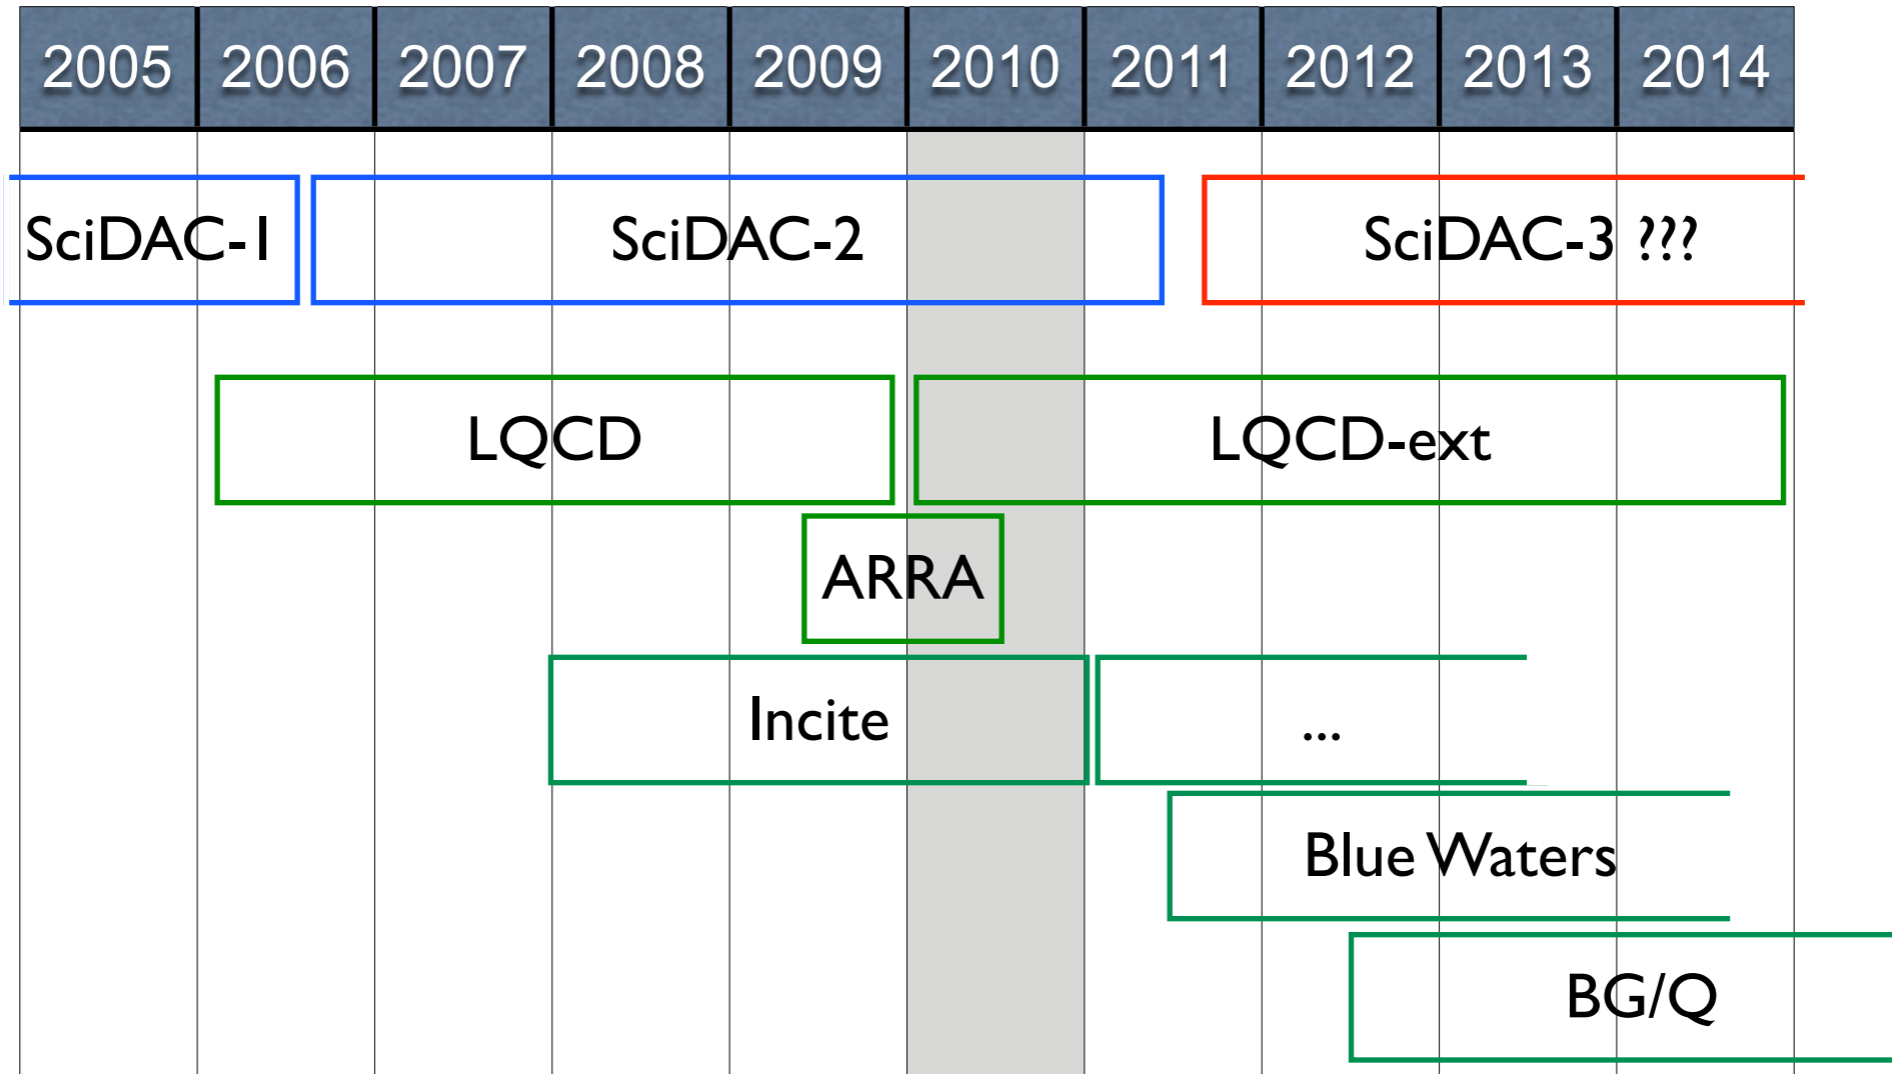
USQCD All Hands' Meeting
BNL
April 16-17, 2010

Outline

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



The shape of things to come



The LQCD-ext Project, 2010-2014

- Continues to operate hardware from the LQCD project and before.
 - QCDOC (-2011), 6n, Kaon, 7n, and JPsi clusters acquired under LQCD.
- New hardware budget of between \$1.60-\$2.46 M/year, operations budget of \$1.14 M in 2010 rising to \$1.64 M in 2014.
 - Total budget of \$18.15 M.
- 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 planned for Fermilab later this year.
 - Combines FY10 and FY11 hardware funds.
 - Will contain a significant component of GPUs. 11+12 TF + ???
- We need metrics for several GPU-related quantities.
 - What fraction of GPU-enabled hardware should be contained in new purchases.
 - Must take into account fraction of program that can use GPUs, GPU cost, ...
 - How should we report the CPU power of a system including GPUs to the DoE?
 - How should GPUs be allocated?
 - Separately from Infiniband systems? Or...
 - In service units that equate a core-hour on a new Infiniband system with a cost-equivalent amount of hardware on a new GPU-enabled system? Or...
 - ...?



The LQCD-ARRA Project

- Separate project from LQCD-ext; resources to be managed 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.)
- Phase I hardware in operation for several months
 - 10 racks of 32 nodes, dual quad core Infiniband cluster.
 - Several types of GPU-enabled nodes.
 - 65 nodes: 16 dual GPU, QDR Infiniband, 10 dual GPU, 37 quad GPU, 2 R&D.
- Phase II hardware
 - 7 racks of 32 node Infiniband cluster now arriving.
 - New GPU hardware just purchased. (Chip Watson's talk.)



USQCD Incite Award

- Time on the DOE's leadership class computers, the Cray XT4 at ORNL and the BlueGene/P at ANL, is allocated through the Incite Program. USQCD has received a three year grant from the Incite Program from Jan. 1, 2008 to Dec. 31, 2010. Ours is the largest allocation for 2010. It consists of:
 - 67 M core-hours on the ANL BlueGene/P,
 - 40 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, mostly of Early Science time, ~1/3 of BG/P cycles in 2008.
- Thanks James Osborn and Software Committee.
- At ALCF in 2009, we had access to a low priority queue instead of an ES grant.
- Used 268 M core-hours in 2009.
- We will submit a new Incite proposal for 2011-13 at the end of June.



Annoying ALCF low-priority policy

- In 2010, low-priority queue went away. We were required to use up our allocated time before continuing with zero-priority time.
- The allocated time has now been used up (somewhat bumpily) and we are running on zero-priority time again.
- The SPC is now working on the plan for running in this mode for the rest of the year.



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 April 29-30, 2010 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 supported by the Offices of High Energy and Nuclear Physics of the DOE;”



Allocations and Scientific Priorities

- A new Incite proposal will be submitted in June of this year. In this and future proposals, 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.



SciDAC-2 Grant

- Grant runs from 2006-2011. On January 8-9, 2009, we received a favorable mid-term review.
- We received \$2,359,000 this year, and we are getting a small cost-of-living increase every 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.
- Grant ends in 2011.
 - SciDAC-3 is being considered at DoE. HEP is happy with the project and expects it to continue. Discussions now underway between HEP, NP, and ASCR.



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

- As was indicated at last year's All-hands Meeting, limited travel funds are available for use by USQCD members.
- 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.



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.



Coming peta-scale hardware

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

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



Blue Waters, NCSA

- Expected 2011? 25,000 eight-core POWER7 CPUs (Wikipedia).
- Acceptance tests: 1 petaflop delivered on scientific applications including MILC asqtad configuration generation.
- Nothing is known 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

- Due late 2012? 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.)
- Early science time in late 2012 (?).
 - Writing proposal now.
 - Argonne is aware that we can't be definite about what science will have highest priority 2 1/2 years from now. We will say that we definitely plan to be ready to do QCD configuration generation, as we were for the BG/P in 2008, and we will give examples of other types of projects that we expect to have high priority.

History of USQCD resources

Year	Total (TF yrs)	Clusters	QCDOC	Incite-Tot (TF yrs)	Blue Waters	BG/Q	Comments
2006	2006	7.2	2.7	4.451			
2007	2007	10.2	5.8	4.372			
2008	2008	40.3	7.7	4.507	28.09		
2009	2009	42.8	15.4	4.534	22.85		
2010	2010	36.4	26.4	4.5	5.49		Projected
2011	2011.5	152.9	47.4	5.49	100.00		"
2012	2012.5	252.9	47.4	5.49	100.00	100.00	"


These figures should not be used in situations which might present a risk to human life or property.

1 TF yr = 3.90 M 6n nh
 = 12.21 M bgp ch
 ARRA 5+6 TF
 Fermi 10 11+10 TF

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%).



Review of LQCD-ext and LQCD-ARRA projects, April 29-30.

The review committee of the 2009 LQCD review recommended that USQCD establish policies on sharing of lattice data and on data curation.

- Data curation strategy.
 - The Executive Committee proposes that USQCD adopt the policy that all large gauge ensembles be stored in two geographically separate locations.
- Public data release policy.
 - USQCD requires that groups generating large gauge ensembles share them within the collaboration.
 - The Executive Committee proposes that USQCD adopt the policy that all ensembles of gauge configurations generated with USQCD resources be made publicly available through the International Lattice Data Grid no later than six months after the first publication in a refereed journal of a paper that makes use of them.



Extra 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.

