LQCD-ext II Project
2016 Annual Review
Overview and USQCD Management

Paul Mackenzie
Fermilab
mackenzie@fnal.gov

For USQCD
http://www.usqcd.org

LQCD-ext II Project 2016 Annual Review
Jefferson Lab
June 28-29, 2016

Synopsis

- Annual review of the LQCD-ext II lattice computing hardware project, 2015-19.
  - Hardware is located at BNL, JLab, Fermilab.
  - Projects funded jointly by DoE's offices of HEP and NP.
    - LQCD-ext II, total budget $14.0 M - under review today.
    - Follow-on to the LQCD and LQCD-ext Projects.
- The LQCD Project is one of several hardware and software efforts overseen by USQCD.
- USQCD is an umbrella group consisting of most US lattice gauge theorists. Its purpose is to develop the software and hardware infrastructure required for lattice gauge theory calculations.

USQCD

- Represents almost all of the lattice gauge theorists in the US; ~ 160 people.
    - ~ 100 participating in physics proposals in a typical 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.
- SciDAC software grants since 2001.
  - Essential for making effective use of Leadership Computing Facilities and our dedicated hardware, and for accomplishing our physics objectives.
- LQCD, LQCD-ext, and LQCD-ext II Projects capacity hardware grants from HEP and NP since 2005.
  - Being reviewed today.
  - Installed at JLab, Fermilab, and BNL.
- INCITE grants since 2008.
  - At DOE's Leadership-Class Computing Facilities.
  - For our largest-scale jobs; jobs that can't be done on smaller computers.

Achievements 2005-2015

- In HEP, lattice-QCD calculations played a critical role in making the flavor-physics program of the B factories and the Tevatron a success.
- In NP, lattice QCD played a key role in motivating the experimental physics programs of RHIC and the 12-GeV upgrade, and it solidly nailed down the QCD deconfinement temperature, a key quantity in interpreting the results of heavy-ion collisions at RHIC.

Opportunities, 2015-2019

- In 2015, our opportunities are broader and deeper than ever before.
- In nuclear physics,
  - Solid calculations of the hadronic resonance spectrum are needed to guide the search for hybrid states in the resonance region.
  - Calculations of the equation of state at non-zero baryon density and cumulants of conserved charge fluctuations are needed to guide the exploration of the QCD phase diagram with the heavy ion program at RHIC.
  - We are now in a position to envision the calculation of the spectra and structure of nuclei from first principles, with lattice QCD at its foundation. It will take years to make this a reality.
- In particle physics, lattice QCD calculations are needed in applications large and small throughout the coming experimental program, like perturbative QCD.

USQCD

LQCD-ext II Project

LQCD Federal project director:
John Kogut (HEP)
LQCD Federal project monitor:
Elizabeth Bartosz (NP)

USQCD Executive Committee:
Paul Mackenzie (chair), Rich Brower, Norman Christ, Will Detmold, Robert Edwards, Frithjof Karsch, Julius Kuti, Kostas Orginos, Martin Savage, Bob Sugar

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

USQCD Software Committee:
Rich Brower (chair)

USQCD Scientific Program Committee:
Anna Hasenfratz (chair)

Green: USQCD present today.

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

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

Organization

- 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 a collaboration like CMS or GlueX.
  - 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 individuals and allocated by SPC. (Role of EC in this process is analogous to that of lab director.)

See the charter of USQCD, http://www.usqcd.org/documents/charter.pdf

- This model has worked very well.
  - 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.
    - People who got junior faculty or staff jobs in the last couple years almost always served as PIs of their own proposals.
  - When groups disagree on methods, they can compete.
  - We have been very productive under this model.

Executive Committee

Paul Mackenzie (chair), Rich Brower, Norman Christ, Will Detmold, Robert Edwards, Frithjof Karsch, Julius Kuti, Kostas Orginos, Martin Savage, Bob Sugar

Present today.

- Provides overall leadership for USQCD 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.

Executive Committee

- Rotates new members at ~ one/year.
    - Close to full rotation over ~ 10 years is planned. As of this week, 2/3 of the founding members will have rotated.
    - 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 largest physics collaborations, one member from each of the three partner labs, a few members from outside these groups.
    - This year, David Richards➜Robert Edwards,
      on July 1, Bob Sugar➜Carleton Detar.

Executive Committee

- This year, we instituted a new type of position on the Executive Committee.
  - filled by election,
  - a period of two years.
- Goals:
  - provide a window into the Executive Committee for younger people,
  - provide the Executive Committee improved input from the community,
  - provide management experience for younger members of USQCD.
- Election was managed by the SPC prior to the All Hands Meeting.
- ⇨ Will Detmold was elected to the Executive Committee.

Scientific Program Committee

Anna Hasenfratz (chair, Colorado), Tom Blum (Connecticut), Will Detmold (MIT), Steve Gottlieb (Indiana), Aida El-Khadra (Ilinois), Swagato Mukherjee (BNL), Kostas Orginos (William and Mary).

Present today.

Former chairs: Robert Edwards, Frithjof Karsch, Andreas Kronfeld, Claudio Rebbi.

The SPC creates a program to accomplish the goals set forth in USQCD's computing proposals.
- It may also advise us on needed evolution of the goals.
- It examines submitted physics 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.

35 people have served so far as members of the EC and/or the SPC.

SPC allocations process

- After approval from the EC, the SPC issues the Call-for-Proposals.
- The SPC collects and reviews the proposals. Further information is often requested from the proposers.
- After deliberation, the SPC arrives at an allocation through an internal vote.
  - About 80 hours total is spent in this process.
- Recommendations for allocation are submitted to the EC for approval. The facility managers are also consulted.
- The SPC notifies the PIs and gives them a written report.
- This year, to encourage smooth use of resources, we instituted system of quarterly allocation reductions for projects that are late in beginning running (like at NERSC).

Allocations

Total by Field

All

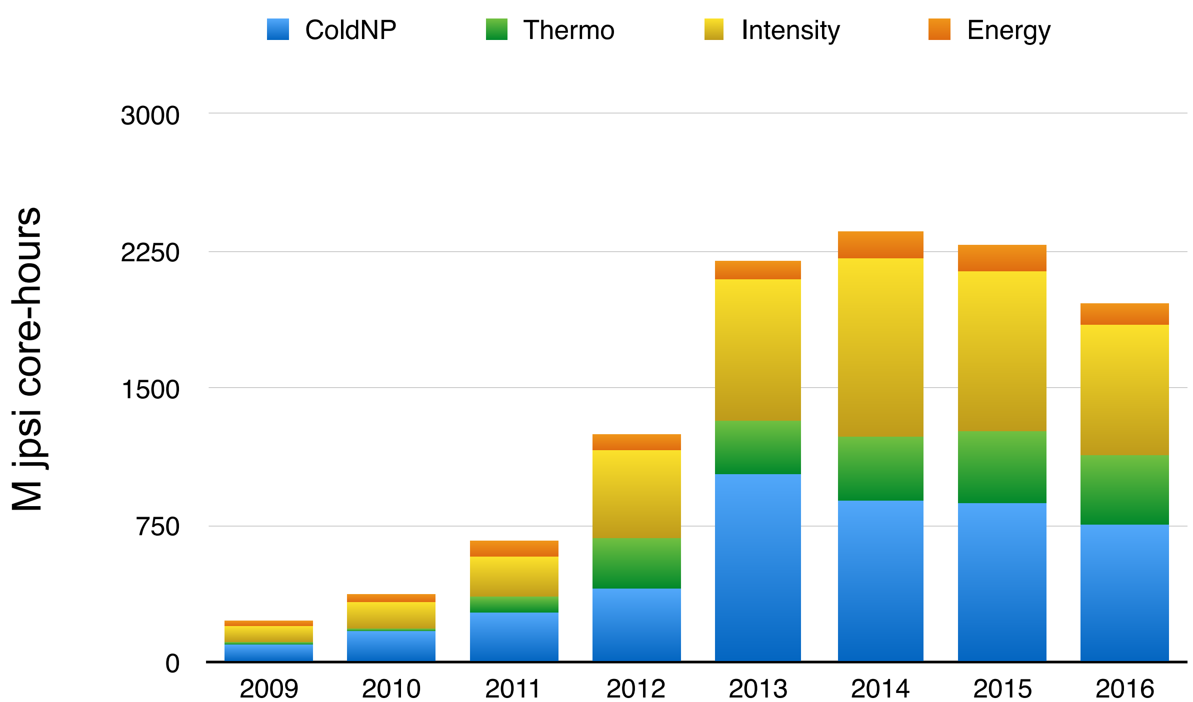

LQCD hardware

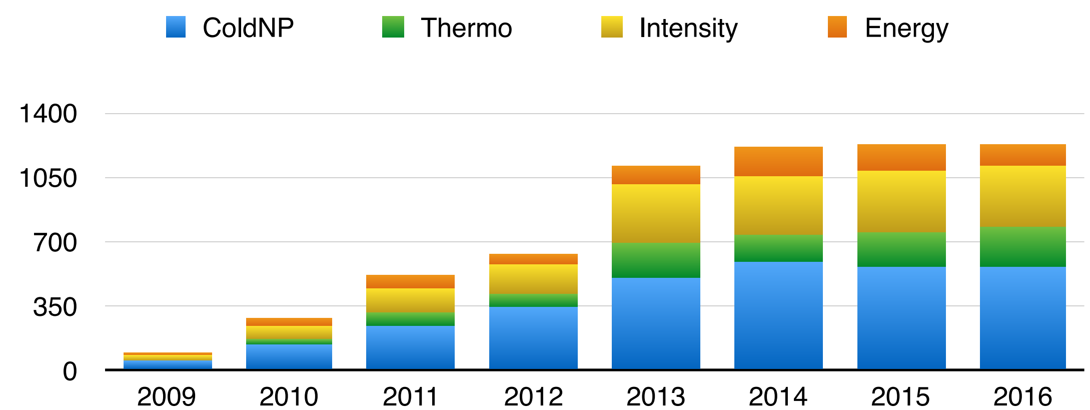

Total by resource type

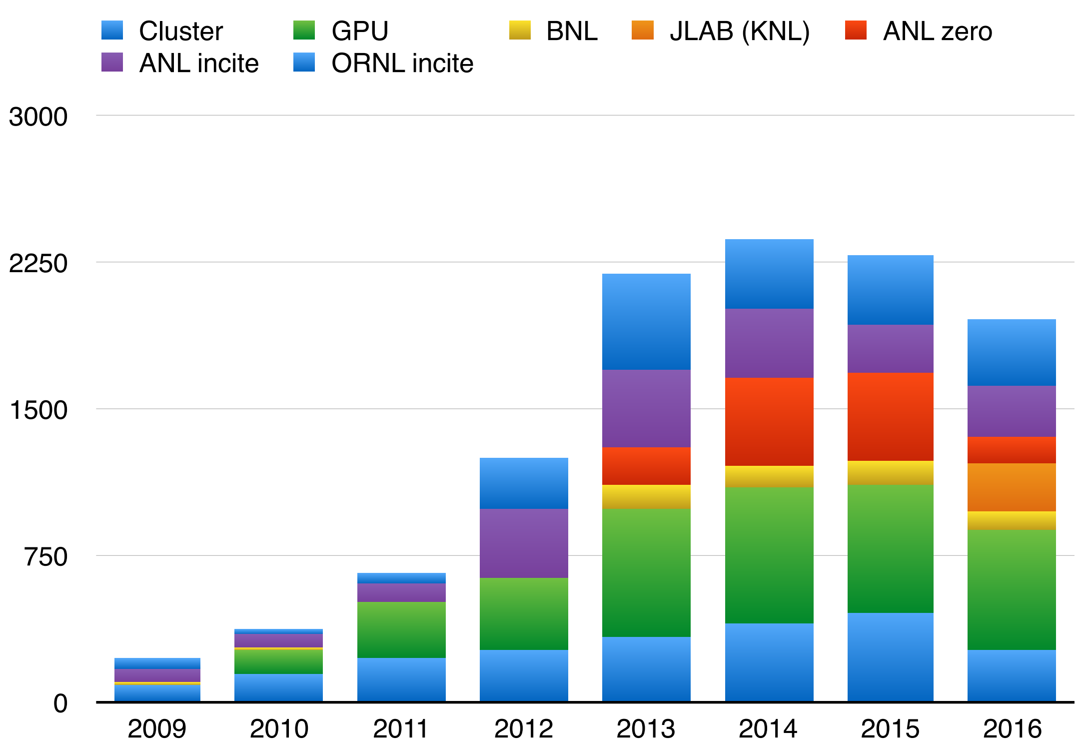

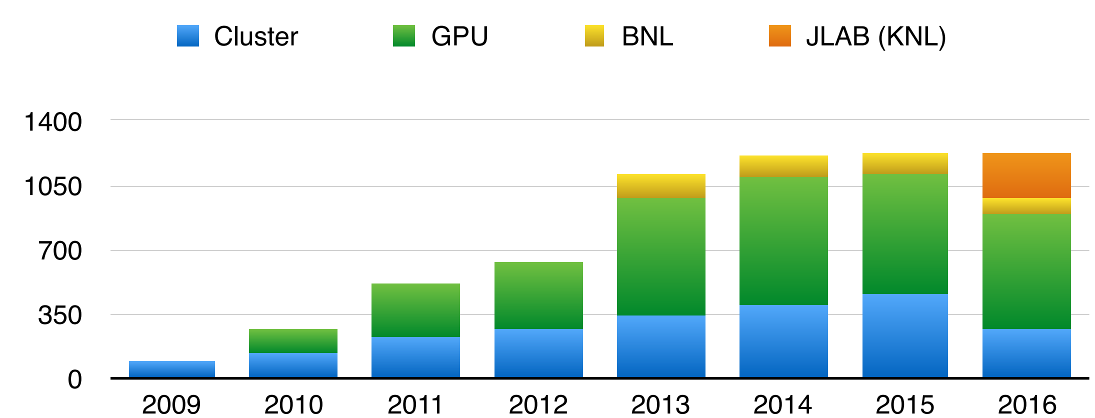

Allocations

Type A conventional cluster allocations in 2016 by project.

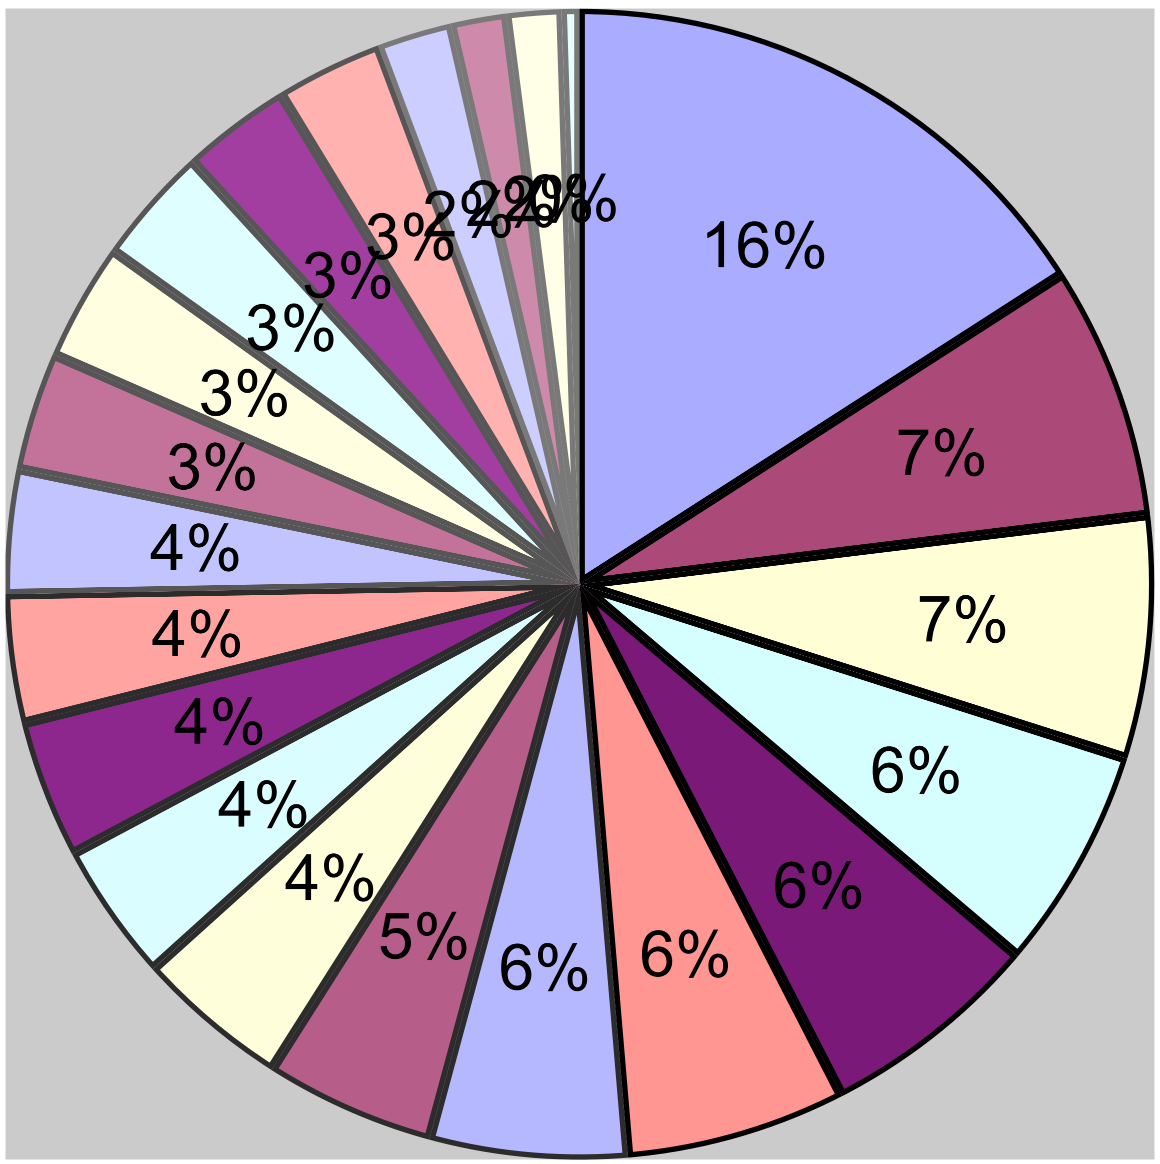

Projects are judged by:
- relevance to the central goals of USQCD;
- size and competence of project team;
- validity and efficiency of methods proposed.

Less high priority projects are typically not zeroed out, but are given less resources. About half of these allocations went to the six highest priority projects (an HEP project producing a dozen high-priority weak matrix elements, an NP project calculating resonance spectra for Gluex, two g-2 projects ...).

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

USQCD total hardware resources

USQCD resources in June, 2016.

	2016 USQCD resources					
			M units	Units	Total M core-hours	Grand total
LQCD Project	clusters	DOE/ HEP&NP			263	
	GPUs	"	9.5	GPU hrs	664	
	BNL BGQ	"			116	
	Jlab KNL	"			250	
Leadership Class	LCF INCITE	DOE/ASCR			280	
	LCF zero priority					
	LCF ALCC	DOE/ASCR			598	
	Blue Waters	NSF	17	node hrs	272	
General purpose	NERSC	DOE/ASCR				
					227	
						2670

The LQCD Project, INCITE, and Blue Waters were applied for by USQCD as a whole.
The physics collaborations making up USQCD also apply for time at NERSC, NSF XSEDE, ALCC ..., independently of USQCD.

LQCD hardware is at Fermilab, JLab, BNL

- Fermilab and JLab have focused on conventional and accelerated clusters.
- BNL
  - 1990s-2016 hosted the Columbia-designed machines that gave rise to IBM's Blue Gene line, and then the Blue Genes themselves.
  - 2015, BG line will not be continued; BNL establishes Computational Sciences Initiative, will host GPU-based institutional cluster.
- We would like BNL to begin helping host USQCD's dedicated clusters.
  - We're enthusiastic: will make our team broader and deeper.
  - Will require an official change request — see Rob's talk.
  - We hope you agree.

GPUs, many-core architectures, ...

- GPUs have supplied a significant fraction of our capacity computing needs in the last few years.
  - Were a disruptive technology: for the projects that can use them, they provide outstanding price/performance in $, but require significant investments in software and physics brain power.
  - A significant amount of high-level GPU code has now been created.
  - Many of our projects are mature on GPUs, but many are not.
  - Price performance varies much more by project than is true for ordinary clusters.
    - Harder to define a standard candle for price/performance — effective price/performance varies depending on job mix.
    - ⇒ Changed reporting from X effective teraflops delivered by GPU resources to Y GPU-hours delivered for GPUs, left conventional benchmarks the same.
    - But…

Knight's Landing

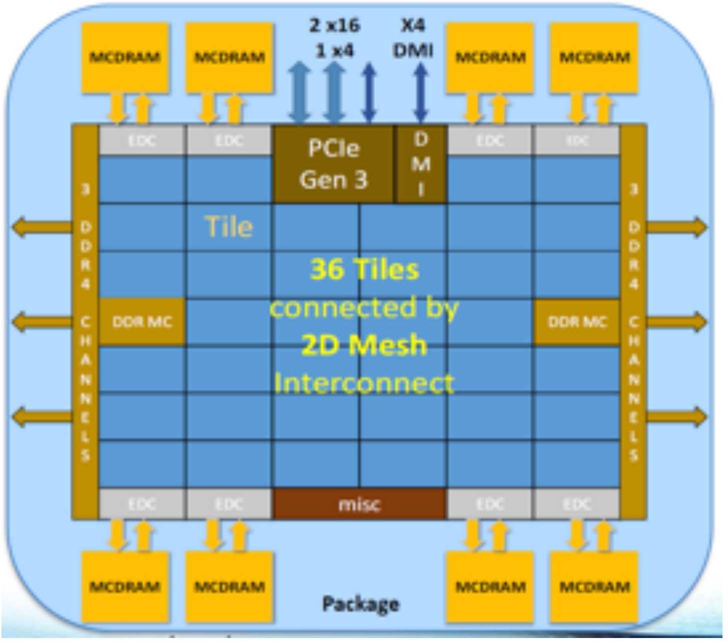

- For the 2016 LQCD hardware purchase, Intel's many-core Knight's landing processor was a contender (and the ultimate winner).
- Has some advantages of both conventional and accelerated hardware, can serve the place of both.
  - Runs the Intel instruction set, non-guru code runs adequately.
  - Also profits from guru-level code like a GPU.
  - "Accelerated benchmarks" just described are inconsistent with this kind of "accelerator". To make this kind of chip consistent with our milestones, need to go back to straight teraflops-based milestones. Also requires an official change request.

- Coming chips will continue to become more painful to deal with than the commodity chips of the 2000s.
  - Intel many-core chips are here, a new disruptive technology.
  - Future chips will continue to use more and more cores/chip, more and more layers of memory hierarchy, etc.
- We are investing a lot of effort in keeping on top of this.
- USQCD profits from close relationships with vendors.
  - Former particle physicist Al Gara is now Intel chief exascale architect; was originally part of Norman Christ's lattice QCD computing project at Columbia; was also chief architect of the Blue Gene computers.
  - USQCD member Kate Clark is a full time NVIDIA employee; still a lead USQCD developer and tests future NVIDIA architectures for performance on lattice QCD algorithms.

Software development programs

- At supercomputing labs:
  - NERSC: Three US groups have NESAPs, work with Intel and NERSC scientists and engineers to get code ready for KNL-based Cori II in 2016.
  - Argonne Leadership-class Computing Facility: Theta early science program to get code ready for KNL-based Theta.
- Exascale Application Program
  - ASCR has ~$160 M this year for its Exascale Program. Part of president's Strategic Computing Initiative.
  - EAP envisions funding 8-9 application projects @ ~$2.5 M/year.
  - USQCD was one of 25 projects asked to submit full proposal. Disposition expected in a month or two.

DOE Leadership-class 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 the INCITE Program and through the ALCC 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 INCITE grant running from calendar year 2014 though calendar year 2016.
    - Submitted new proposal for 2017-2019.
    - USQCD had 2nd largest INCITE allocation in 2016,
    - 100 M core-hours on OLCF Titan, 180 M core-hours on ALCF Mira, plus zero-priority time.
- Physics collaborations apply for ALCC time individually, received almost 600 M hours this year.

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.

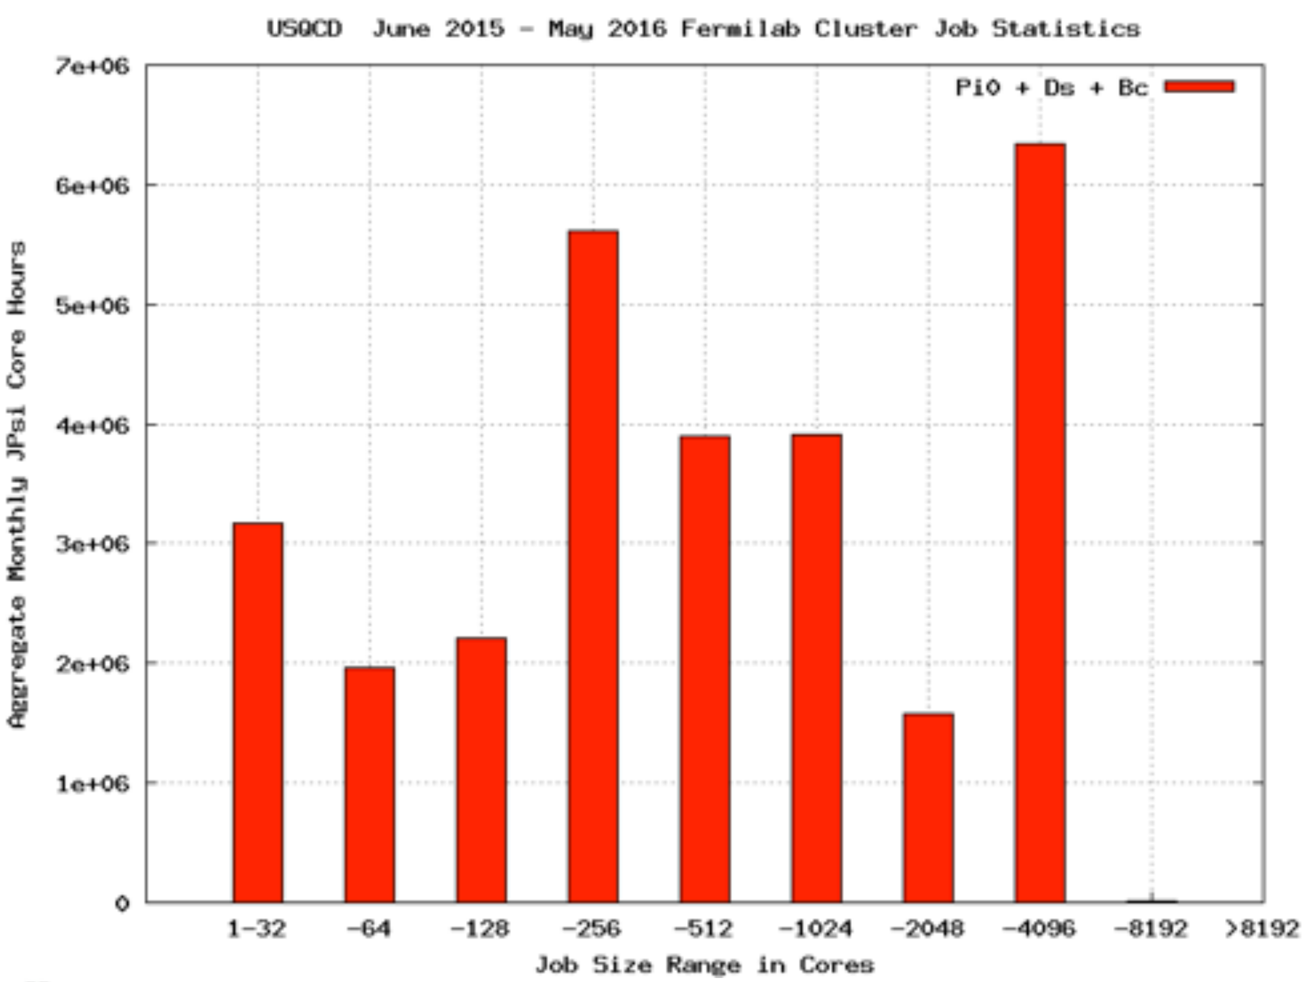

We have large computing needs at all scales of job-sizes, from one-node jobs to hundred thousand node jobs on a log scale in job-size.

Job size distribution on USQCD 2015/16 conventional clusters at Fermilab.

Capacity and capability computing

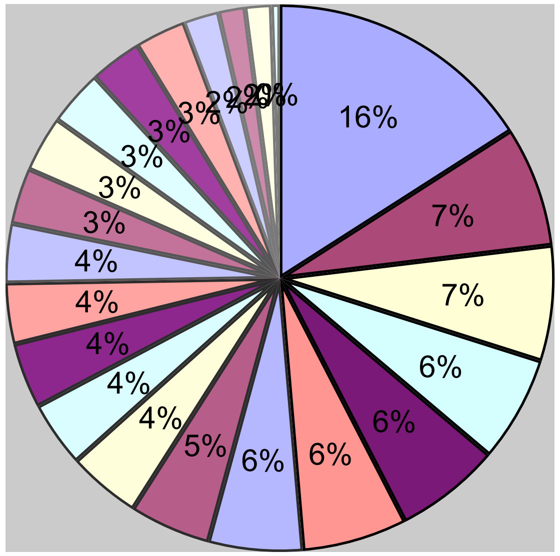

- The LCFs are only set up to handle a small number (~30) of very large projects.
  - Only simple queues are set up, user support assumes sophisticated users.
- NERSC can handle smaller projects, but competition is much worse so lattice QCD gets much less time there.
- Most of the smaller projects in our portfolio would be close to impossible to do at LCFs.
  - These include many of our most important projects:
    - Innovative calculations on new quantities. E.g, g-2 calculations 5-10 years ago.
    - Projects developing new algorithms and methods that will ultimately make the flagship project more efficient.

Capacity and capability computing

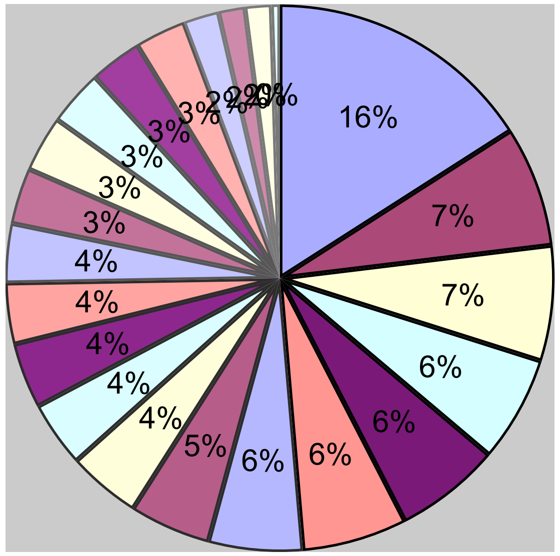

- Some of our projects consist of a large number of medium-sized jobs. By bundling medium-sized jobs together it may be possible to create jobs suitable for the LCFs.
  - This is now allowed at the LCFs — policy is evolving here.
- Inefficient in several ways:
  - Diverts human time away from physics into software and painful work-flow wrangling.
  - A hit from unnecessary large scale communications degrades performance.
  - Reliability degrades since a failure in one core crashes the entire suite of jobs.
- We're working to understand the proper role for this in our program.
- The LCFs are great for parts of our program but not for all.

USQCD hardware resources by year

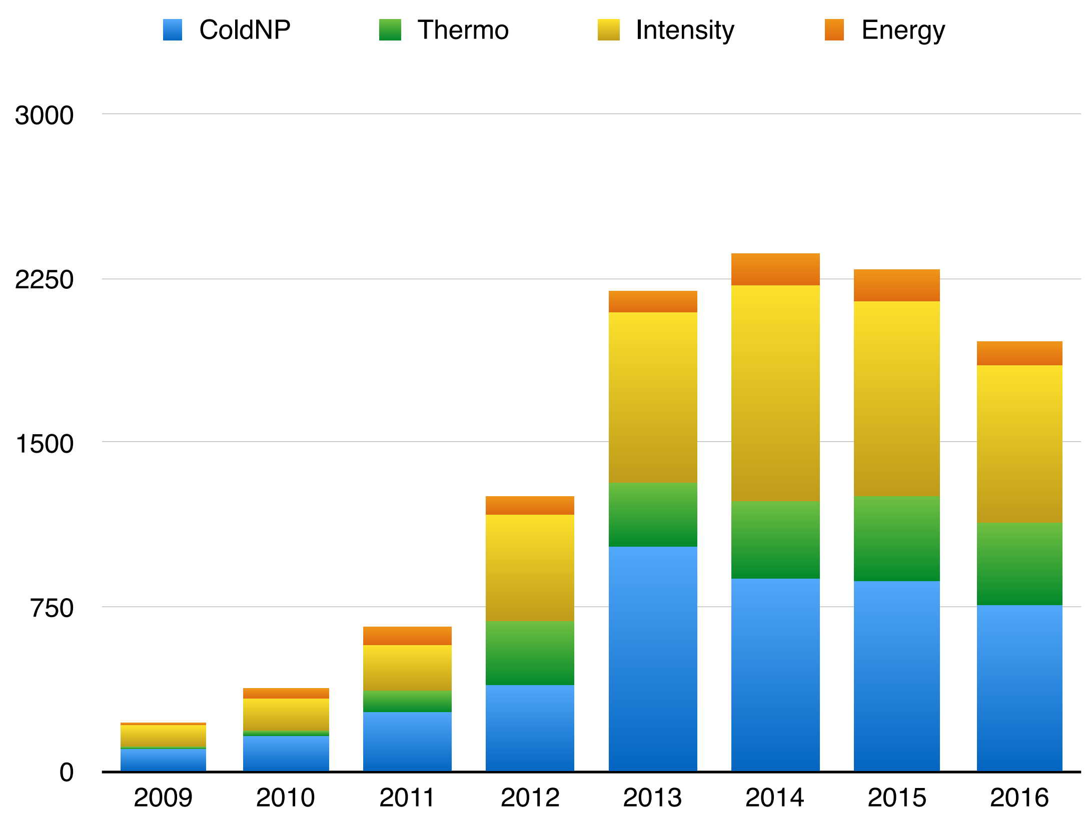

- Total USQCD hardware resources rose exponentially following Moore's law up to around 2013.
- This year's trend is flat, and will rise only slowly for the next few.

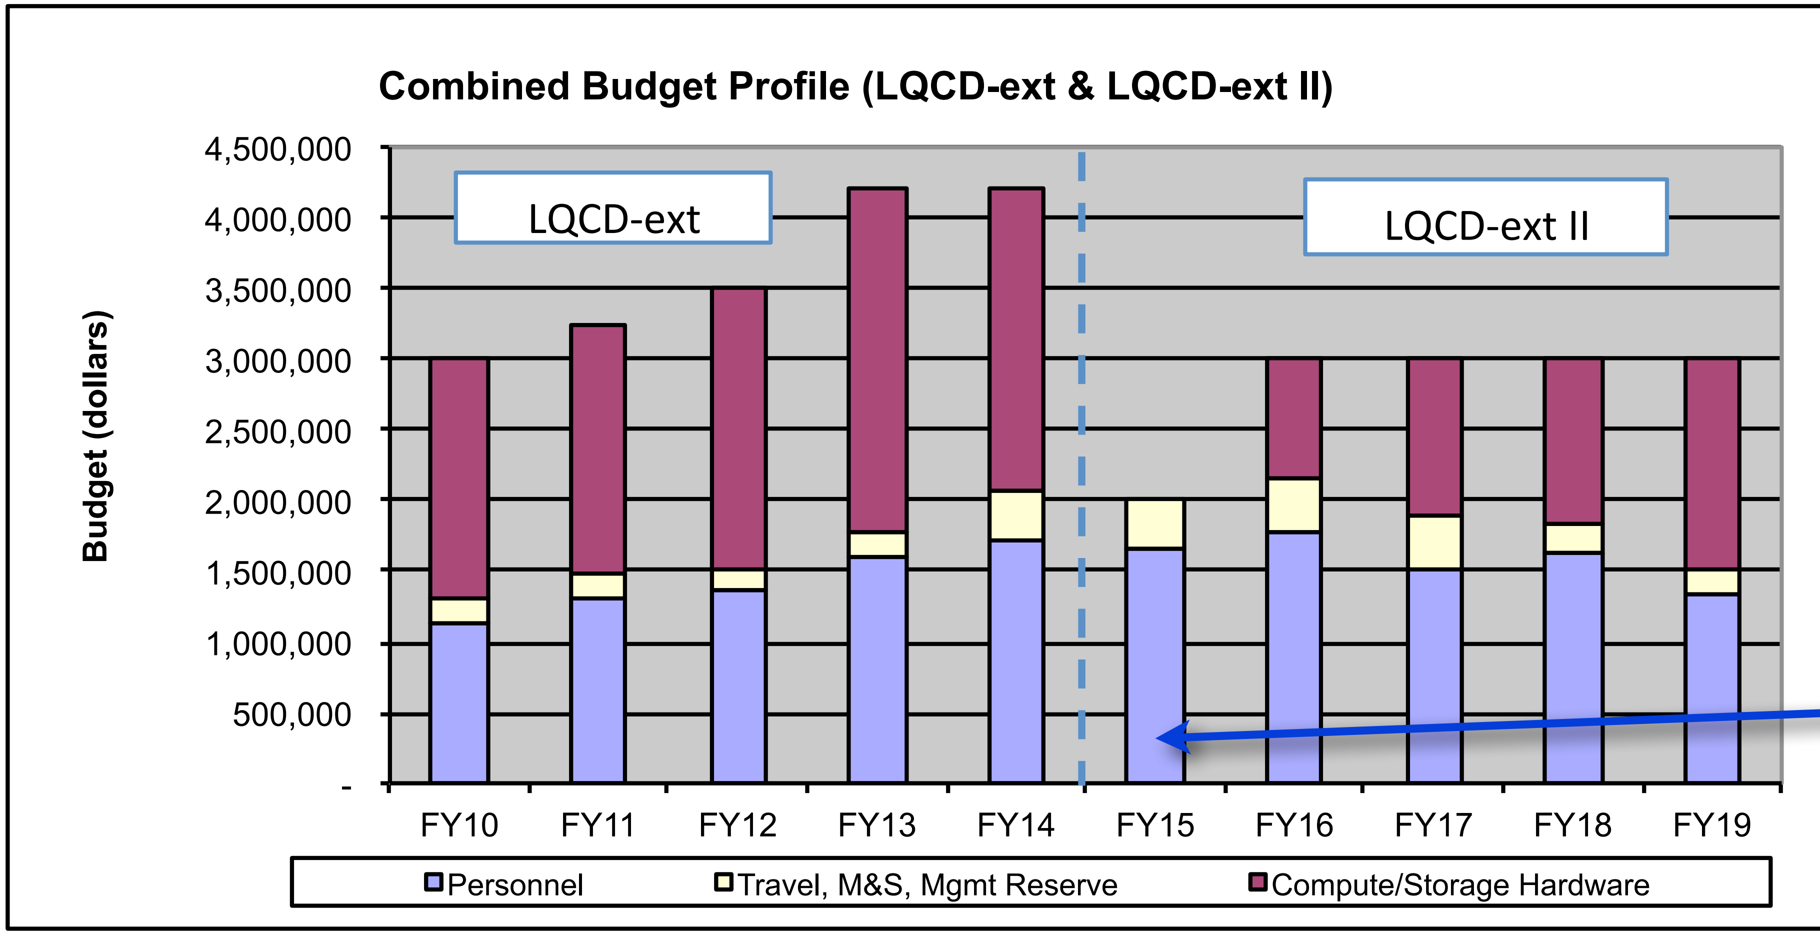

Under LQCD-ext, we spent ~ 60% of our funds on hardware, about 40% on operations. FY15 funding to allow only continued operations - no new hardware purchases.

- Current plans call for reduced and back-loaded funding for LQCD-ext II compared with LQCD-ext, particularly in 2015, a terrible year for HEP.
- The Leadership-class computers at Argonne and Oak Ridge upgraded their resources by a factor of ~ten in 2013; the next major upgrades planned for ~2017-19.
- USQCD science program will have to adjust for this fact, stretch out science goals, make sure our most important deliverables remain on track.
- Improvements in algorithms and methods will become even more critical.

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 in NP and $0.5 M in HEP 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.
  - Will become even more important over next few years.

USQCD science goals

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.
  - 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; had 23 authors.
  - Continually evolving, in consultation with the SPC.
  - Discussed by USQCD members at All Hands meetings.

Community planning

- Community planning documents are influential in formulating our goals.
  - We participate in creating them, and community needs influence our lattice program.
- We have been active in creating the NSAC long-range plan.
  - See http://science.energy.gov/~/media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf and https://www.phy.anl.gov/nsac-lrp/Whitepapers/ , especially EIC, FSNu, HotQCD, QCDHadron, LENP and RHIC-Spin documents for lattice-related material.
- We were active in HEP community planning exercises Snowmass 2013 leading up to the P5 process.
  - See http://www.slac.stanford.edu/econf/C1307292/docs/ComputingFrontier/Lattice-44.pdf .

Science Advisory Board

- We have formalized our many interactions with experimental and theory colleagues by naming a Science Advisory Board.
  - Brendan Casey (Fermilab, g-2), Daniel Cebra (UC Davis, STAR), Jesse Thaler (MIT), David Kaplan (U. Washington), Curtis Meyer (Carnegie Mellon, GlueX), Alan Schwartz (Cincinnati, Belle), Volker Koch (LBL).
  - Among the most useful advisors on white papers and proposals.
- Developed a more detailed process this year.

1. Ask for draft written report based on USQCD documents and reviews.
2. Conference call between members of SAB, EC, and SPC at which SAB can interrogate each other and USQCD leadership.
3. Revised and final written reports.
4. We did a trial procedure recently which was very successful and will do a full version in the fall before allocations begin.

Input from the SABs

- 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.
- 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."
- 2016: role of the SAB, and the goals of the BSM program.
  - "As someone who already has limited time available for traveling, I am reluctant to commit to additional activities."
  - There was much discussion on the phone conference on what were the most important goals of the BSM program. Lattice BSM work has several kinds of purposes - what are the most important?

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

2016 USQCD all-hands meeting

- Took place April 29-30 at Brookhaven. ~49 members attended, plus ~6 remote participants. (http://www.usqcd.org/meetings/allHands2016/ .)
- Reports from the Executive Committee, the LQCD-ext Project Manager, the SPC, and the hardware site managers.
- In each science domain, reports from
  - representative physics projects,
  - members of the SPC on the relation between the allocated projects and the long-term goals.

Lattice meets experiment meetings

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”.
- BNL, Feb. 26-27, 2015, “Theory and Modeling for the Beam Energy Scan”.
- LLNL, April 23 to 25, 2015, “Lattice for Beyond the Standard Model Physics”.
- BNL, June 10, 2015, “RHIC users' meeting: Beam Energy Scan workshop”.
- ANL, April 21-22, 2016, “Lattice for Beyond the Standard Model Physics”.

Junior faculty and staff job creation

	Year	Research institution, HEP	Research institution, NP	Computational scientist	Teaching college	Industry	Foreign
Sergei Syritsyn	2016		Stony Brook				
Martha Constantinou	2016		Temple				
Andrea Schindler	2016		MSU				
Huey-Wen Lin	2016		MSU				
Alexei Bazavov	2016	MSU					
Ethan Neil ***	2015	Colorado/BNL					
Christoph Lehner	2014	BNL					
Mei-Feng Lin	2014			BNL			
Stefan Meinel ***	2014	Arizona/BNL					
Hiroshi Ohno	2014						Tsukuba
Heng-Tong Ding	2013						CCNU
Andre Walker-Loud**,****	2013		Wm & Mary/JLab→LBL				
Jack Laiho	2013	Glasgow→Syracuse					
Will Detmold **	2013		Wm & Mary →MIT				
Ethan Neil ***	2013	Colorado/BNL					
Christopher Thomas	2013						Cambridge
Ruth Van de Water	2012	BNL→Fermilab					
Brian Tiburzi ***	2011		CUNY/BNL				
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/JLab					
Nilmani Mathur	2007						Tata Institute
Joel Giedt	2007	RPI					
Matthew Wingate	2006						Cambridge
Jozef Dudek **,****	2006		Old Dominion/JLab→William&Mary				
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/JLab				
George Fleming	2005			Yale			
Tom Blum ***	2003	Connecticut/BNL					
Silas Beane *	2003		UNH→U Wash.				
Total		11	13	4	3	3	7

* NSF Early Career Award
** DoE OJI/Early Career

*** RIKEN/BNL bridge positions
**** JLab joint positions

Good job creation in the last few years.

Five new jobs in last year.

In 2015: Michigan State created three new positions in lattice gauge theory.

Opportunities and Challenges, 2015-2019

- Lattice gauge theory stands today with wider and more important opportunities than ever before.
- But ...
  - In 2005-2014, our support from DoE rose. This will not continue in the coming five years.
  - In the previous ten years, Moore's law progress in chip speeds continued rapidly. This is slowing down.
  - In the previous ten years, commodity hardware performed very well for scientific calculations. Coming chips for scientific computing are becoming increasing complex and hard to fully exploit.

Backup slides

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.
    - ...

The computational challenge of lattice QCD

Lattice	Quark			Gauge	ensembles		Analysis propagators, correlators		
spacing *a (fm)*	mass *ml/ms*	Volume (sites)	Configurations	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.

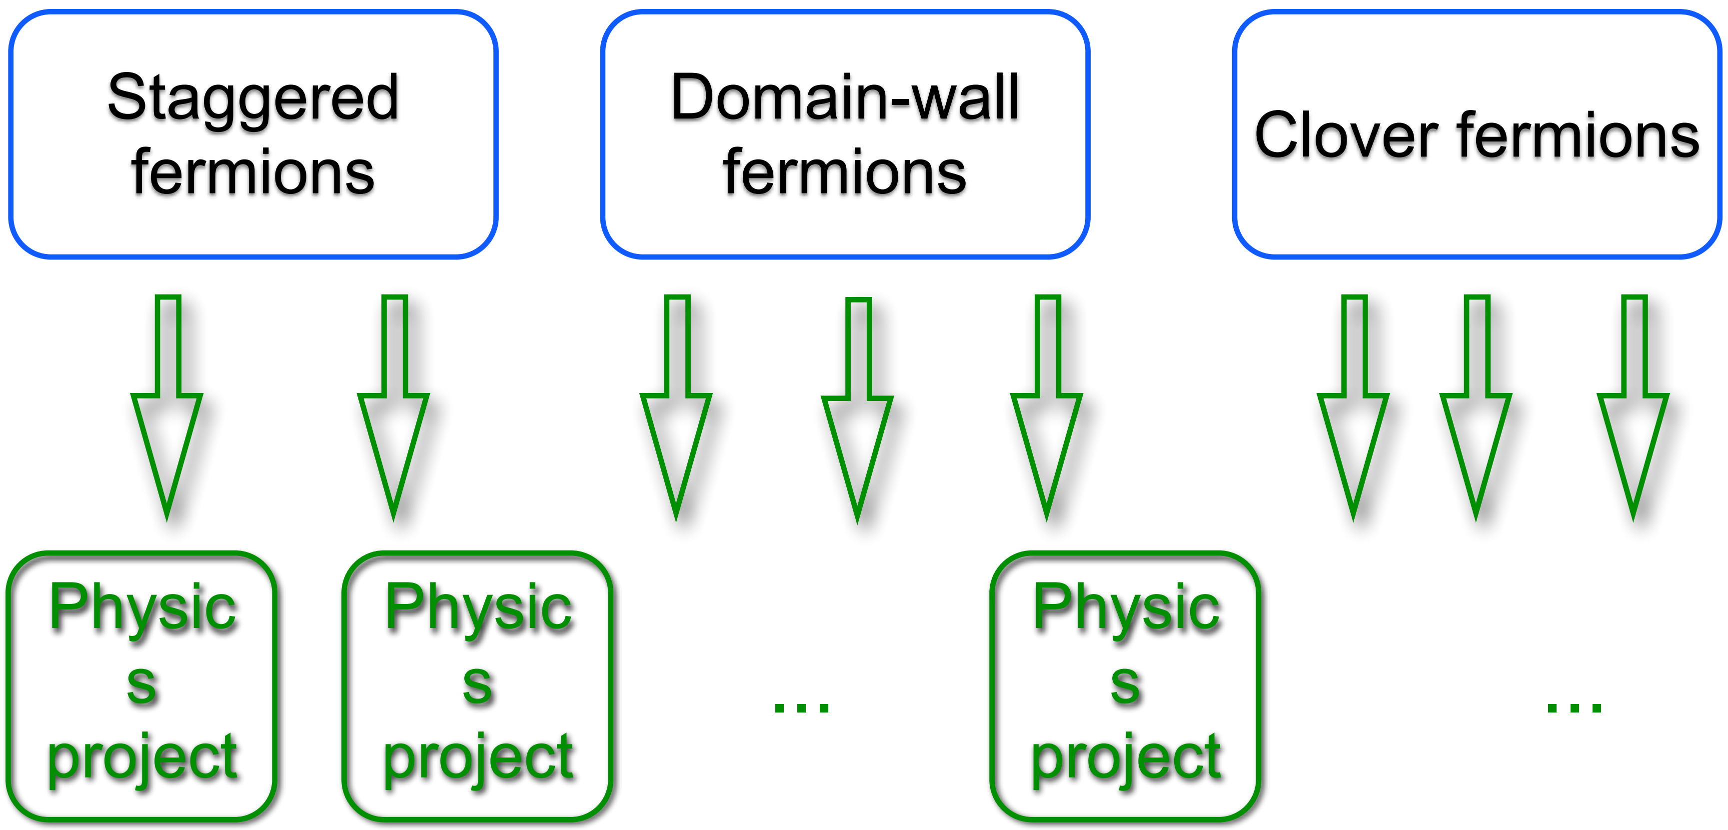

Shared among a couple of dozen groups, in both HEP and NP.

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.

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.

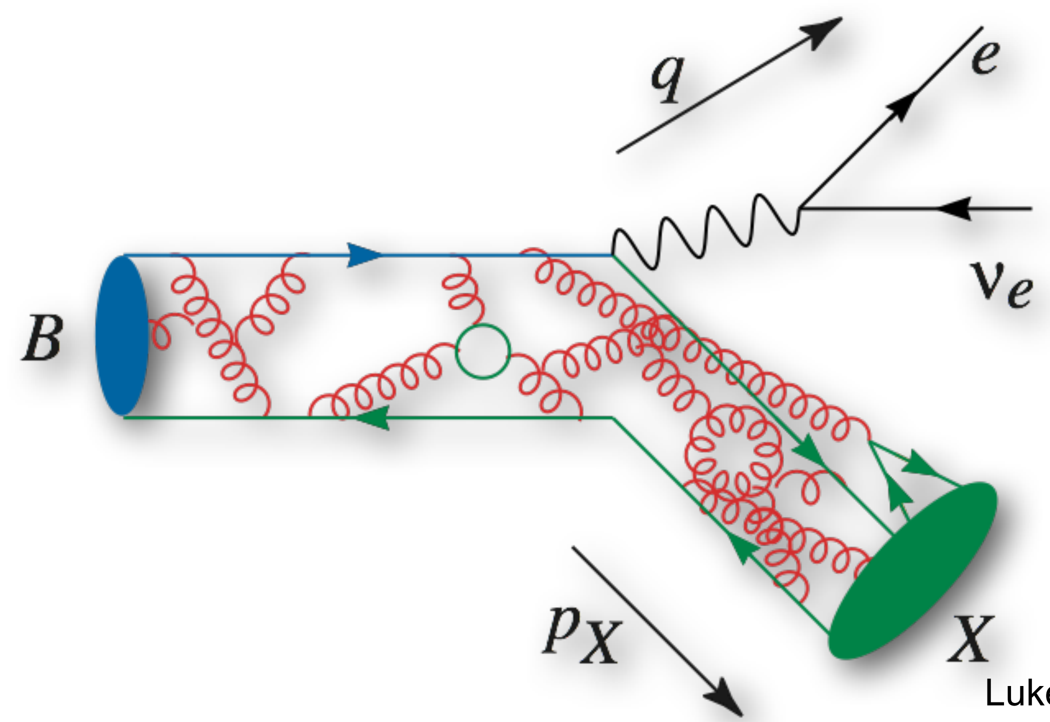

$B \to \pi l \nu$ semileptonic decay

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.

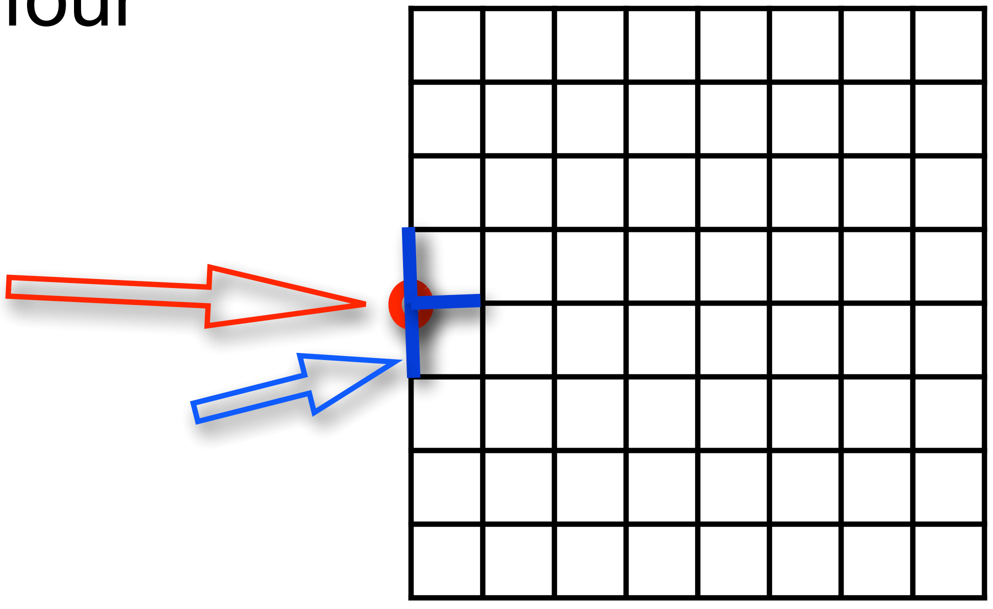

Monte Carlo methods are used to generate a representative ensemble of gauge fields. Relaxation methods are used to calculate the propagation of quarks through the gauge field.

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.