LQCD-ext Project
2014 Annual Review

Overview and USQCD Collaboration Management

Paul Mackenzie
Fermilab
mackenzie@fnal.gov

For the USQCD Collaboration
http://www.usqcd.org

LQCD-ext Project 2014 Annual Review
Fermilab
May, 2014

Synopsis

- Final annual review of the LQCD-ext lattice computing hardware project, 2010-14.
  - Hardware is located at BNL, JLab, Fermilab.
  - Projects funded jointly by DoE's offices of HEP and NP.
    - LQCD-ext, total budget $18.05 M - under review today.
      - Operations of our sister ARRA project, LQCD-ARRA, 2009-12, were folded into the LQCD-ext Project. LQCD-ARRA total budget $4.96 M.
      - Follow-up project, LQCD-ext II, is now under review.
- The LQCD Project is one of several hardware and software efforts overseen by the USQCD Collaboration.
- USQCD is a collaboration consisting of most US lattice gauge theorists. Its purpose is to develop the software and hardware infrastructure required for lattice gauge theory calculations.

Plan of talks

(Detailed schedule at http://projects.fnal.gov/lqcd/reviews/May2014Review/agenda.shtml .)

May 15

08:30	Executive session (45 min)	
09:15	Welcome (15 min) – *Nigel Lockyer*	
09:30	Logistics and Introductions (5 min) – *Bill Boroski*	
09:35	LQCD-ext Overview & USQCD Governance (50 min) - *Paul Mackenzie*	Overview
10:25	Break (15 min)	
10:40	USQCD Allocation Process (20 min) – *Robert Edwards*	Scientific achievements
11:00	Science Talk 1: Lattice QCD for HEP: Standard-Model Parameters and Matrix Elements (40 min) – *Ruth Van de Water*	Scientific achievements
11:40	Science Talk 2: Lattice Field Theory Beyond the Standard Model (20 min) – *Rich Brower*	Scientific achievements
12:00	Lunch / Executive Session	Scientific achievements
1:00	Science Talk 3: Hadron Spectroscopy, Structure and Interactions (30 min) – *Martin Savage*	Scientific achievements
1:30	Science Talk 4: High Temperature/Density QCD (30 min) – *Frithjof Karsch*	Scientific achievements
2:00	LQCD-Ext Project: Management and Performance (60 min) – *Bill Boroski*	LQCD Project management and technology
3:00	Coffee Break (20 min)	LQCD Project management and technology
3:20	LQCD-Ext: Technical Performance of FY2013 BG/Q Deployment (20 min) - *Bob Mawhinney*	LQCD Project management and technology
3:40	LQCD-Ext: Technical Performance of FY2013 Cluster Deployment (20 min) – *Amitoj Singh*	LQCD Project management and technology
4:00	LQCD-Ext: FY2014 Hardware Plan & Cluster Deployment Status (30 min) – *Don Holmgren*	LQCD Project management and technology
4:30	LQCD-Ext: Future Planning - FY2015 and Beyond (15 min) – *Bill Boroski*	LQCD Project management and technology
4:45	Executive Session (60 min)	
5:45	Committee request for additional information - *Committee/Project Leadership*	
6:00	Adjourn	
7:00	Dinner	

The USQCD Collaboration

- Organizes computing hardware and software infrastructure for lattice gauge theory in the US.
- Represents almost all of the lattice gauge theorists in the US; ~ 163 people.
    - ~ 100 participating in physics proposals this year.
- Physics calculations are done by smaller component collaborations within USQCD:
    - Fermilab, HotQCD, HPQCD, HadSpec, LHPC, LSD, MILC, NPLQCD, RBC, ...
    - These are the core entities of the US lattice community.

USQCD timeline

- USQCD formed in 1999.
- LQCD-ext Project capacity hardware grants from HEP and NP.
  - Being reviewed today.
  - Installed at JLab, Fermilab, and BNL.
- INCITE grants since 2008.
  - For our largest-scale jobs; jobs that can't be done on smaller computers.
- SciDAC software grants since 2001.
  - Essential for making effective use of Leadership Computing Facilities and our dedicated hardware, and for accomplishing our physics objectives.

USQCD Collaboration

LQCD-ext Project

Executive Committee:
Paul Mackenzie (chair), Rich Brower, Norman Christ, Frithjof Karsch, Julius Kuti, John Negele, David Richards, Martin Savage, Bob Sugar

Federal project director:
John Kogut (HEP)
Federal project monitor:
Kawtar Hafidi (NP)

Contract project manager:
Bill Boroski
Associate project manager:
Rob Kennedy

Scientific Program Committee:
Robert Edwards (chair)

Software Committee:
Rich Brower (chair)

Green: present today.

The USQCD collaboration is funded through SciDAC, through the LQCD project, and through base HEP and NP funds at BNL, Fermilab, and JLab.

USQCD collaboration web page: http://www.usqcd.org.

USQCD total hardware resources

USQCD resources in January, 2014, normalized to JPsi core-hours.

	2013/14 lattice QCD resources of USQCD and its contituents					
	Red: DoE resources	Subtotal	Core hours (M jpsi) allocated	Core hours (M jpsi) zero-priority	Comment	Policies
Lattice QCD	LQCD-ext, conventional.		411			
	LQCD-ext, GPUs		646		GPU-based*	
		1057				
Leadership-class	OLCF (DoE), Titan		182		GPU-based	Jobs using < 20% of the machine (nodes) don't make it through the queue.
	ALCF (DoE), Mira		325	413		Proposed calculations using < 4-8 racks (64-128 K cores) are discouraged. 16 K core-jobs are allowed in zero-priority time.
	ALCF Intrepid		20	52	(Being retired)	
		992				Total will fall in 2014 and (presumably) beyond.
	NCSA(NSF), Blue Waters		656		Partly GPU-based	No firm allocation policies announced yet. USQCD PRAC grant is entirely leadership-class, runs through April, 2016. Unknown whether part of a long-term program, and so is not part of USQCD long-term planning.
		656				
General purpose	NERSC (DoE)		55			
	XSEDE (NSF)		109			
		164				

*** The power of GPU-based systems is highly dependent on the calculations being run.**
The rating of the LQCD GPUs was based on the particular mix of projects being run when the rating was done.

The physics collaborations making up USQCD also apply for time at NERSC, NSF XSEDE, ..., independently of USQCD, but these come far short of meeting our physics needs.

Incite resources

for capability computing.

- The DOE allocates time at its leadership class facilities (LCFs), the Cray/GPU system (Titan) at ORNL and the BlueGene/Q (Mira) at ANL, through its INCITE Program.
  - These are essential for that part of our program requiring high capability computing, such as generation of large gauge configuration ensembles.
- USQCD currently has a three year grant running from calendar year 2014 though calendar year 2016.
- The USQCD allocation for 2014:
  - 100 M core-hours on Titan, the largest at the OLCF.
  - 240 M core-hours on Mira, the largest at the ALCF.

Capacity and capability computing

- Leadership class computing is essential for generating large ensembles of gauge configurations. This computing cannot be done any other way.
- We have an even greater need for flops analyzing these configurations.
    - Can often be done very efficiently (cheaply in $/flop) in parallel on much smaller systems.

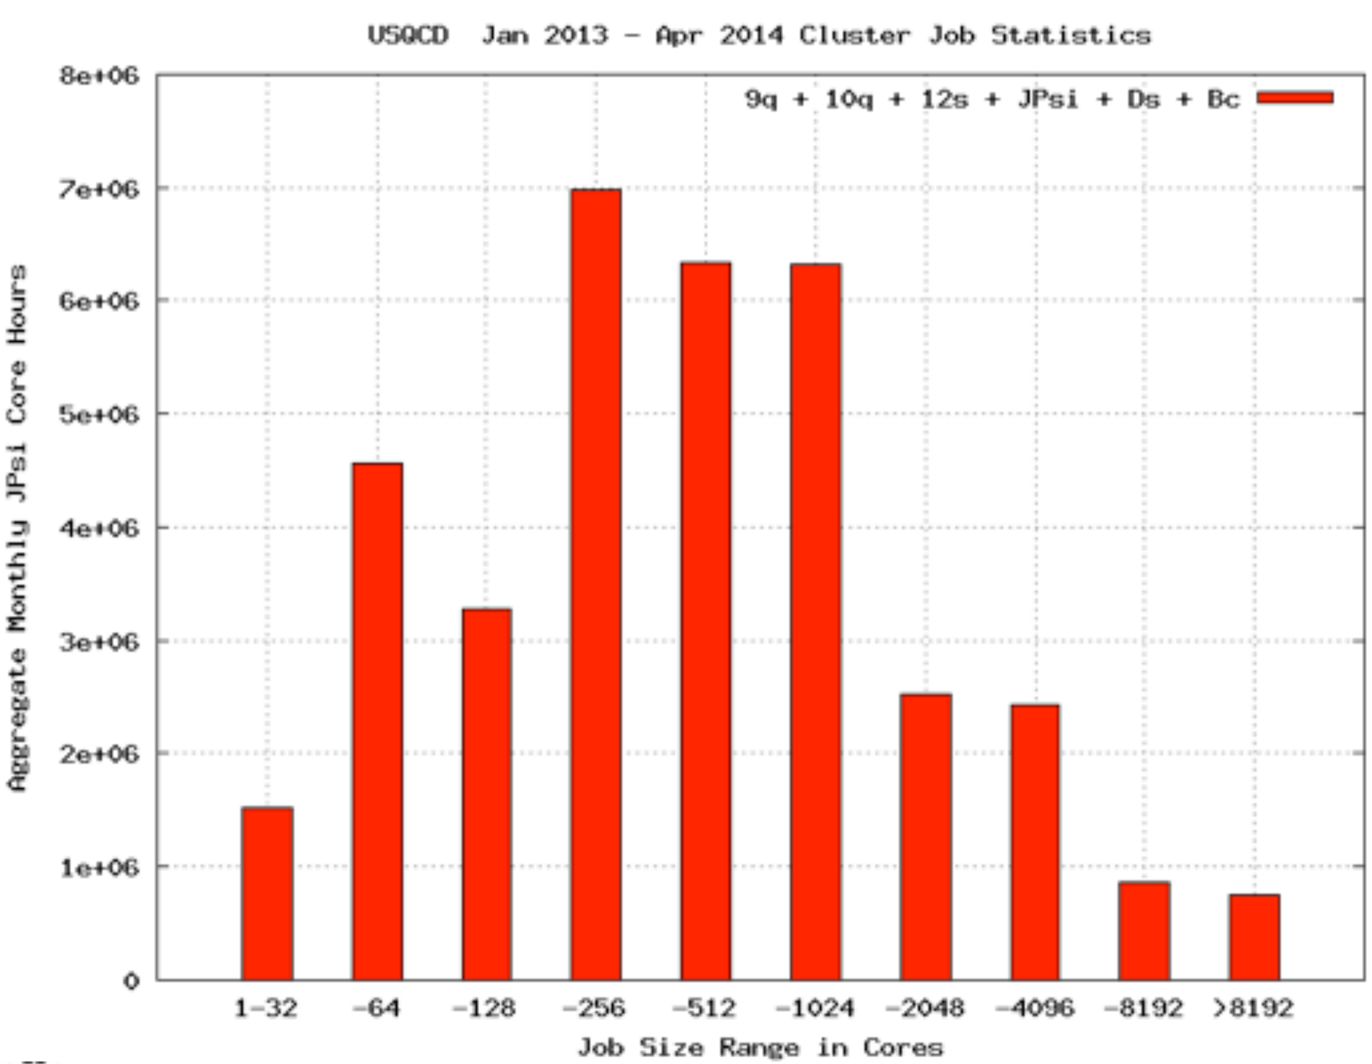

We have an approximately flat distribution of job-size needs from one-node jobs to hundred thousand node jobs on a log scale in job-size.

This year we have a (temporary) deficiency of cycles between the 1,000-node jobs that run efficiently on clusters and the 100,000-core jobs that are mandated at the LCFs. Gap filled by small BG/Q at BNL.

Job size distribution on USQCD 2013/14 conventional clusters.

Capacity and capability computing

- The LCFs are mandated to supply only capability computing needs, jobs that can't be done on any other machines.
    - >128,000 core jobs at ALCF, >3,600 GPU jobs at OLCF.
- Lattice QCD has an even greater need for capacity cycles that LCFs are mandated *not* to supply.
    - The LCFs are well aware that this need exists and are very happy that it is being met because it gives value to the large ensembles we generate there.
- This capacity need can be supplied much more efficiently on dedicated lab clusters than at multipurpose computing facilities.
    - A case study examined by USQCD showed that in one case the USQCD hardware was a factor of three more cost-effective than LCF hardware.

HEP and NP labs are well suited to supply this need.

- LQCD clusters leverage lab capabilities.
  - Cluster expertise (reconstruction farms and real time triggers)
  - Storage (networks, file systems, data movement)
  - Lattice clusters exist in an environment of similar but larger experimental clusters for reconstruction and analysis; cycles can be traded back and forth between projects during critical periods, benefiting all concerned.
- LQCD hardware often provides design ideas and prototyping that is useful to other programs at labs and universities.
  - E.g., at Fermilab, for several years Ds-type machines (quad-socket Opterons) have been the standard used for Run 2, FermiGrid, CMS Tier 1;
  - other programs at the labs are now becoming very interested in GPUs and Intel MIC architecture.

GPUs

- GPUs have supplied a significant fraction of our capacity computing needs in the last few years.
  - A disruptive technology: for the projects that can use them, they are very price performant in $, but require significant investments in software and physics brain power.
  - Some of our projects are mature on GPUs, but a large fraction is not.
  - Price performance varies much more by project than is true for ordinary clusters.
    - Harder to define a standard candle for price/performance.
  - The speedup we find has been enabled by the terrific work of our software committee.

USQCD hardware resources by year

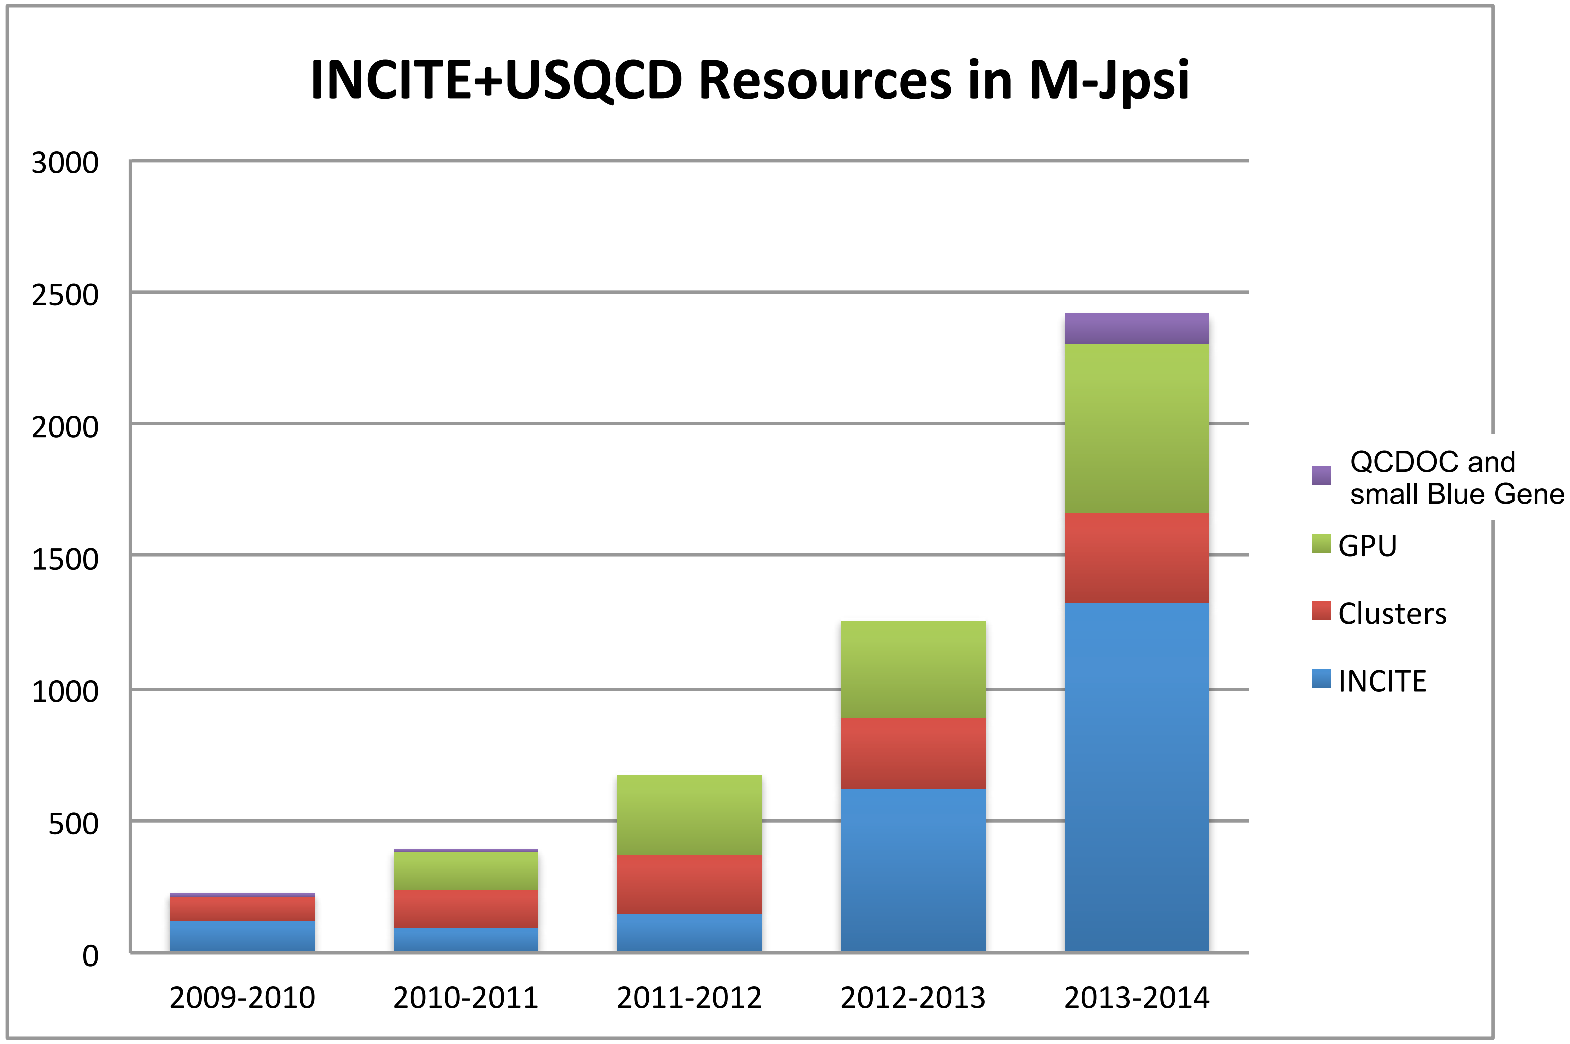

GPUs 20% of $, but of 50% of cycles
on USQCD clusters (red and green bars);
but less general purpose cycles:
much not error correcting,
many single GPU jobs.

- Total USQCD hardware resources have risen exponentially following Moore's law for the last few years.
- This will not continue for the next few.

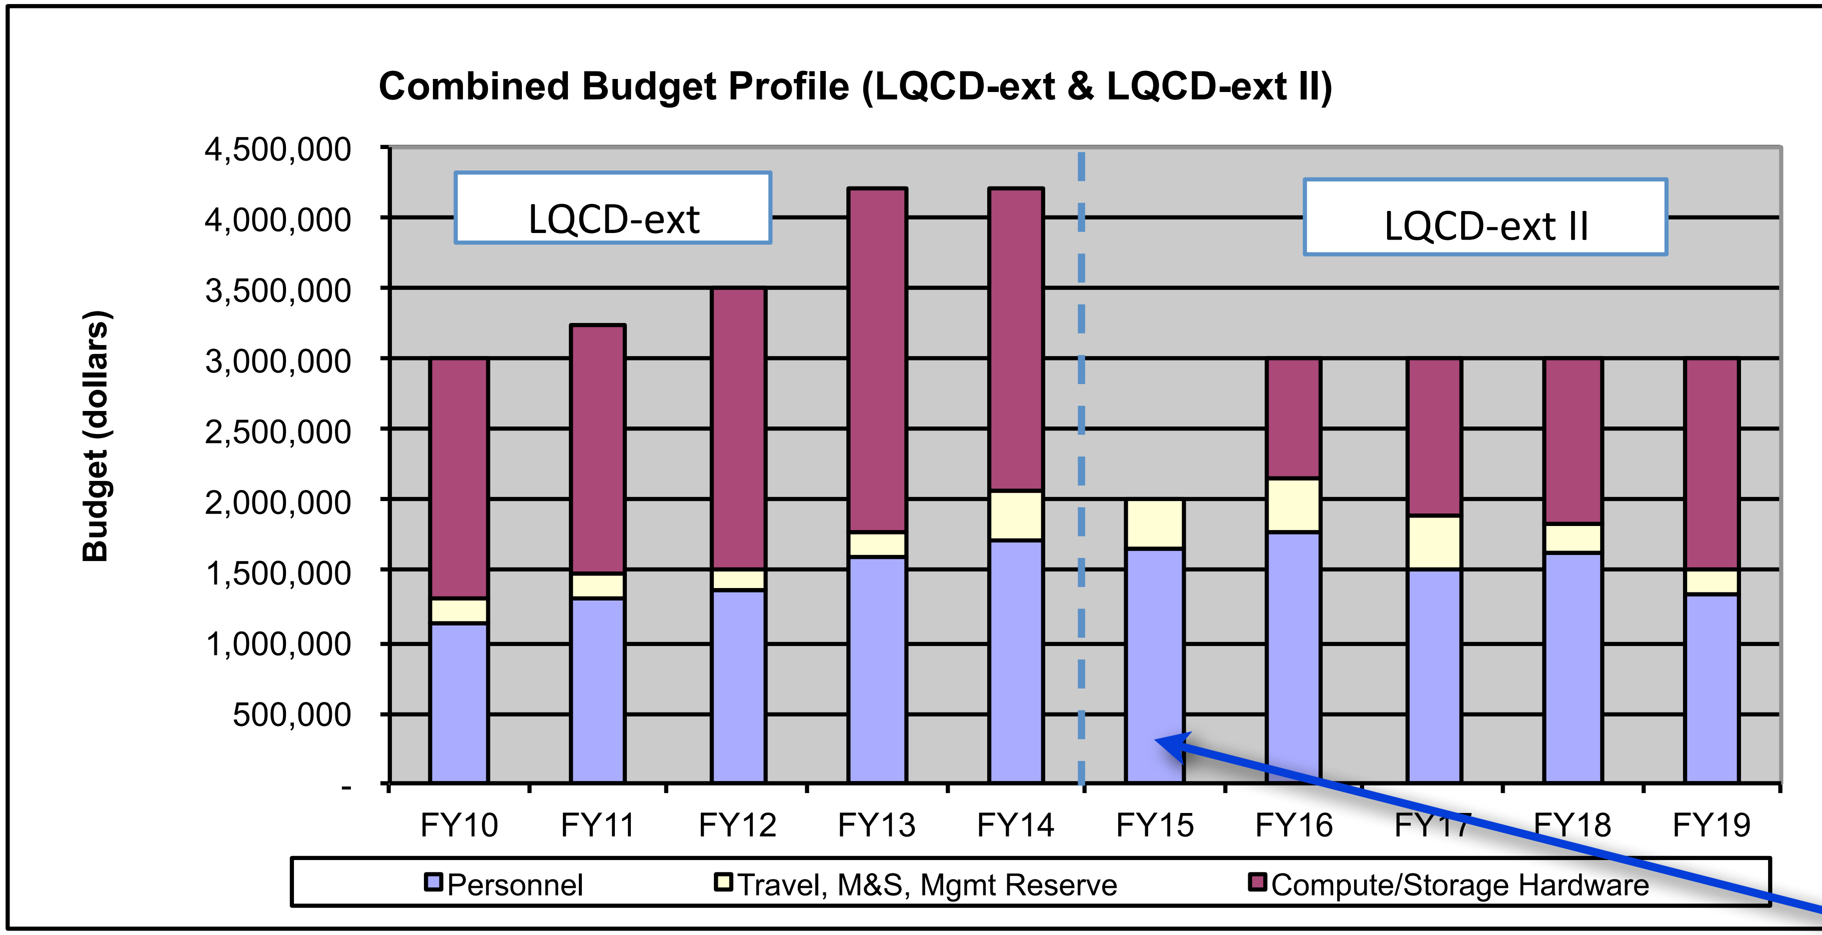

Under LQCD-ext, we have spent ~ 60% of our funds on hardware, about 40% on operations. Current guidance is for FY15 funding to allow only continued operations - no new hardware purchases.

- Current DoE guidance is to expect reduced and back-loaded funding for LQCD-ext II compared with LQCD-ext, particularly in 2015, a terrible year for HEP. (CD 2/3 on June 5, 2014.)
- The Leadership-class at Argonne and Oak Ridge upgraded their resources by a factor of ~ten in 2013; the next major upgrade is planned for ~2018.
- USQCD science program will have to adjust for this fact, stretch out science goals, make sure our most important deliverables remain on track.

SciDAC lattice QCD software R&D

The third critical component of our computational infrastructure.

Software Committee: Richard Brower (chair).

Regular Thursday phone conferences for people working on USQCD software.

USQCD has SciDAC-3 grants from HEP+ASCR and NP+ASCR for about $1 M each for creating lattice QCD software infrastructure: community libraries, community codes, optimization and porting to new architectures, implementation of up-to-the-minute algorithm advances...

- The QCD API and community libraries
  - Lower entrance barriers to lattice QCD.
  - Enable postdocs to run major projects without being part of major collaborations.
- Porting and optimizations for new platforms
  - Critical to efficient use of new hardware.

Organization, goals, allocations

- In 2003 when USQCD hardware funding began, Peter Rosen (head of HEP and NP) made it clear that DoE expected the hardware to be operated as a national facility.
  - Open to all in US to submit proposals.
  - USQCD is like Fermilab fixed-target facilities, not like CMS or LBNE.
  - Overall physics goals are set by USQCD in our white papers and proposals for hardware and software, but specific projects are developed by component collaborations like MILC, RBC, NPLQCD, HOTQCD, ..., or by individual postdocs and allocated by SPC. (Role of EC in this process is analogous to that of lab director.)
- We think this model has worked very well.
- A different model: USQCD could function as a physics collaboration like CMS or LBNE.
  - Not open to all in US; individuals would apply to join and specify their contribution in advance.

- Advantages of the facility model.
  - Young people can be PIs of their own physics programs as soon as they are able to formulate a project and a proposal that is convincing to the Scientific Program Committee.
    - They can be recognized for their own scientific programs much more easily than as part of a hundred-member collaboration.
    - The five people who got junior faculty or staff jobs in the last couple years all served as PIs of their own proposals; two of them with no senior collaborators.
  - When groups disagree on methods, they can compete.
- Possible advantages of the collaboration model?
  - Could more straightforwardly enforce that the most important goals are implemented in the allocation process? We think that the current allocation process does this well.
- We think that the advantages outweigh the disadvantages of the facility model.

Executive Committee

Paul Mackenzie (chair), Rich Brower, Norman Christ,
Frithjof Karsch, Julius Kuti, John Negele, David Richards,
Martin Savage, Bob Sugar

Present today.

- Provides overall leadership for the collaboration and point of contact for the DoE.
- Organizes the writing of the proposals for hardware and software and of the white papers and chooses the members of the other committees.
- Rotates new members at ~ one/year.
  - Close to full rotation over ~ 10 years is planned. A about half has rotated already.
  - The EC rotates in a way that preserves rough balance. Current practice: approximate balance between HEP and NP, one member from each of the half dozen most important physics collaborations, one member from each of the three partner labs, a few members from outside these groups.

Executive Committee

- Last years review committee suggested we consider:
  - implementing mechanism for more regular turnover of EC members,
  - electing one or more members at All Hands meetings.
- This year, we have instituted specific terms (with possibility of reappointment) for Executive Committee members, with the terms of most senior members expiring first.
  - Seniority defined by years service on committee, years from PhD in case of ties.
  - Bob Sugar (MILC) and Norman Christ (RBC) were considered this year. The EC consulted with members of MILC and RBC, and the collaborations consulted among themselves.
  - Result: the EC asked Christ to continue on the committee; asked Carleton DeTar of Utah to replace Sugar. DeTar requested that the rotation be deferred for two years when his term as Utah department chair ends, which the EC accepted.
  - For EC members not representing large collaborations, we are continuing to consider other ways of making rotations, such as elections of at-large members at All Hands meetings.

Scientific Program Committee

Robert Edwards (chair), JLab, Simon Catterall, Syracuse, Will Detmold, MIT, Taku Izubuchi, BNL, Peter Petretzky, BNL, Doug Toussaint, Arizona, Ruth Van de Water, Fermilab

Present today.

Each year, the many smaller physics collaborations within USQCD submit proposals to the Scientific Program Committee for allocations of time on USQCD's LQCD and Incite resources.

The SPC creates a program to accomplish the goals set forth in the USQCD Collaboration's proposals.

- It may also advise us on needed evolution of the goals.
- It examines submitted proposals in light of the desired program.
- In principle, it could state in the Call For Proposals that proposals in a certain area would be welcome; has not seen the need to do that yet.
- The SPC is programmatic without being top-down.

The Executive Committee seeks the advice of the SPC on physics priorities when writing new proposals for DoE computing resources.

Chair rotates every two years. Members rotate every four years, at a rate of about two/year.

30 people have served so far as members of either the EC or the SPC.

Goals: proposals and white papers

2013 USQCD white papers at http://www.usqcd.org/collaboration.html

- The physics goals of USQCD are set out in our proposals and white papers organized by the Executive Committee in consultation with the SPC.
  - 2013 white papers had 23 authors.
  - Continually evolving, in consultation with the SPC.
  - Discussed by the Collaboration at All Hands meetings.
- 2013 white papers are the most recent statement of our view of our most important goals and opportunities, and our view of our highest impact results.
  - Outlined possible physics program assuming approximate continuation of current support, $23 M over five years from HEP and NP. (Current guidance is for lower level of support in renewal project, $14 M over five years. Progress will be slower than hoped.)
- In setting and updating our goals we have always relied on informal input from numerous experimenters and phenomenologists.

Science Advisory Board

- This year, we are formalizing this process by naming a Science Advisory Board.
  - Brendan Casey (Fermilab, g-2), Marina Artuso (Syracuse, LHC-b), Jesse Thaler (MIT), David Kaplan (U. Washington), Curtis Meyer (Carnegie Mellon, GlueX), Nu Xu (LBL, Star), Volker Koch (LBL).
  - Among the most useful advisors on white papers and proposals.
- At the beginning of each year's allocation process, they are asked to
  - Comment and suggest revisions of our general goals,
  - Read and comment on the year's physics proposals and allocations,
    - We are exploring how closely our advisors would like to be involved with the allocations process.

Input from the 2014 SAB

- Useful comments on USQCD program,
  - as food for thought. E.g., the lattice HEP program would be stronger if it reflected more faithfully the HEP experimental program of the next few years,
  - or because they reflected a need to improve our message.
- USQCD allocations process
  - “I find the proposals I read mostly pretty well written, with a science justification in the intro, the abstracts are all remarkably of the same format: brief science justification, goals, requested allocation, which is pretty accessible (without being asked to judge whether the project is realistic)...I do not actually imagine that the SAB is going to have much useful feedback for you, but sharing this information might impress the people on the board about what a diverse and active community this is.”

Detailed comments of SAB members are on the review web site.

Allocations

Robert Edward's talk

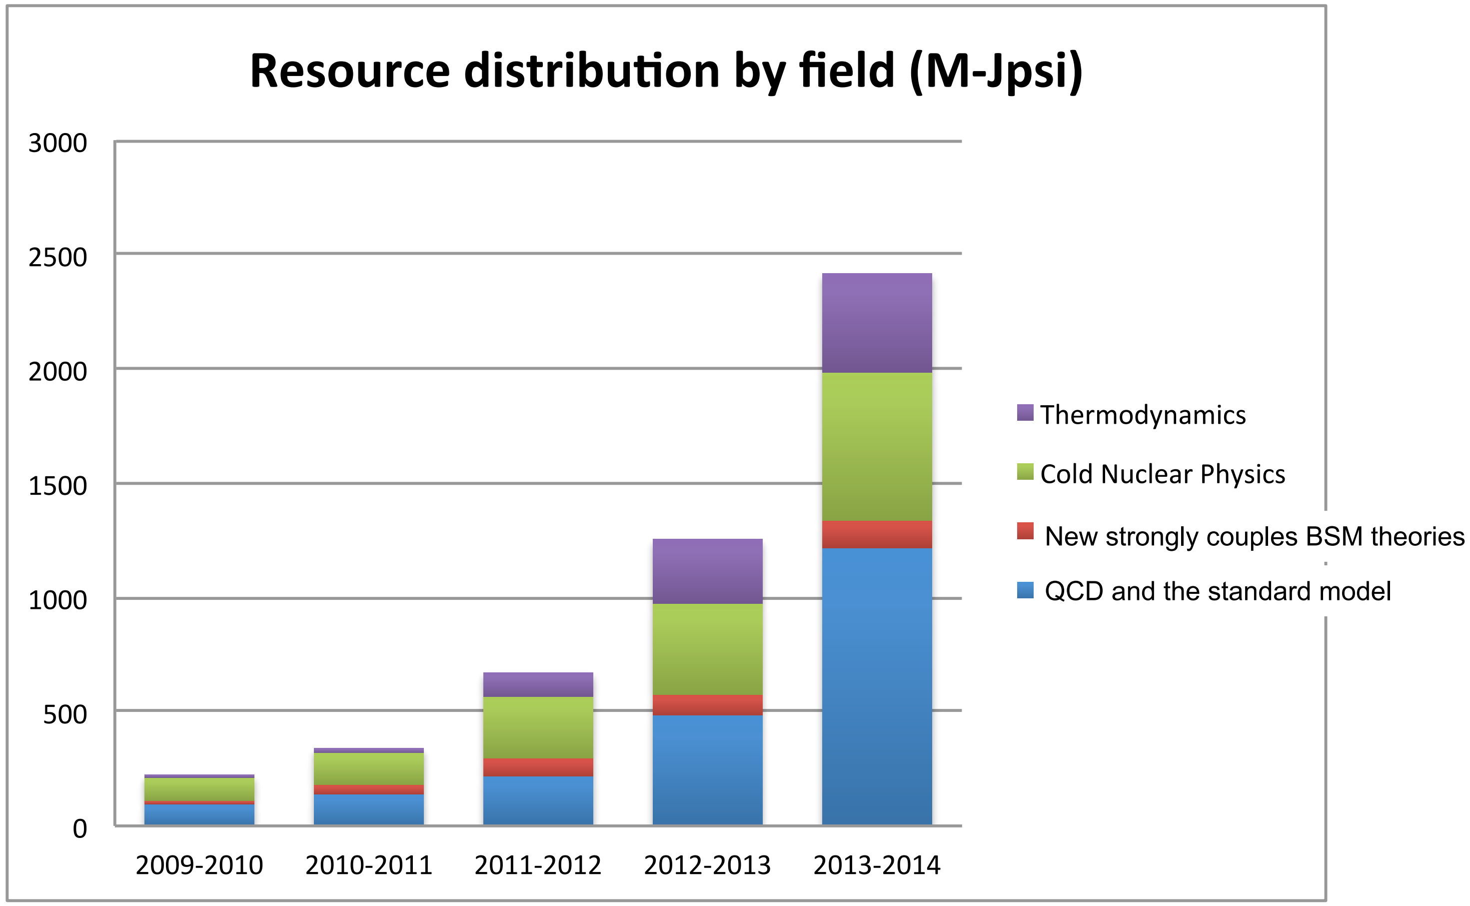

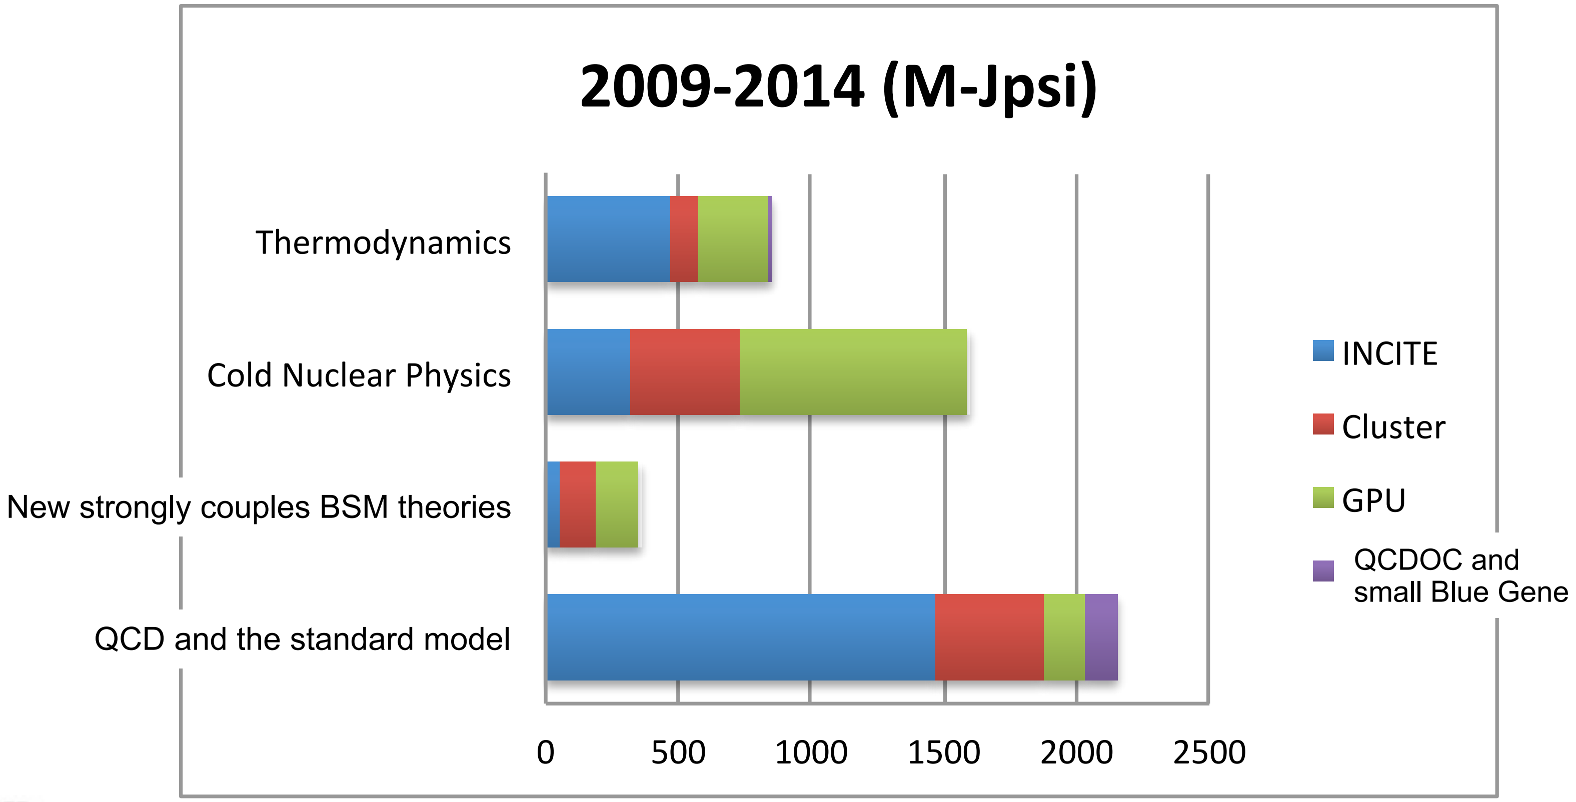

Allocations

Type A cluster allocations in 2013 by project.

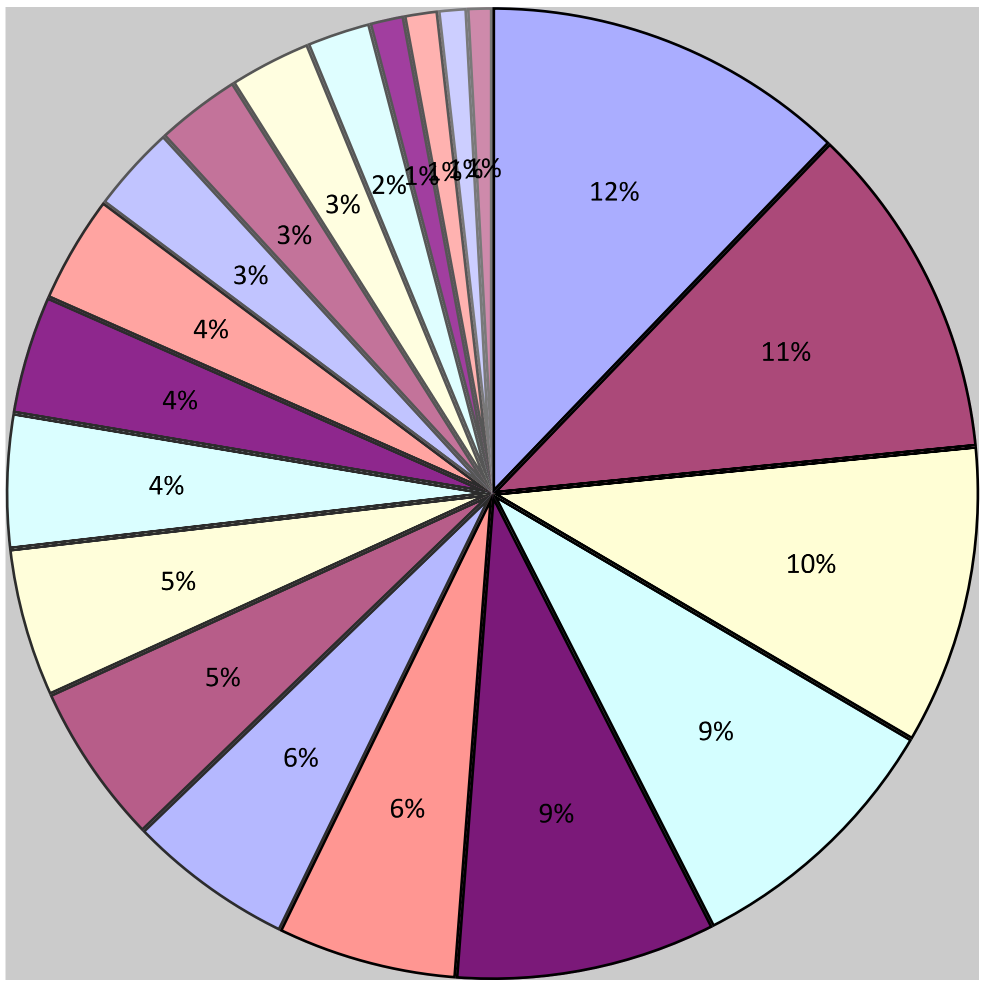

Projects are judged by:
- relevance to the central goals of USQCD;
- size and competence of project team;
- validity and efficiency of methods proposed.

Less high priority projects are typically not zeroed out, but are given less resources. About half of these allocations went to the five highest priority projects (an HEP project producing a dozen high-priority weak matrix elements, an NP project calculating resonance spectra for Gluex, ...).

(Different from experimental programs, where experiments must be voted either up or down.)

2014 USQCD all-hands meeting

- Took place April 18-19, 2014 at JLab. ~56 members attended. (http://www.usqcd.org/meetings/allHands2014/.)
- Reports from the Executive Committee, the LQCD-ext Project Manager, the SPC, and the hardware site managers.
- In each science domain, reports from
  - white paper authors,
  - representative physics projects,
  - members of the SPC on the relation between the allocated projects and the long-term goals.

- Collaboration discussion on
  - USQCD scientific program in light of expected tightening of funding in the follow-on period 2015-19.

Lattice meets experiment meetings

To increase the interaction between lattice gauge theory and experiment and phenomenology, members of USQCD have organized a series of workshops with experimenters and phenomenologists.

- SLAC, Sept. 16, 2006, Standard Model physics. With BaBar.
- Fermilab, December 10-11, 2007, “Lattice Meets Experiment” in flavor physics.
- Livermore, May 2-3, 2008, “Lattice Gauge Theory for LHC Physics”.
- JLab, Nov. 21-22, 2008, “Revealing the Structure of Hadrons”, Nuclear.
- BNL, June 8-9, 2009, “Critical Point and Onset of Deconfinement”, QCD thermodynamics.
- BU, Nov. 6-7, 2009, “Lattice Gauge Theory for LHC Physics”. BSM.
- Fermilab, April 26-27, 2010, “Lattice Meets Experiment” in flavor physics.
- BU, 8-10 September 2010, “Sixth Workshop on QCD Numerical Analysis, Boston.
- JLab, Feb. 23-25, 2011, “Excited Hadronic States and the Deconfinement Transition”.
- BNL, Oct. 3-5, 2011, "Fluctuations, Correlations and RHIC low energy runs".
- Fermilab, Oct. 14-15, 2011, “Lattice Meets Experiment: Beyond the Standard Model”.
- Boulder, Oct 28, 2012, “Lattice Meets Experiment 2012: Beyond the Standard Model”.
- George Washington University, Aug. 21-23, 2012, “Extreme QCD”.
- BNL, December 5-6, 2013, “Lattice Meets Experiment 2013: Beyond the Standard Model”.
- Fermilab, March 7-8, 2014, “Lattice Meets Experiment, 2014”.

International collaboration

- Lattice QCD is an international field with very strong programs in Germany, Italy, Japan and the United Kingdom, and elsewhere.
- Non-US lattice theorists are welcome to contribute to USQCD projects in collaboration with US theorists.
- Groups within USQCD have formed a number of international collaborations:
  - The USQCD effort using DWF quarks is an international effort between the United States based RBC, the Edinburgh, and Southampton members of the UKQCD Collaboration, and RIKEN.
  - The Fermilab Lattice, HPQCD and MILC Collaborations have worked together in various combinations to study heavy quark physics using improved staggered quarks. HPQCD includes physicists in both USQCD and UKQCD.
  - Members of the BNL Nuclear Physics lattice gauge theory group have a long term collaboration with physicists at the University of Bielefeld, Germany.
  - Members of USQCD working on the hadron spectrum using Clover quarks on anisotropic lattices have close ties with colleagues in Trinity College, Dublin, the Tata Institute, Mumbai, Cambridge U.
  - ...

Membership survey and demographic information

- We are starting to collect membership and demographic information in a more organized way.
  - New membership list. Currently, ~ 163 members.
  - Demographic survey.

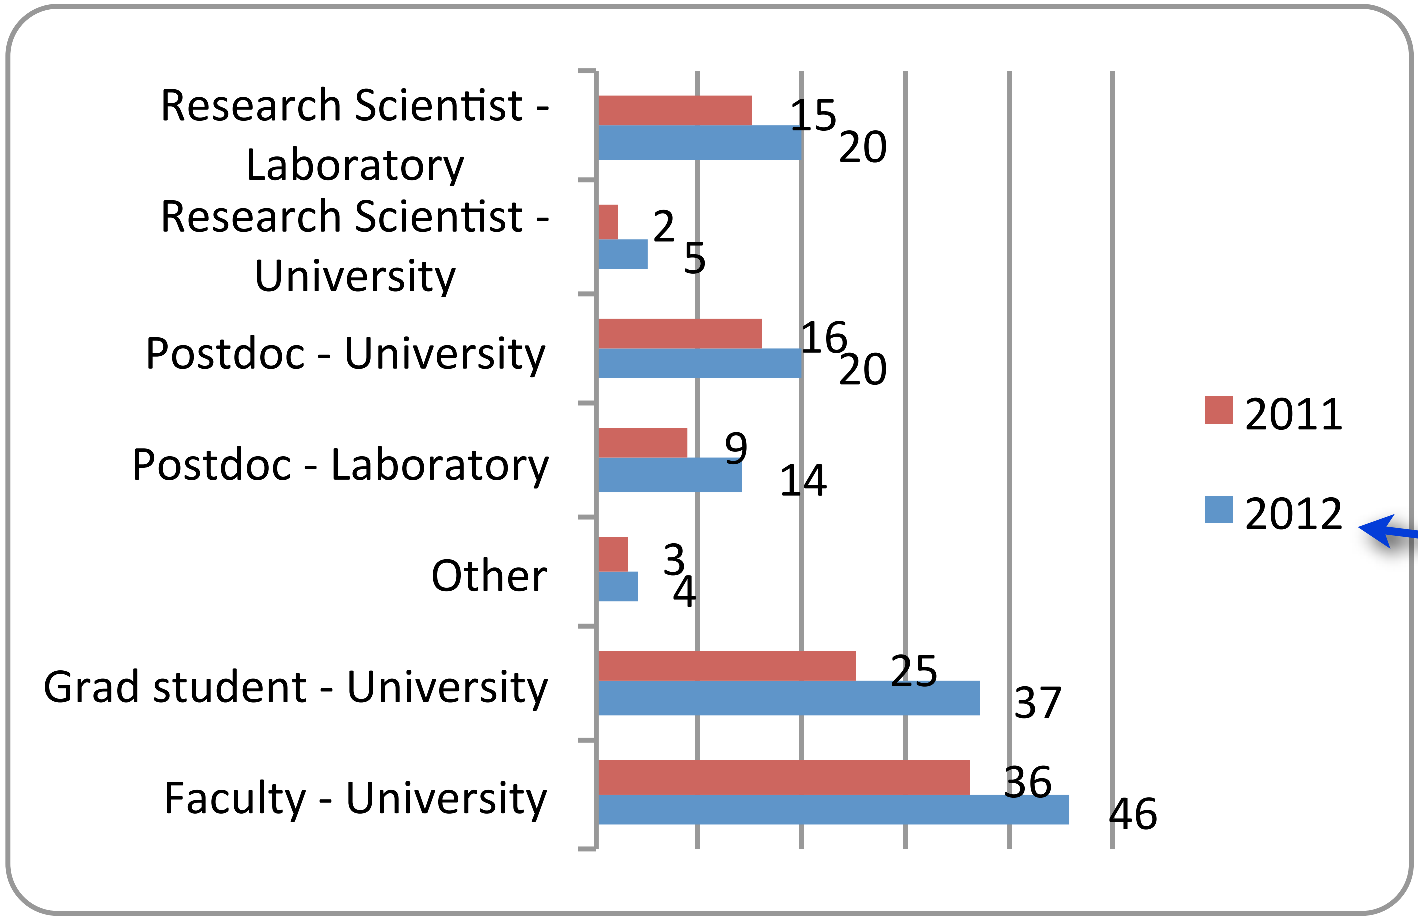

Better response to the survey in 2012, not a increase in the field.

Intend to make it annual. Now overdue, planned for this year.

We've grown from about 90 people in 2000 to about 163 today.

Junior faculty and staff job creation

	Year	Research institution, HEP	Research institution, NP	Computational scientist	Teaching college	Industry	Foreign
Mei-Feng Lin	2014			BNL			
Stefan Meinel ***	2014	Arizona					
Hiroshi Ohno	2014						Tsukuba
Heng-Tong Ding	2013						CCNU
Andre Walker-Loud	2013		William & Mary				
Jack Laiho	2013	Syracuse					
Will Detmold **	2013		MIT				
Ethan Neil ***	2013	Colorado					
Christopher Thomas	2013						Cambridge
Ruth Van de Water	2012	Fermilab					
Brian Tiburzi ***	2011		CUNY				
Andrei Alexandru *	2011		GWU				
Elvira Gamiz	2011						Granada
Mike Clark	2011					NVIDIA	
Ron Babich	2011					NVIDIA	
Christopher Aubin	2010				Fordham		
Swagato Mukherjee	2010		BNL				
Changhoan Kim	2010					IBM	
Enno Scholz	2009						Regensburg
Taku Izubuchi	2008	BNL					
James Osborn	2008			Argonne			
Chris Dawson	2007	Virginia					
Nilmani Mathur	2007						Tata Institute
Joel Giedt	2007	RPI					
Matthew Wingate	2006						Cambridge
Jozef Dudek **	2006		Old Dominion				
Jimmy Juge	2006				U. of the Pacific		
Peter Petreczky	2006		BNL				
Balint Joo	2006			JLab			
Kieran Holland	2006				U. of the Pacific		
Kostas Orginos **	2005		Wm & Mary				
George Fleming	2005			Yale			
Tom Blum ***	2003	Connecticut					
Silas Beane *	2003		UNH				
Total		8	9	4	3	3	7

* NSF Early Career Award
** DoE OJI/Early Career
*** RIKEN/BNL joint positions

USQCD PhDs produced

since 2004.

	2013	2014
Kuti	3	3
Savage	4	5
Sharpe	4	4
Beane	2	2
Kaplan	1	2
MILC	10	11
Columbia	13	15
Karsch	3	4
BU	3	3
MIT	5	6
Maryland	3	3
CMU	3	5
Willam & Mary	3	4
Kentucky	3	3
	60	70

The computational challenge of lattice QCD

Lattice	Quark			Gauge	ensembles		Analysis propagators, correlators		
spacing *a (fm)*	mass *ml/ms*	Volume (sites)	Configu rations	Core-hours (M)	TB/ ensemble	Files/ ensemble	Core-hours (M)	TB/ ensemble	Files/ ensemble
0.15	1/5	16^3*48	1000	1	0.1	1,000	1	4	155,000
0.15	1/10	24^3*48	1000	2	0.2	1,000	2	12	“
0.12	1/5	24^3*64	1000	3	0.3	1,000	3	16	155,000
	1/10	32^3*64	1000	8	0.6	1,000	8	39	“
	1/27	48^3*64	1000	26	2.0	1,000	26	130	“
0.09	1/5	32^3*96	1000	10	0.9	1,000	10	58	155,000
	1/10	48^3*96	1000	35	3.1	1,000	35	196	“
	1/27	64^3*96	1000	46	7.2	1,000	46	464	“
0.06	1/5	48^3*144	1000	38	4.6	1,000	38	294	155,000
	1/10	64^3*144	1000	128	10.9	1,000	128	696	“
	1/27	96^3*144	1000	218	36.7	1,000	218	2,348	“
0.045	1/5	64^3*192	1000	135	14.5	1,000	135	928	155,000
	1/10	88^3*192	1000	352	37.7	1,000	352	2,412	“
	1/27	128^3*192	1000	1083	116.0	1,000	1,083	7,422	“
0.03	1/5	96^3*288	1000	685	73.4	1,000	685	4,697	155,000
				2,770					

Example gauge ensemble library.

CPU times normalized in JPsi core-hours.

Planned MILC HISQ ensembles of guage configurations.
$m_l = 1/27 m_s = m_{phys}$

Operationally, lattice QCD computations consist of

1) Sampling a representative set of gauge configurations with Monte Carlo methods,

    E.g., the Metropolis method, the hybrid Monte Carlo algorithm, ...
    Consists of one long Markov chain. A capability task.

2) Calculating the propagation of quarks through the gauge configurations,

    Solve the Dirac equation on each configuration with relaxation methods, e.g., biconjugate gradient algorithm, etc. A capacity task.

3) Constructing hadron correlation functions from the quark propagators (smaller task).

Two main components of a typical lattice calculation

multi-TB
file sizes

Generate O(1,000) gauge configurations on a leadership facility or supercomputer center. Hundreds of millions of core-hours.

Transfer to labs for analysis on clusters. Larger CPU requirements.

Gauge configuration generation:
a single highly optimized program,
very long single tasks,
"moderate" I/O and data storage.

Hadron analysis.
Large, heterogeneous analysis code base,
10,000s of small, highly parallel tasks,
heavy I/O and data storage.

Two comparably sized jobs with quite different hardware requirements.

US lattice gauge theory work flow

Zero-temperature QCD:

Currently three main streams of QCD gauge configurations are being generated by USQCD for different physics goals:

These high-value ensembles are data-rich resources that are shared among all of USQCD.

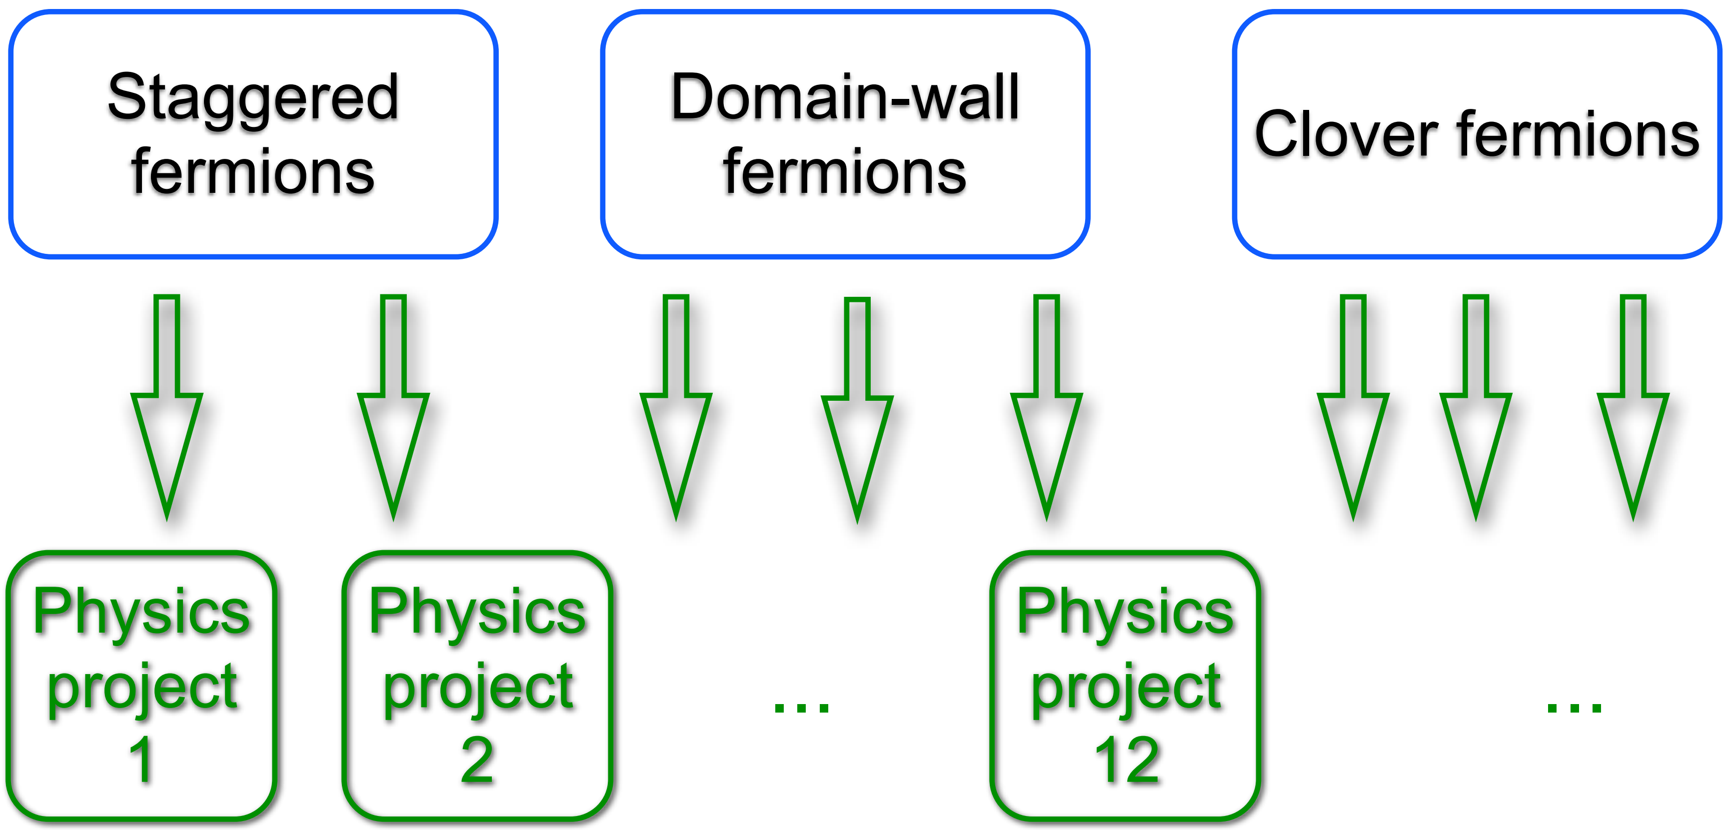

Shared among a couple of dozen groups, in both HEP and NP.

Physics projects are done on these configurations by smaller groups of 5-15 members within USQCD.

Around 90 of the 163 members of USQCD have submitted jobs to USQCD hardware.

QCD thermodynamics and BSM projects generate their own configurations tailored to specific goals.

2013 USQCD physics projects

Projects for configuration generation

Projects initiated by post-docs or students

PI	Project
Aubin	Hadronic contributions to the muon g-2 using staggered fermions
Bazavov	The Equation of State and fluctuations of conserved charges in 2+1 flavor QCD using Highly Improved Staggered Quarks
Catterall	Lattice study of N = 4 Super Yang-Mills
Christ	Generating ensembles with 2+1 flavors of domain wall fermions
DeTar	Quarkonium Physics with Open Charm and Radiative Transitions
Detmold	Light Nuclei, Hypernuclei, and their Electromagnetic Properties
Ding	Universal properties of the chiral phase transition in 2+1 flavor QCD using Highly Improved Staggered Quark action
Edwards	Dynamical Anisotropic-Clover Lattice Production for Hadronic and Nuclear Physics
Engelhardt	Nucleon transverse momentum dependent parton distribution functions on a large isotropic clover fermion ensemble
Fleming	Lattice Gauge Theory for Physics beyond the Standard Model on Leadership Class Machines
Hasenfratz	Eight Flavor SU(3) gauge theory with nHYP smeared fermions
Izubuchi	Hadronic vacuum polarization and hadronic light-by-light contributions to the muon anomalous magnetic moment using statistical error reduction techniques
Kelly	Lattice Determination of the Delta I= 1/2, K -> pi pi Amplitude
Kuti	Can the Higgs impostor hide in the BSM sextet model
Lin	Probing TeV Physics through Neutron-Decay Matrix Elements
Liu	Nucleon Structure with Overlap Fermion
Mackenzie	CKM Physics from B, D, and K Mesons with HISQ Fermions
Mawhinney	Pion and Kaon Physics from 2+1 Flavor DWF Lattices with m_pi = 140 MeV and V=(5.5 fm)^3, II
Meinel	Disconnected Contributions to Nucleon Ground State Structure
Mukherjee	Continuum limit of higher-order charge-fluctuations at the physical point
Syritsyn	Nucleon Structure from Domain Wall Fermions at physical pion mass
Neil	Two color gauge theories in the LHC era
Orginos	Isotropic Clover Fermions
Richards	Excited Meson and Baryon States using Anisotropic Clover Lattices
Shigemitsu	High-Precision Heavy-Quark Physics
Sugar	QCD with Four Flavors of Highly Improved Staggered Quarks
Witzel	B-meson physics with domain-wall light quarks at their physical mass and relativistic heavy quarks

Main scientific thrusts

- Testing the Standard Model and the search for new physics in high-precision experiments. Talk by Ruth Van de Water.
- Understanding the properties of new strongly interacting gauge theories. Talk by Rich Brower.
- The structure and interactions of nucleons. Talk by Martin Savage.
- The behavior of QCD in extreme conditions. Talk by Frithjof Karsch.

Our timeline for delivering on our scientific program will be significantly stretched out by the reduction and backloading of funding expected in our current guidance for LQCD-ext II.

Backup slides

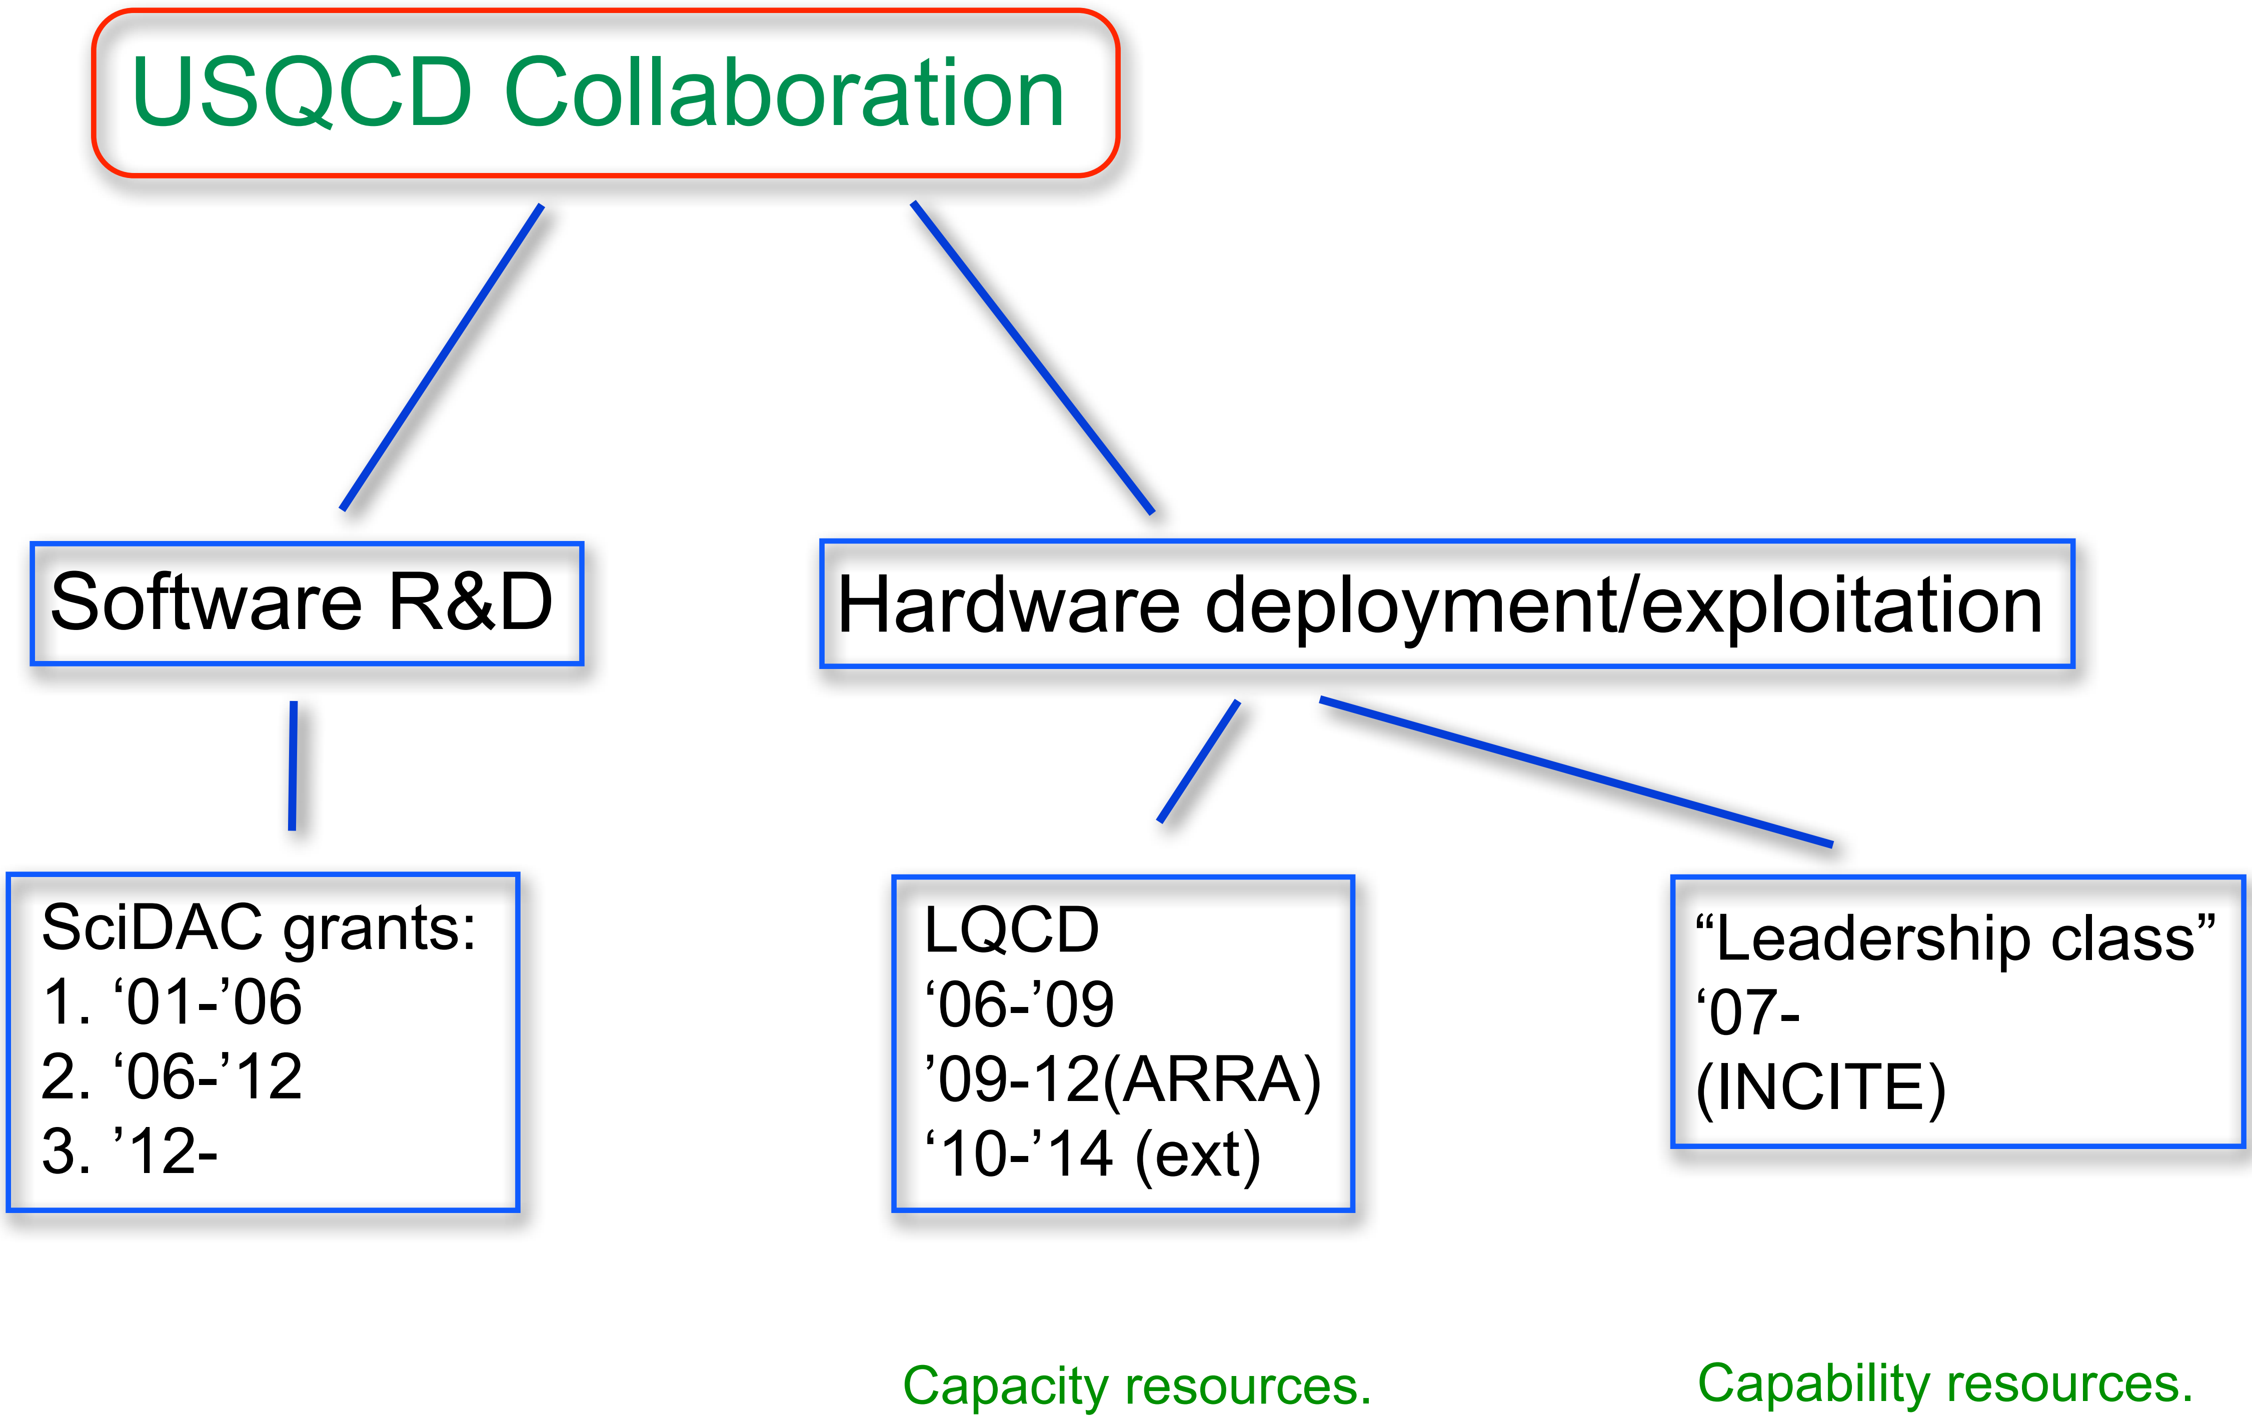

USQCD has grants for
- R&D for software development through the SciDAC program.
- Hardware deployment and use from several sources, including the current LQCD-ARRA and LQCD-ext projects.

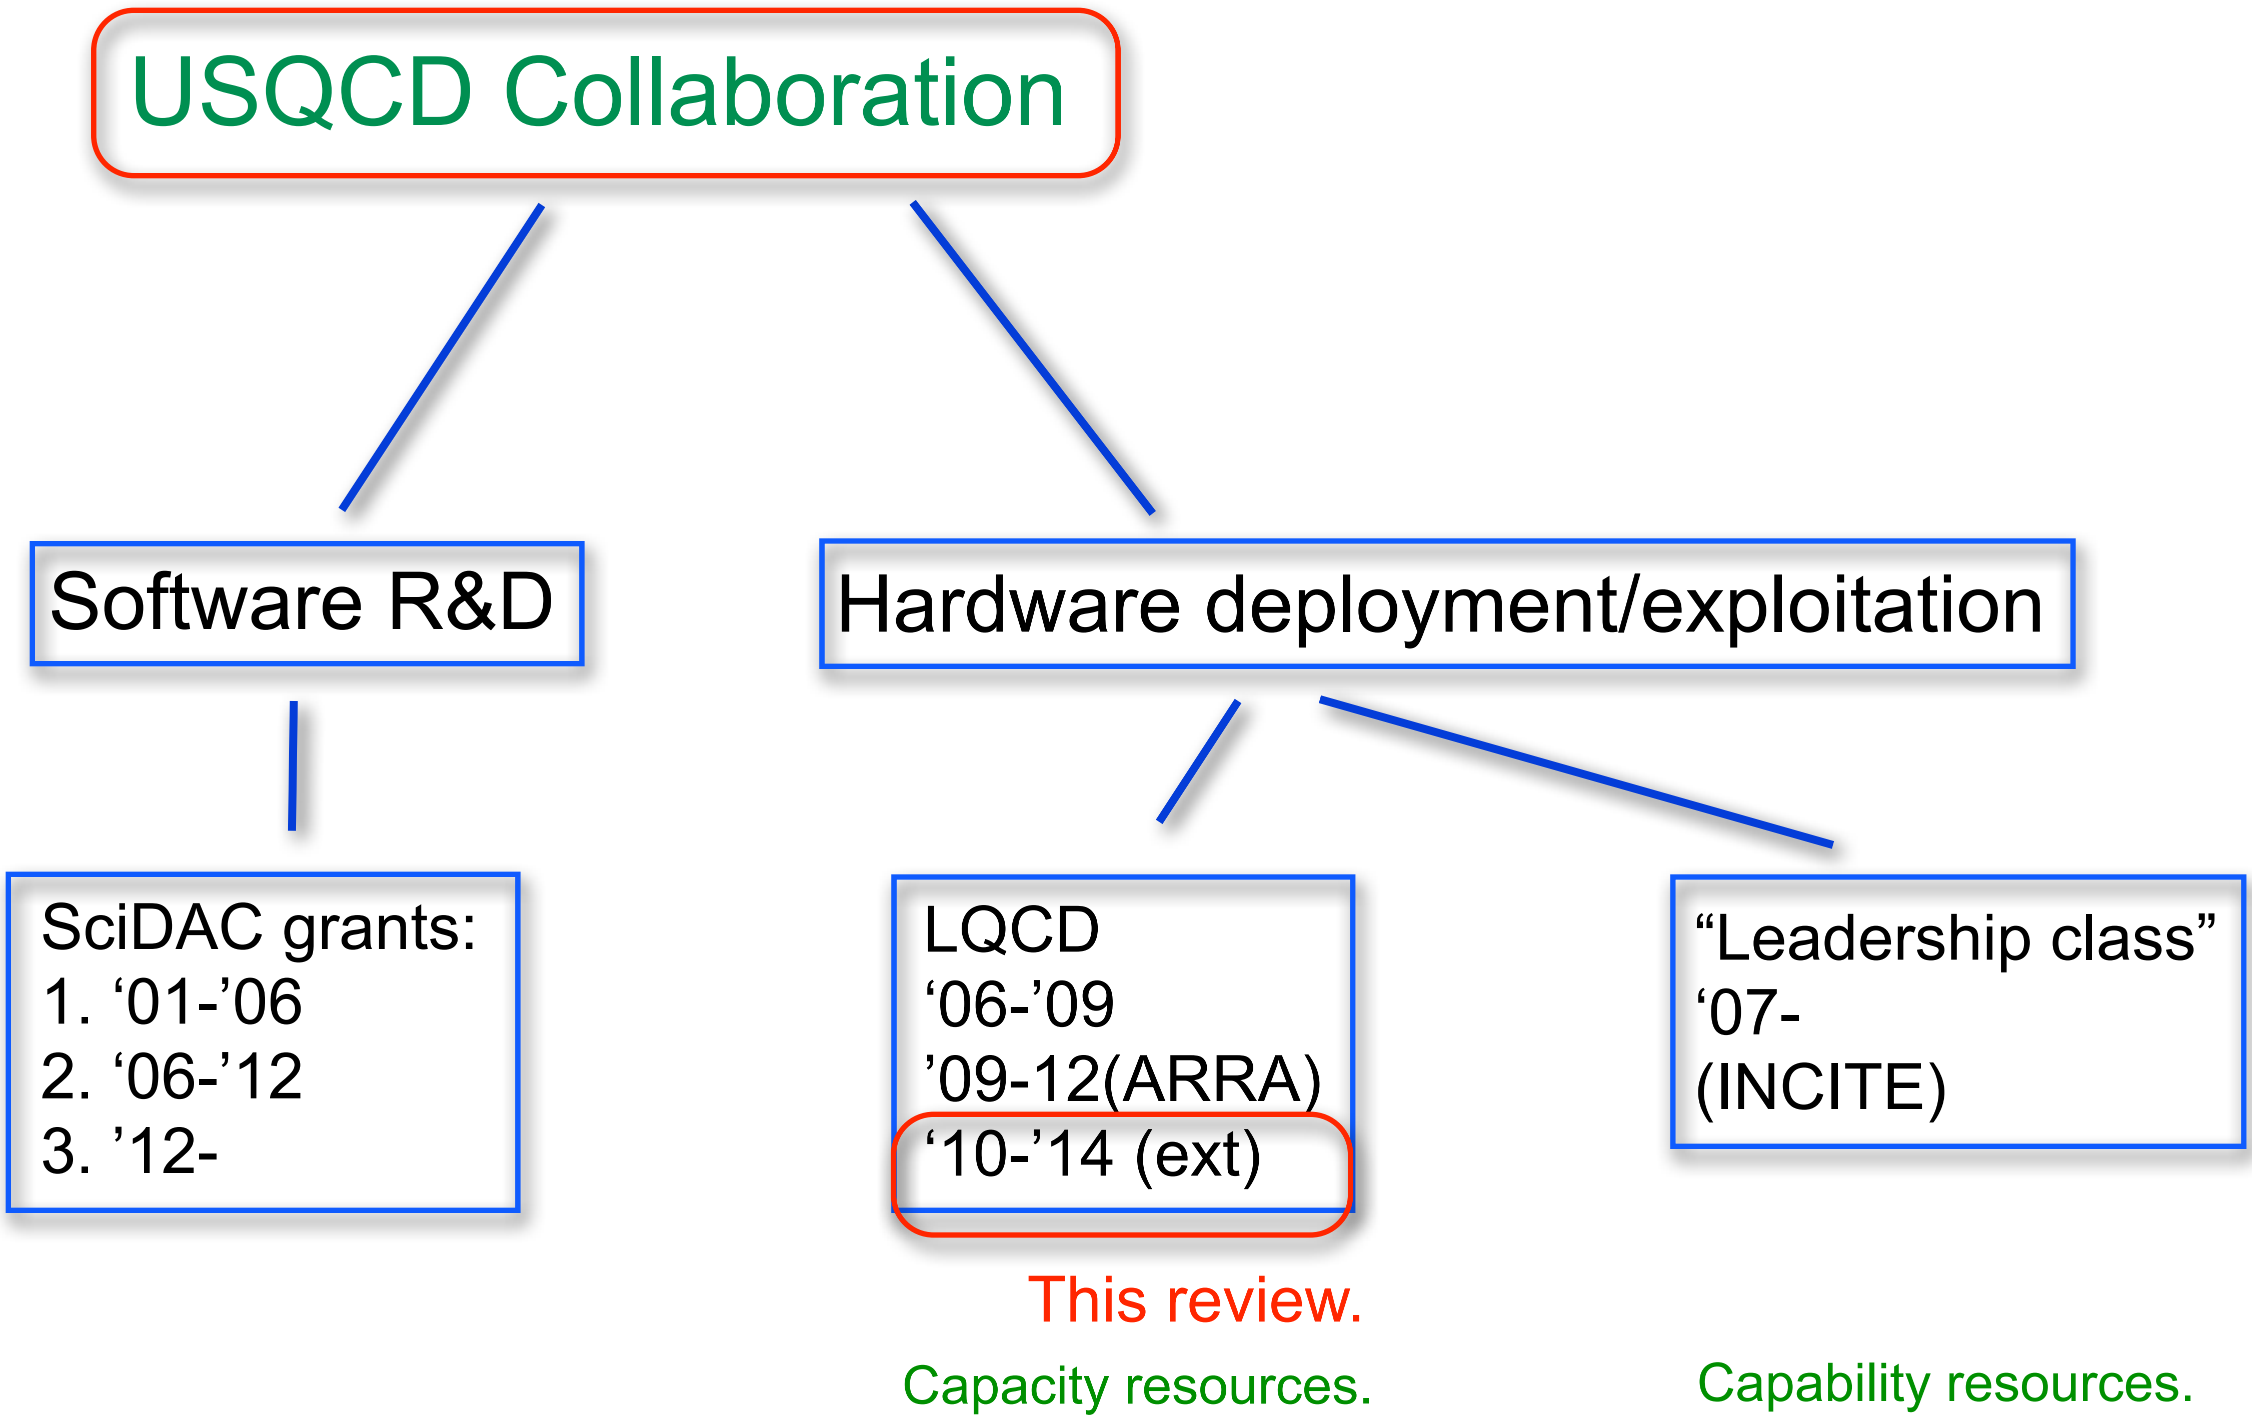

USQCD has grants for
- R&D for software development through the SciDAC program.
- Hardware deployment and use from several sources, including the current LQCD-ARRA and LQCD-ext projects.

Main scientific thrusts

- Testing the Standard Model and the search for new physics in high-precision experiments. Talk by Ruth Van de Water.
  - Experiments impacted: Belle, T2K, HyperK (KEK), BES III (Beijing), g-2, Mu2e, Nova, LBNE (FNAL), NA62 (CERN), KOTO (J-Parc), Belle II (KEK), and LHC-b, ATLAS, CMS (CERN), ILC, super CDMS, LUX.
- Understanding the properties of new strongly interacting gauge theories. Talk by Rich Brower.
  - Facilities impacted: LHC (CERN), ILC.
- The structure and interactions of nucleons. Talk by Martin Savage.
  - Experiments and facilities impacted: GlueX (JLab), RHIC (BNL), FRIB (MSU), g-2 (FNAL), FAIR (GSI), JPARC (Japan), SNS neutron EDM (ORNL), COMPASS and HERMES (DESY).
- The behavior of QCD in extreme conditions. Talk by Frithjof Karsch.
  - Facilities impacted: RHIC (BNL), FAIR, SPS and LHC (CERN).

2013 USQCD physics projects

Distribution of projects among sub-fields

HEP: QCD and SM

HEP: Non-QCD, BSM

Cold NP, hadron structure

Heavy-ion physics

Not a clear distinction between sub-fields:
Some QCD calculations are also looking for BSM effects in hadrons.
Some topics (hadron structure, symmetries of nature, ...) are of interest in both NP and HEP.
Some gauge ensembles are useful in both NP and HEP.

PI	Project
Aubin	Hadronic contributions to the muon g-2 using staggered fermions
Bazavov	The Equation of State and fluctuations of conserved charges in 2+1 flavor QCD using Highly Improved Staggered Quarks
Catterall	Lattice study of N = 4 Super Yang-Mills
Christ	Generating ensembles with 2+1 flavors of domain wall fermions
DeTar	Quarkonium Physics with Open Charm and Radiative Transitions
Detmold	Light Nuclei, Hypernuclei, and their Electromagnetic Properties
Ding	Universal properties of the chiral phase transition in 2+1 flavor QCD using Highly Improved Staggered Quark action
Edwards	Dynamical Anisotropic-Clover Lattice Production for Hadronic and Nuclear Physics
Engelhardt	Nucleon transverse momentum dependent parton distribution functions on a large isotropic clover fermion ensemble
Fleming	Lattice Gauge Theory for Physics beyond the Standard Model on Leadership Class Machines
Hasenfratz	Eight Flavor SU(3) gauge theory with nHYP smeared fermions
Izubuchi	Hadronic vacuum polarization and hadronic light-by-light contributions to the muon anomalous magnetic moment using statistical error reduction techniques
Kelly	Lattice Determination of the Delta I= 1/2, K -> pi pi Amplitude
Kuti	Can the Higgs impostor hide in the BSM sextet model
Lin	Probing TeV Physics through Neutron-Decay Matrix Elements
Liu	Nucleon Structure with Overlap Fermion
Mackenzie	CKM Physics from B, D, and K Mesons with HISQ Fermions
Mawhinney	Pion and Kaon Physics from 2+1 Flavor DWF Lattices with m_pi = 140 MeV and V=(5.5 fm)^3, II
Meinel	Disconnected Contributions to Nucleon Ground State Structure
Mukherjee	Continuum limit of higher-order charge-fluctuations at the physical point
Syritsyn	Nucleon Structure from Domain Wall Fermions at physical pion mass
Neil	Two color gauge theories in the LHC era
Orginos	Isotropic Clover Fermions
Richards	Excited Meson and Baryon States using Anisotropic Clover Lattices
Shigemitsu	High-Precision Heavy-Quark Physics
Sugar	QCD with Four Flavors of Highly Improved Staggered Quarks
Witzel	B-meson physics with domain-wall light quarks at their physical mass and relativistic heavy quarks

Junior staff job creation

- In the last few years, job creation has been good. About 25 permanent jobs in the US in the last 10 years; 30 or more would be a more comfortable number.
- We are working on strategies to improve it.
  - We have formed a speakers committee to find prominent speaking slots for talented young people.
    - Andreas Kronfeld (chair), David Richards, Peter Petreczky, Simon Catterall.
  - NP has had success creating university jobs with JLab bridge positions. We investigated the possibility for HEP.
    - DoE: DoE theory supports this idea in principle. They asked us to bring it to them again when the research budget stabilizes.
    - NSF: NSF avoids pushing universities in any particular direction.
    - Riken: the Japanese research organization Riken has agreed to create a joint job at the Riken BNL center joint with U of Colorado,
      - Went to Ethan Neil. BNL is invesitgating the possibility of more positions.

II. Lattice QCD

QCD is the theory of quarks and gluons. Quarks and gluons cannot be directly observed because the forces of QCD are strongly interacting.

Quarks are permanently confined inside hadrons, even though they behave as almost free particles at asymptotically high energies.

"Asymptotic freedom", Gross, Politzer, and Wilczek, Nobel Prize, 2004.

Lattice QCD is used to relate the observed properties of hadrons from the properties of their quark and gluon constituents.

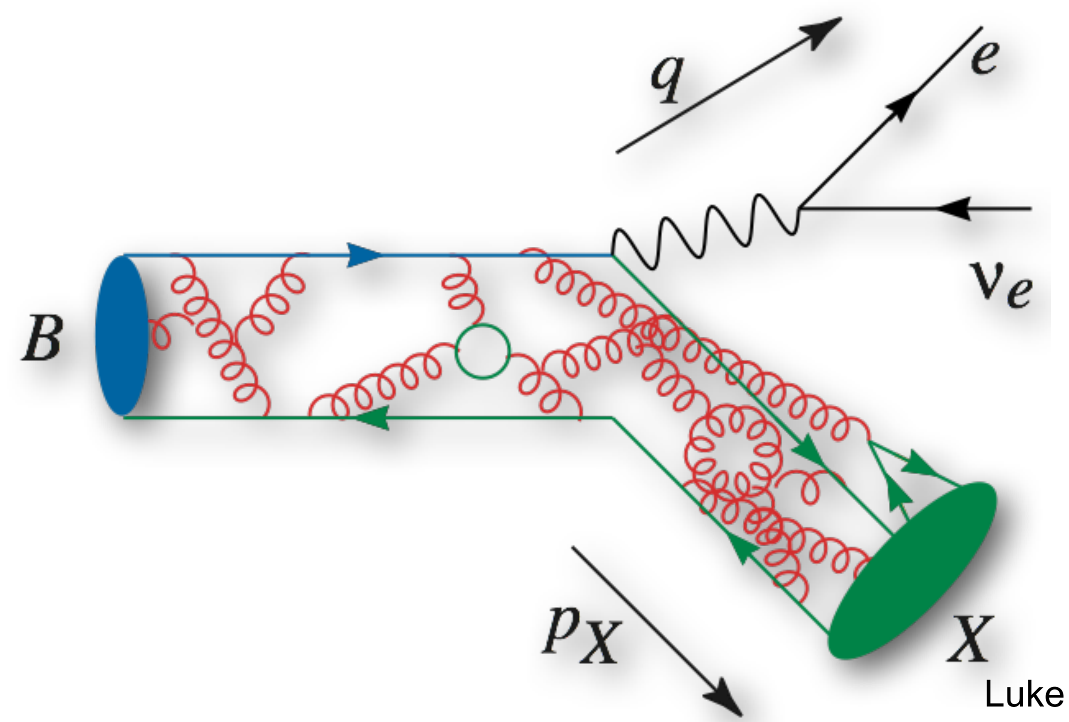

$B \to \pi l \nu$ semileptonic decay

Lattice quantum field theories

Approximate the path integral of quantum field theory by defining the fields on a four dimensional space-time lattice.

Quarks (ψ) are defined on the sites of the lattice, and gluons (U_μ) on the links.

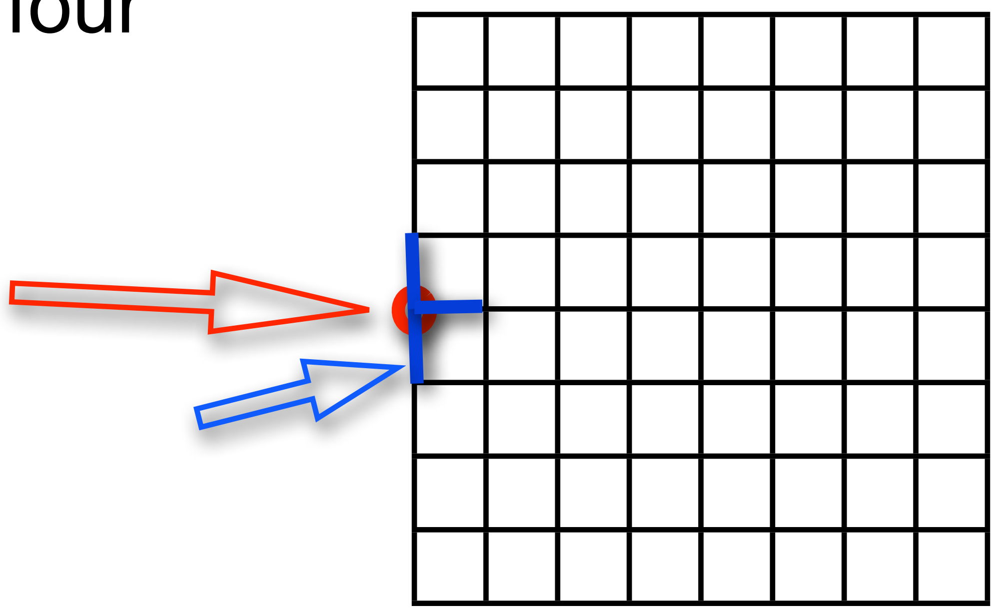

Monte Carlo methods are used to generate a representative ensemble of gauge fields. Relaxation methods are used to calculate the propagation of quarks through the gauge field.

Continuum quantum field theory is obtained in the zero lattice spacing limit. This limit is computationally very expensive.

The Dirac, or "Dslash", operator

The fundamental operation that consumes the bulk of our cycles is the solution of the Dirac equation on the lattice.

The fundamental component of the Dirac operator is the discrete difference approximation to the first derivative of the quark field on the lattice.

$$\partial_\mu \psi(x) \to \Delta_\mu \psi(x) \approx \frac{1}{2a}(\psi(x+\hat{\mu}a) - \psi(x-\hat{\mu}a)) + \mathcal{O}(a^2)$$

Quarks in QCD come in three colors and four spins.
The color covariant Dslash operator of lattice QCD is

$$D_\mu \gamma_\mu \psi(x) \equiv \frac{1}{2}\left(U_\mu(x)\gamma_\mu\psi(x+\hat{\mu}) - U_\mu^\dagger(x-\hat{\mu})\gamma_\mu\psi(x-\hat{\mu})\right)$$

γ operates on spin four-vector.

U operates on color three-vector of the quark.

The bulk of the flops envisioned in this project are consumed in multiplying complex 3-vectors by 3x3 complex matrices.