

Report from the Executive Committee

Paul Mackenzie
mackenzie@fnal.gov

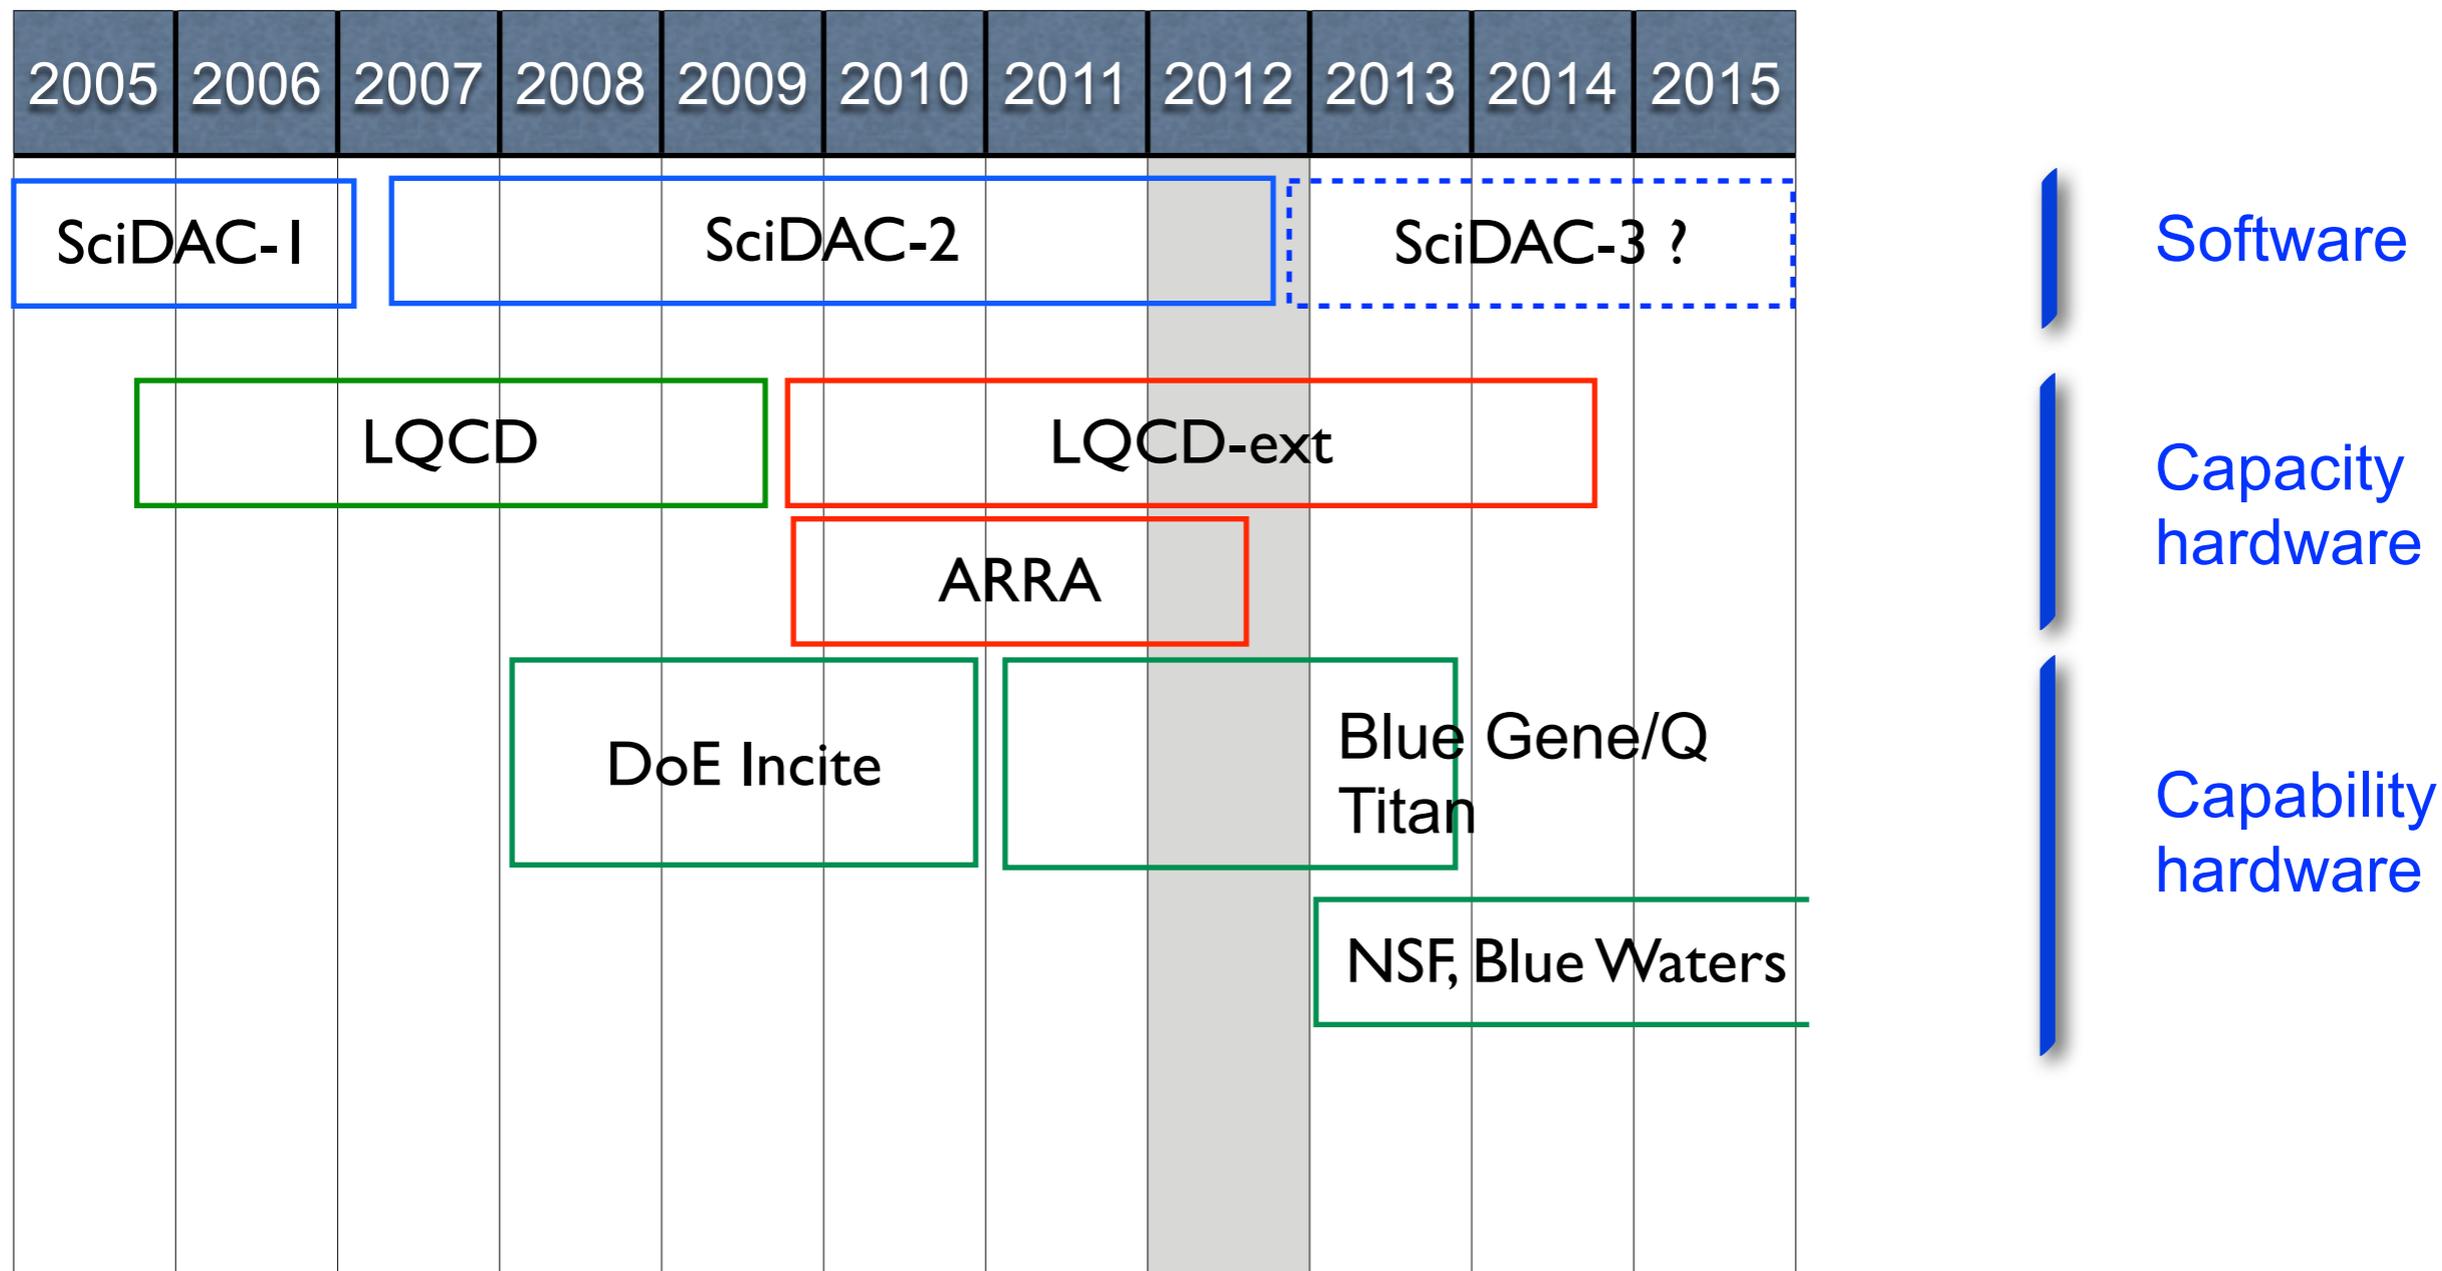
USQCD All Hands' Meeting
Fermilab
May 4-5, 2012

Outline

- LQCD-ext Project, 2010-2014
- LQCD-ARRA Project, 2009-2012
- Current INCITE Grant
- SciDAC-2 Grant, 2006-2011
- Surveys
- Travel Funds
- Coming INCITE and NSF resources



USQCD projects



The LQCD-ext Project, 2010-2014

- Project 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

- 2010/11 hardware at Fermilab.
 - Ds: 421-node, 13,440-core, quad-socket, 8-core Infiniband cluster.
 - Dsg: 76 nodes, 152 Fermi GPUs. (Don's talk.)
- 2012 hardware at JLab.
 - 12s: 212 nodes, each dual-socket eight-core 2.0 GHz Intel,
 - additional mixture of nodes and Kepler (we hope) GPUs. (Chip's talk.)
- 2012 hardware at BNL
 - Use of 10% of a Blue Gene/Q rack at BNL. (Bob M.'s talk.)
- 2012 Project annual review in two weeks at BNL.
 - We need from each physics project PI
 - updated publication lists,
 - updated project web pages.

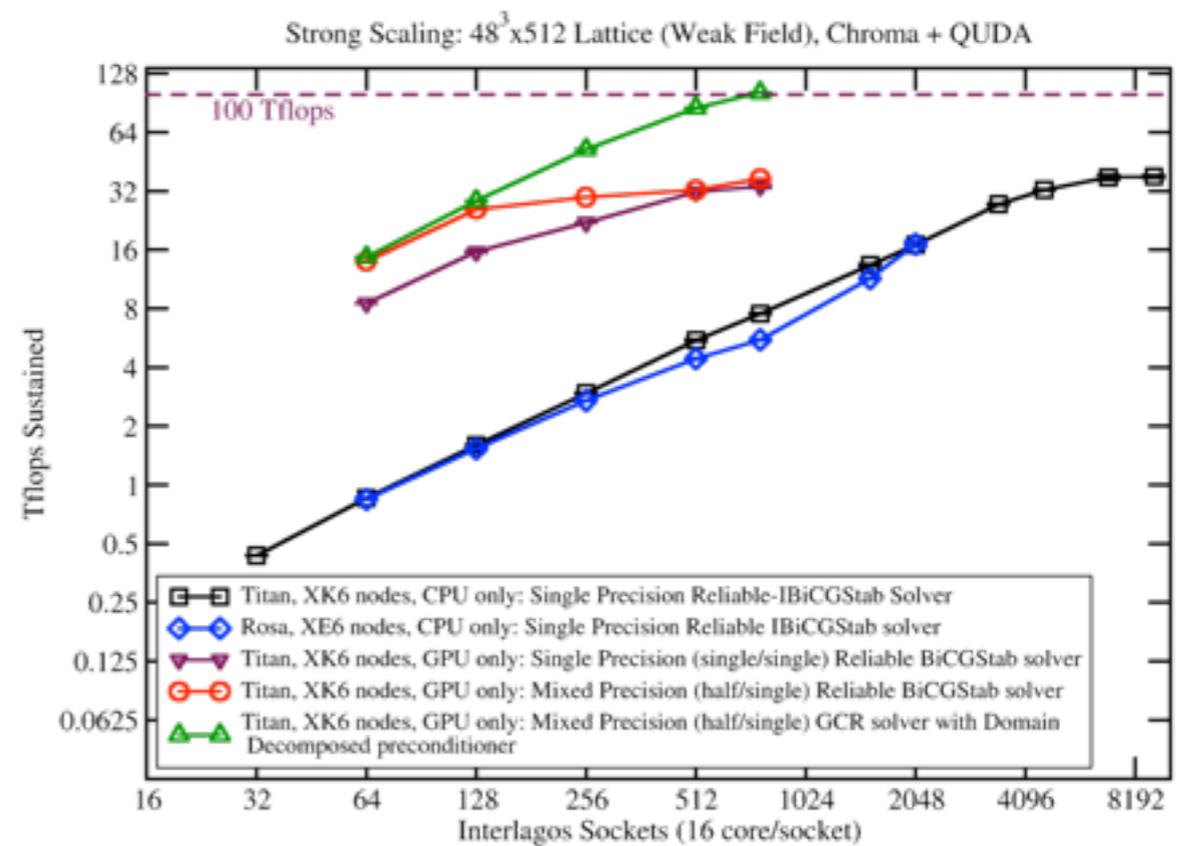


The LQCD-ARRA Project

- Separate project from LQCD-ext;
 - project management have been separate and parallel to LQCD-ext.
 - Resources have been managed for science as a coherent whole.
 - Project will be brought to close in 2012, operations folded into LQCD-ext.
- 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.
 - ~500 nodes, dual quad core Infiniband cluster.
- GPUs
 - 480 GPUs of several types.
 - Both Tesla (scientific) and gaming cards



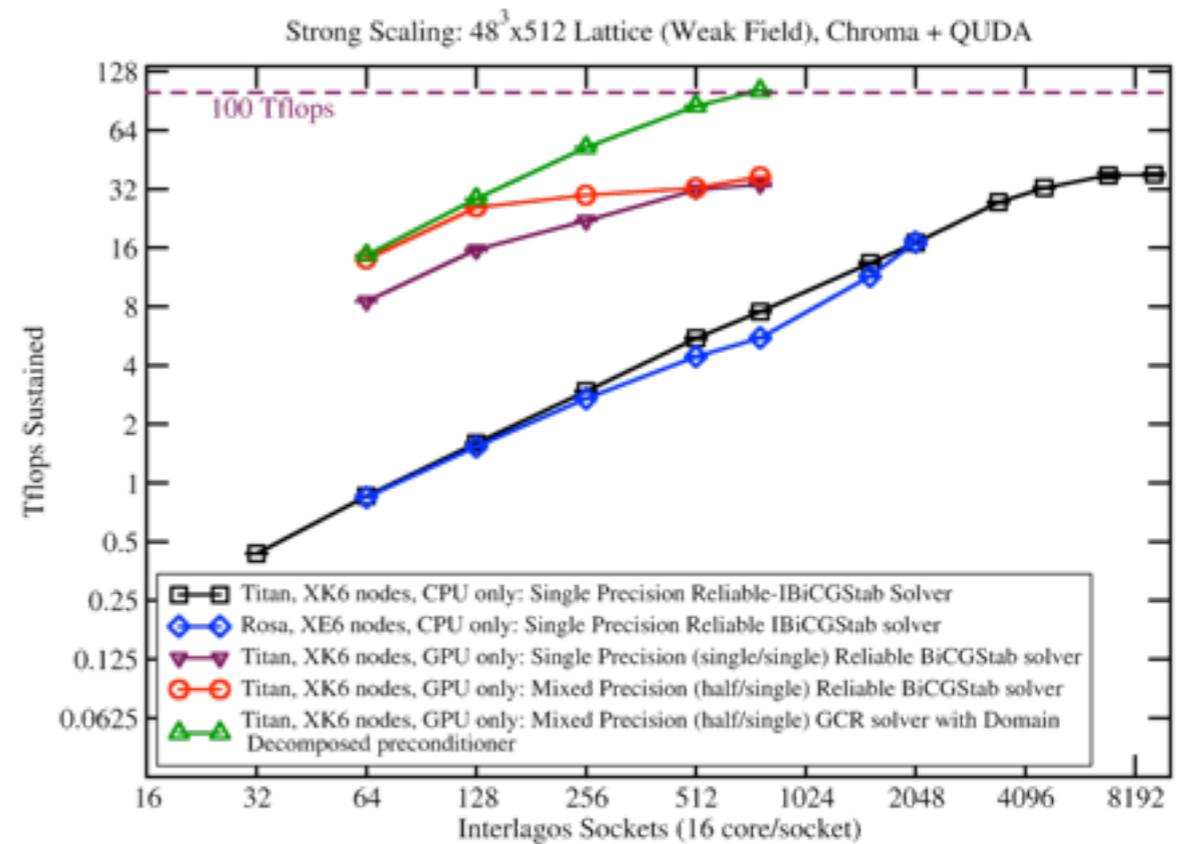
GPU progress



Clark and Joo, ACS Symposium, 2012

- Much progress with GPU codes this year.
 - Decent strong scaling on $48^3 \times 512$ run with 4-D decomposition.
- It's clear that GPUs can handle part of our capacity needs very well. How big is that part?
 - Current plan is for the FY12 12s to be supplemented with additional GPUs.
 - FY13 purchase could include clusters, accelerated clusters, or BG/Q. Benchmarking information by June would have maximum usefulness.

GPU progress



Clark and Joo, ACS Symposium, 2012

- The Project needs community input 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.



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.
- USQCD has a three-year grant from Jan. 1, 2011 to Dec. 31, 2013.
- Ours is one of the three largest allocations for 2012. It consists of:
 - 50 M core-hours on the ANL BlueGene/P, plus zero-priority time (130 M ch in 2012),
 - 46 M core-hours on the ORNL Cray XT5.
- In 2011 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.
- New INCITE-managed resources coming in 2013 (later).



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, 187 M in '10, 180 M in '11.
- 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 16-17, 2012, [LQCD annual review](#) panel is as usual 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.



Committee Members

- Current Executive Committee is Paul Mackenzie (chair), Rich Brower, Norman Christ, Frithjof Karsch, Julius Kuti, John Negele, David Richards, Steve Sharpe, and Bob Sugar.
- Current Scientific Program Committee is Robert Edwards (chair), Simon Catterall, Martin Savage, Taku Izubuchi, Doug Toussaint, Peter Petreczky, Ruth Van de Water



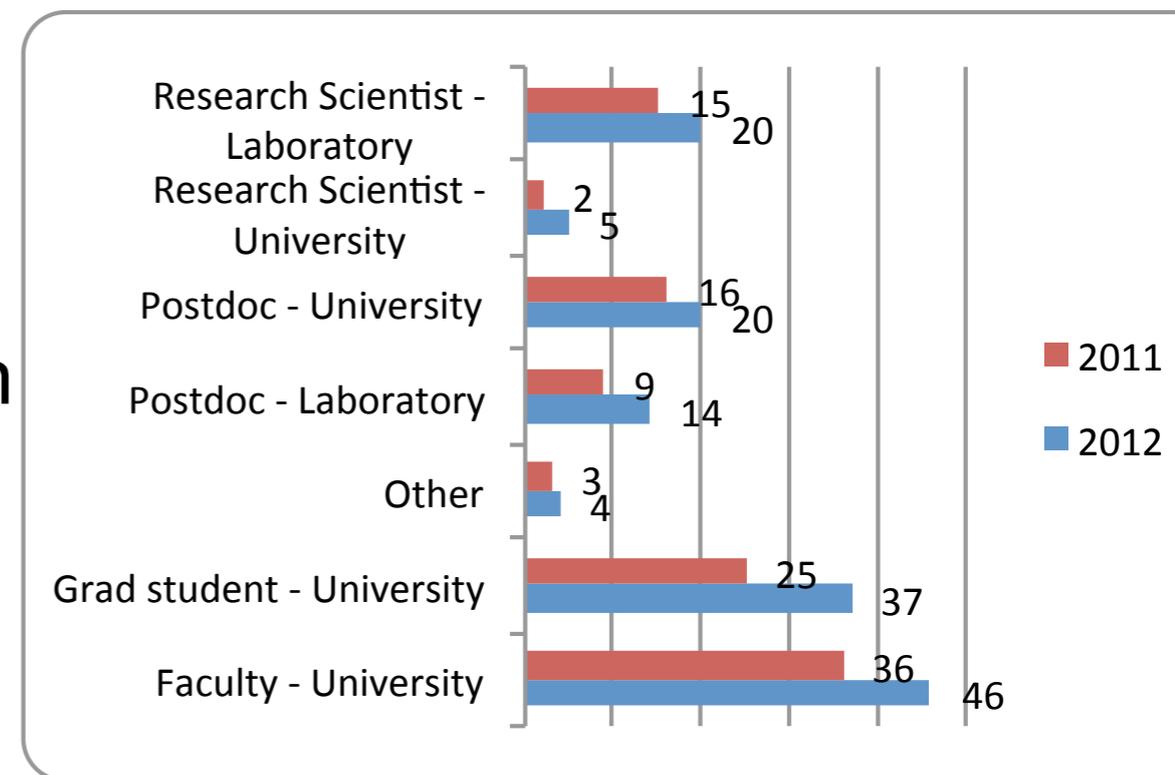
SciDAC-2 Grant

- Grant runs from 2006-2012.
- We received \$1,817,000 this year.
- Recent efforts have focused on USQCD codes for the BlueGene/P and Cray XTs as well as methods to meet the challenges of GPU and many-core hardware and multi-level algorithms. Rich Brower will give an overview of these activities for the Software Committee.
- SciDAC-3 beginning in late 2012 is under review. Project was split into an HEP project and an NP project. News is expected soon.



Membership, demographic, and user surveys

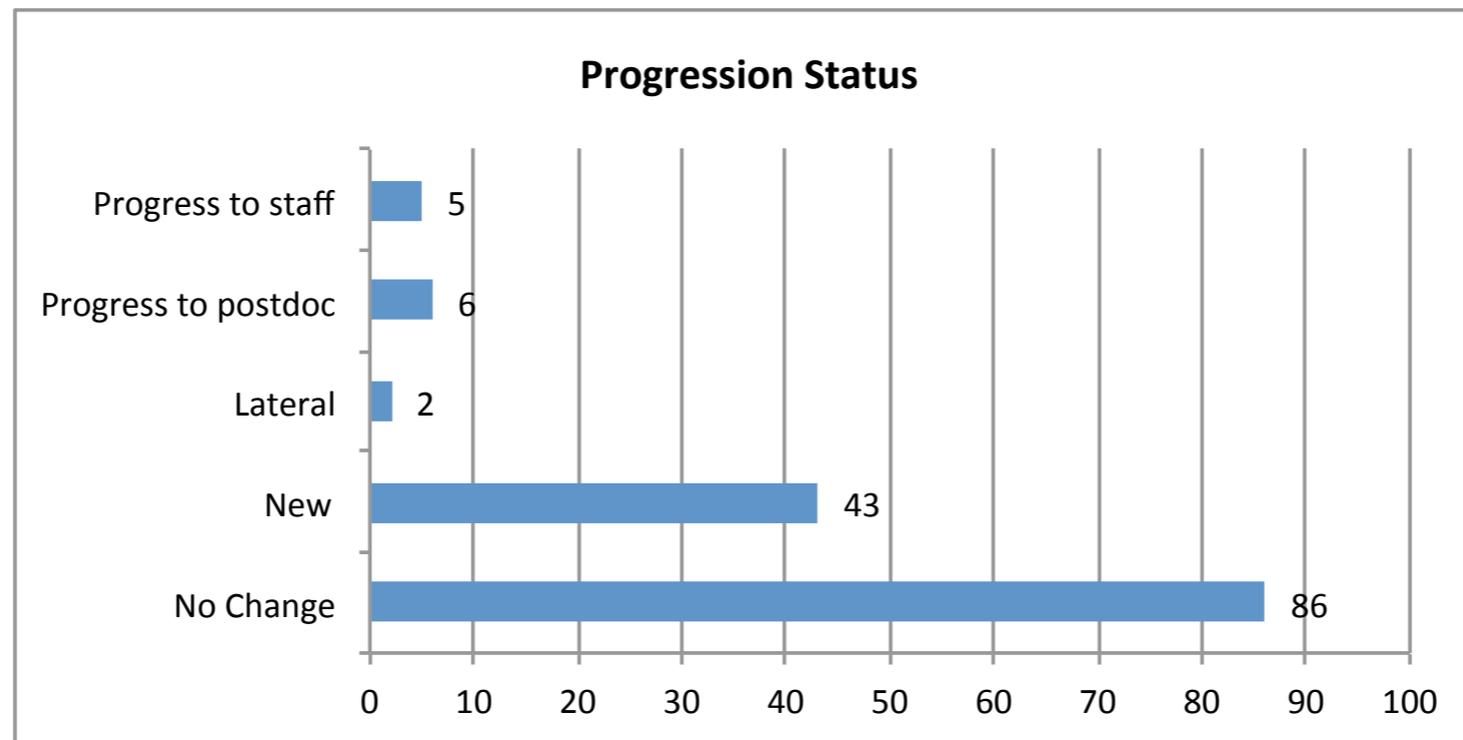
- DoE asks the collaboration to take regular surveys on various topics.
 - We understand that this is a pain in the neck, but the information is important to the DoE.
- DoE has asked the project to keep regularly updated **demographic information** on our field. New postdocs and students, new faculty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



Demographic progression

Our project managers at DoE have expressed particular interest in the progress and promotions of young people.

Our information collected so far is clearly incomplete; we will be interacting with you in the next week to try to get more complete information before the hardware review.



Membership, demographic, and user surveys

- Membership list and member email list.
- 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 year.

- INCITE:
 - “Mira”, IBM BG/Q at Argonne
 - “Titan”, Cray with GPU accelerators at Oak Ridge
- The new INCITE Resources
 - are for projects that can run on 20% of the machine (a partition > full BG/P),
 - projects that can't be done on any smaller machine (like BG/P or Jaguar).
- NSF:
 - IBM Cray Blue Waters at NCSA



BG/Q at Argonne

- 48 racks, 48 K 16-core nodes.
- INCITE will allocate 0.768 B - 3 B ch in 2013, starting perhaps spring-summer 2013.
 - USQCD through Columbia involved in design. (Peter Boyle dslash was the first realistic code running on simulator. Chulwoo, James Osborn, ... working on higher level codes, QLA, QDP, ... on the BG/Q.)
- Early science proposal.
 - Awarded 150 M core-hours, beginning late 12.
- Prototype BG/Q hardware at BNL and Argonne now.



Oak Ridge 2012 machine, Titan

- Upgraded Jaguar, + Fermi (→Tesla) GPUs.
 - 299,088 AMD Interlagos cores + 14,592 GPUs.
 - Currently, TitanDEV has 960 nodes with Fermi GPUs.
- INCITE will allocate 2 B ch for 2013.
- Formal collaboration with NVIDIA to prepare for it.
 - Mike Clark, Ron Babich→NVIDIA.
 - NVIDIA has decided that lattice QCD is an application they should support.



NSF: Blue Waters at NCSA

- AMD “Interlagos” nodes, >380,000, >3,000 GPUs.
- Chroma and MILC have been run.
- Trial projects going on a small amount of early science time.

- Not much 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.

