Report from the Executive Committee

Andreas S. Kronfeld
Fermilab

USQCD All Hands' Meeting
Brookhaven NL I April 26–27, 2019

USQCD Executive Committee

- Richard Brower
- Norman Christ
- Carleton DeTar
- Will Detmold
- Robert Edwards (Deputy)
- Aida El-Khadra (ex officio, SPC Chair)
- Anna Hasenfratz
- Andreas Kronfeld (Chair ⇔ Spokesperson)
- Christoph Lehner (elected junior member)
- Swagato Mukherjee
- Kostas Orginos

Scientific Program Committee

- Alexei Bazavov
- Aida El-Khadra (Chair)
- Jack Laiho
- Meifeng Lin
- Keh-Fei Liu
- Ethan Neil
- David Richards (Deputy)

- Steve, Swagato, and Tom—thank you for your work on the SPC the past few cycles.

- Type A proposals: this Call.
- Type B proposals: submit to Aida any time; response in ~1 week.
- Type C proposals: submit to site contacts; response asap:
  - BNL: Bob Mawhinney;
  - Fermilab: Jim Simone;
  - JLab: Chip Watson.
- No response? Send follow-up.

USQCD Scientific Advisory Board

- Current members:
  - Ayana Arce (Duke, ATLAS)
  - Daniel Cebra (UC Davis, STAR)
  - Lawrence Gibbons (Cornell, mu2e)
  - Krishna Rajagopal (MIT, theory)
  - Alan Schwartz (Cincinnati, Belle 2)
  - Matthew Shepherd (Indiana, GlueX, BES III)
  - Jure Zupan (Cincinnati, theory)

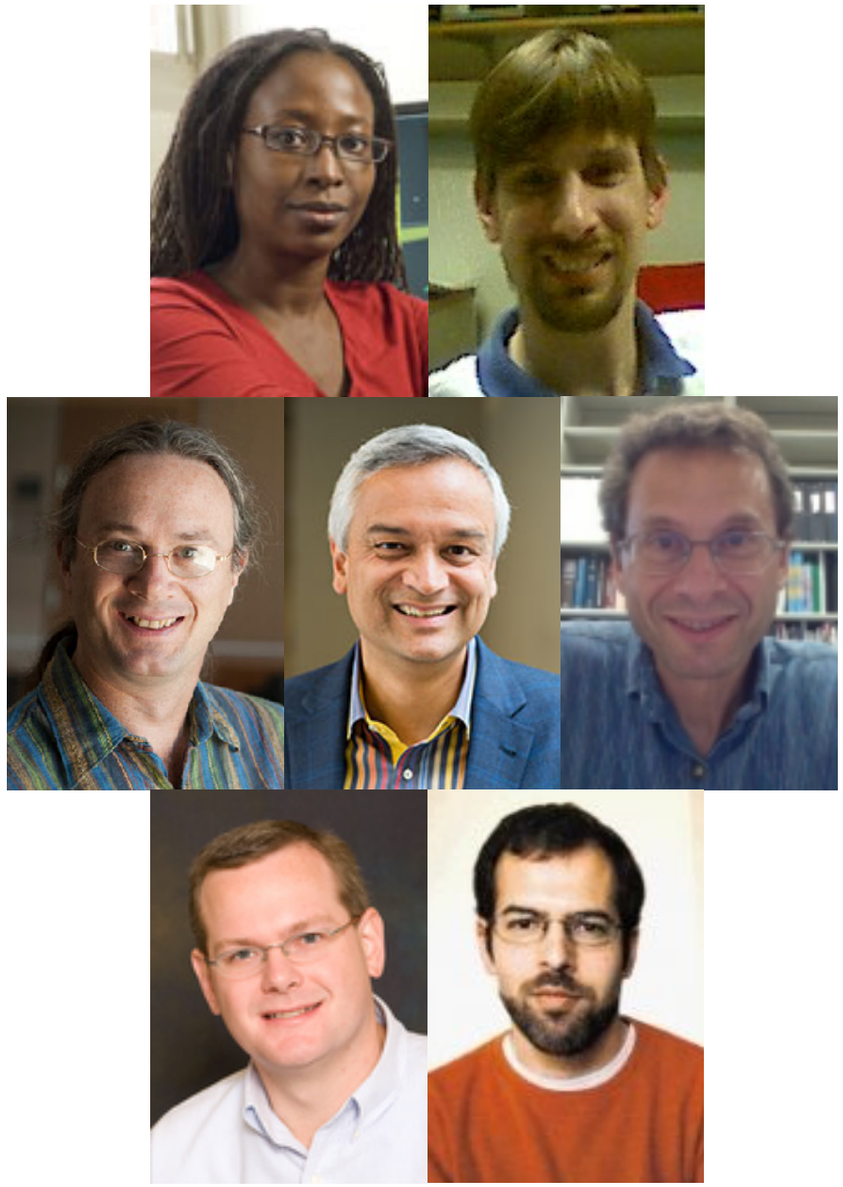

- SAB advised USQCD on the WPs' content and organization.

Nag, Nag, Nag

- When you (as PI) submit a proposal, you tacitly agree that, should you receive an allocation,
  - you and all active users on your project fill out the User Survey;
  - you will set up a web page describing the project's progress and publications;
  - you will acknowledge USQCD resources in publications.
- "Computations for this work were carried out with resources provided by the USQCD Collaboration, [other sources]. USQCD resources are acquired and operated thanks to funding from the Office of Science of the U.S. Department of Energy."

Jargon

- LQCD refers to an infrastructure project; lattice QCD means the science.
- HEP refers to the Office of HEP; particle physics means the science.
- NP refers to the Office of NP; nuclear physics means the science.
- In lattice QCD, the distinction between particle physics and nuclear physics is blurry and can be both unhelpful and helpful.
- We are accustomed to periodic boundary conditions and have to cope with stovepipe boundary conditions.

Structure of USQCD

- Executive Committee started with SciDAC support to develop software, and soon became steward of a QCDOC and dedicated clusters.

- It now encompasses

  - Initiative NPPLC;
  - Project LQCD ext. 2 (and research program LQCD ext. 3??);
  - SciDAC (NP+HEP for several cycles; now NP only);
  - INCITE allocations;
  - Exascale Computing Project (in practice, has subsumed Software Committee).

same SPC
overlapping procurement committees—
(R. Edwards, R. Mawhinney, J. Osborn)

Outline

- Not in this talk:
  - inventory of all USQCD computing resources (see Aida's talk, Bill's talk, Chip's talk, the Call for Proposals); ECP Software (see Carleton's talk).
- In this talk:
  - sharing our expertise;
  - budget FY2019, USQCD whitepapers, FY2020-2024 proposal;
  - INCITE.

Sharing our Expertise

Expertise on Computing

- Some LGT algorithms have a wide reach: HMC is used in Bayesian inference and machine learning; cross-fertilization with solvers (ECP).

- Experimenters will have to port codes to HPC platforms with GPU, Xeon Φ, Power9, etc. Collaborate on FPGAs, QPUs?

- We have experience and valuable expertise.

- There have not been many places to communicate the experience and expertise. So we should organize a workshop (or two):
  - “Lattice Meets Experiment—Computational Techniques”, or even
  - “Lattice Teaches Experiment—Computational Techniques”.
- If you are interested in being part of this, let me know.

Planning and HEP Proposal for 2020–2024

Budget Outlook

- Last year at the AHM we were worried about funding from HEP:
  - successful review led to restoration of LQCD ext. 2 funding;
  - reviewers from HEP emphasized need to do better in formulating milestones (discussion below).

- We have been encouraged to submit a Field Work Proposal from Fermilab (Bill Boroski) to secure funding for institutional clusters at BNL & Fermilab:
  - *e.g.*, $2.03M in FY2020, *= 1.03 escalation each year till FY2024;
  - presentations to HEP July 9, 10;
  - narrative to review panel 4–5 weeks beforehand.

Data Management (DM)

- Lack of DM strategy was a finding in last year's Project review.
- USQCD needs a policy statement concerning preservation and access to data used in publications.
- Robert Edwards and Jim Simone have been formulating a coherent framework for data management plans:
  - USQCD shared data (*e.g.*, gauge-field ensembles);
  - site-dependent issues and policies (*e.g.*, this page);
  - your DM plan (*e.g.*, DOE or NSF require DM plan for funding).
- See Jim's talk for more information.

Physics Narrative

- We are making a pitch to HEP, so we must speak to their interests:
  - the HEP staff has a duty to make HEP-funded experiments succeed;
  - express scientific arguments in a way that helps them do so.
- Much QCD research straddles HEP and NP:
  - no apologies; no naïveté either.
- In a few years, we presumably will propose to renew NPPLC initiative:
  - same philosophy will apply: help our NP sponsors succeed.
- Proposal narrative will be drawn from the whitepapers.

Whitepaper Coordinators

• Hot-dense Lattice QCD	Frithjof Karsch & Swagato Mukherjee
• Hadrons and Nuclei	Will Detmold & Robert Edwards
• Fundamental Symmetries	Zohreh Davoudi & Vincenzo Cirigliano
• Neutrino-Nucleus Scattering	ASK & David Richards
• Quark and Lepton Flavor Physics	Christoph Lehner & Stefan Meinel
• LGT for Physics BSM	Rich Brower, Anna Hasenfratz, Ethan Neil
• LGT Calculations: Exasale and Beyond	Balínt Joó & Chulwoo Jung

Whitepaper Coordinators

• Hot-dense Lattice QCD	Frithjof Karsch & Swagato Mukherjee
• Hadrons and Nuclei	Will Detmold & Robert Edwards
• Fundamental Symmetries	Zohreh Davoudi & Vincenzo Cirigliano
• Ne	
• Qu	
• LGT for Physics BSM	Rich Brower, Anna Hasenfratz, Ethan Neil
• LGT Calculations: Exasale and Beyond	Balínt Joó & Chulwoo Jung

Thanks to coordinators, authors, and those who provided input!

Milestones

- Guaranteeing precision for a specific quantity is problematic:
  - experiments don't do so;
  - hard to foresee how available resources (computers or luminosity).
- Could stipulate percentage of resource:
  - but once target precision is met, heeding %age could be mindless.
- Suggests hybrid:
  - stipulate %age until relevant precision is reached;
  - stipulate %ages for *exploratory* and *non-precision* targets.

Milestone Landscape

- Guaranteeing precision requires big picture:
  - key quantities such as HLbL or nucleon axial radius will require
    - LCFs for larger (finer) lattices;
    - clusters for smaller (coarser) lattices and innovation;
    - single nodes for processing correlator files.
- Thus, USQCD resources obtained via INCITE cannot (in this scheme) be decoupled from clusters.
- ALCC and NERSC (which USQCD doesn't coordinate) also influence how precision milestones are delivered.

Charge

- The critical issues to be examined in the July 9-10 review include:
  - What is the scientific case for continuing simulations of Quantum Chromodynamics (QCD) in high energy physics past 2019? Are the goals of the proposed research program aligned with the experimental and theoretical physics goals of HEP for the period 2020-2024?
  - What is the impact and interplay of lattice QCD simulations on the experimental and theoretical programs of HEP? Will the value of our experimental and theoretical programs be measurably enhanced by such simulations? Give specific examples where LQCD calculations impact the experimental program and add value to its experimental results.
  - Why is an extended project needed if the Office of Advanced Scientific Computing Research is providing the lattice community access to Leadership Class machines? In particular, is mid-scale hardware, such as CPU or GPU Institutional Clusters, essential and cost effective in such an environment? What is the optimal mix of machines, Leadership Class and mid-scale clusters, given realistic budget scenarios?
  - What are the plans at Fermilab and Brookhaven for LQCD Institutional Cluster computing? How are these plans incorporated into your proposal for the LQCD research program in 2020-2024?

Reviewers

- Maarten Golterman, San Francisco State (pheno and lattice QCD)
- Katrin Heitmann, ANL (computational cosmology)
- Patrick Huber, Virginia Tech (neutrino theory)
- Kevin McFarland, Rochester (neutrino experiment)
- Alexey Petrov, Wayne State (flavor physics theory)
- Laura Reina, Florida State (pQCD and collider pheno)
- Lee Roberts, Boston U. (Muon g-2 experimentalist)

INCITE

Facts about INCITE

Innovative and Novel Computational Impact on Theory and Experiment

- Web site http://www.doeleadershipcomputing.org/
- 2020 Call for Proposals closes at 8:00 pm EDT on June 21, 2019.
  - Renewal proposal submittal deadline is 5:00 pm EDT on July 26.
  - INCITE allocates up to 60% of IBM AC922 at ORNL (Summit) and Cray XC40 at ANL (Theta). (Mira and Titan retired from INCITE.)
- Summit has 4608 nodes; each node has 2 Power9 + 6 Volta, 512 GB DDR4 + 96 GB HBM memory; connected by dual-rail Infiniband EDR-IB (25 GB/s); peak performance of 200 Pflop s^{-1}.
- Theta has 281,088 KNL cores, 70 TByte of high-bandwidth MCDRAM, 843 TByte of DDR4 memory, 562 TByte on SSDs; peak performance of 11.69 Pflop s^{-1}.

- Open to US- and non-US-based researchers and research organizations:
  - Wuppertal-Budapest and ETM collaborations have applied.
- INCITE considers requests regardless of funding source.
- Pursue transformational advances in science and engineering.
- The INCITE program encourages community proposals.
- Each project is assigned a staff member to provide scientific support.
- 2019: 30 projects on Summit. Largest are 1050k, 1050k, 1000k, **900k**, 900k, **850k**, 850k, 740k, *700k*, 700k, 700k node-hours:
  - that means 3 of top 11 lattice QCD, 2 of the top 7 USQCD;
  - 900k node-hours (Mukherjee, thermodynamics), 850k node-hours (Christ, quark flavor physics), 700k node-hours (*CalLat*, nucleon form factors).

Proposal Strategy

- Historically, USQCD submitted one proposal covering the whole scope of computational lattice gauge theory (in the US); PI = EC Chair; 3-yr awards.

- In 2017 (INCITE 2018), USQCD submitted a second proposal focusing on hadron structure; PI = Kostas Orginos; 1-year award on Titan:

  - provided evidence supporting a conjecture (from NP) that more USQCD proposals could end up with more total INCITE resource.

- In 2018 (INCITE 2019), the EC decided to submit 6 proposals:

  - renew (now NP-less) omnibus* (PI = Mackenzie)—successful, Mira, Theta;

  - quasi-renew hadron structure (PI = Orginos)—unsuccessful;

  - 4 brand-new proposals to use Summit (see next slide).

* "Omnibus" means covering more than one of flavor, BSM, thermo, cold NP.

Four Proposals

- The four brand-new proposals were:
  - Quark flavor (Fermilab/MILC + RBC, PI = Christ)—successful;
  - BSM LGT (PI = Brower)—unsuccessful;
  - Hadrons and nuclei (PI = Edwards)—unsuccessful;
  - Hot, dense QCD (PI = Mukherjee)—successful.
- Other lattice-QCD proposals:
  - CalLat Collab.—successful;
  - ETM Collab.—unsuccessful.

Observations

- Successful proposals:
  - half of USQCD;
  - half of lattice QCD;
  - 2/3 that focusing on particle physics;
  - 2/5 that focusing on nuclear physics;
- Focused proposals were successful.
- Share of Summit going to lattice QCD is in line with past shares of LCFs.

Investigative Journalism

- The review panels are constructed to preserve institutional memory.

- A large community proposal helps reviewers understand USQCD priorities but leads to a sentiment "we know we have to fund USQCD even though [because of the 15 pp limit] we don't know exactly what they'll do"

- Last year's reviewers really liked understanding the priorities of the various lattice-QCD thrusts.*

- The review committee wrestled to identify where to cut support (for QCD and for non-QCD in HEP & NP).

- They found it too difficult to fund any of the four proposals at half strength —it wasn't clear how a half-funded proposal could still be successful.

* One reviewers told an EC member they felt we were gaming the system.

- The statement "a typical Summit award is 300k–800k node-hours" is meant to be a guide, NOT a cap:
  - a proposal that is significantly (×10) larger than 800k would receive extra special scrutiny.
- Submitting several 400k proposals is not necessarily a bad strategy (to obtain a large total) but do NOT expect all of them to be funded—USQCD is not unique, by the way, there are a few other communities that have to think about the same issues,
- An advantage of a 4M-hour proposal is that if 2M (rounding up Norman's and Swagato's total) is awarded, we can manage.
- For 2020, 3-yr awards on Summit are likely (but policy not completely settled).
- Pointing to the whitepapers in an omnibus proposal could mitigate the disadvantages of "too little information" but requires extra effort for the reviewers to read [and could annoy reviewers—gaming the page limit].

Strategies

- One omnibus proposal for 4M Summit node-hours.
- Two proposals for 2M each—nuclear and particle omnibuses.
- Four proposals for with $\sum_{t=1}^{4} P_t$ = 4M—e.g., the "usual" four thrusts:
  - very unlikely that all will be successful, so ...
  - ... balance is in the hands of the reviewers, whose decision can be altered by the flapping of a butterfly's wings.
- Many proposals, one from every science collaboration within USQCD:
  - with or without USQCD imprimatur.

Advantages and Disadvantages

- One omnibus proposal for 4M Summit node-hour:
  - USQCD will receive a large award;
  - some of us will be "taking a hit for the team";
  - we explicitly rejected this approach, because we thought there was an upper bound to a single proposal.
- Two proposals (nuclear-leaning, particle-leaning):
  - some danger that one is not funded, so each one will have to speak to the needs of the offices;
  - some of us will be "taking a hit for the team".

Advantages and Disadvantages

- Four proposals:
  - some will be successful, some won't;
  - large (800k–1M) proposals may lead to the current situation;
  - if 3-year awards are obtained, the unsuccessful can try again next year.
- More than four proposals:
  - proposals could naturally be focused;
  - natural division is among science collaborations within USQCD:
    - undermines rationale for USQCD, complicates milestone formulation.

Summary and Outlook

- Thanks to John Kogut and Elizabeth Bartosz for their leadership in the Offices, and to George Fai and Bill Kilgore for their support.
- USQCD has to continually sharpen the case for funding: “Be relevant!”
- Understand what “relevant” means by interacting with experimenters not only on physics, but also on computing.
- Collaboration work (thanks in advance for your help):
  - proposal preparation (INCITE and HEP);
  - procurement advice;
  - working groups (DM, QIS).

Questions and Discussion

Backups

Working Groups

- EC is forming two working groups to address some overarching issues.

- Quantum information science (QIS) is becoming prominent throughout the Office of Science, e.g., Dear colleague letter from Stephen Binkley (2017), and signed legislation in January 2019:

  - Martin Savage has kindly agreed to lead a WG to explore synergy between LGT and QIS;

  - see Yannick's talk.

- All proposals these days require a data management plan. Last review of USQCD noted the lack of a collaboration-wide plan:

  - Robert Edwards and Jim Simone are developing a framework.

Allocations with Two Hardware Projects

- Both Offices (according to > 1 person per Office)
  - find the USQCD allocation process to work well;
  - want the SPC to allocate its computer resource together with the other Office's;
    - encourage the SPC to neglect the source of funding.
- That said, both Offices will want to see high-quality, relevant results:
  - "relevant" depends on perspective.
- We have many calculations that are relevant to both Offices, perhaps for slightly different reasons: we need to make the most of this.

NP Funding

- This report has not talked about NP funding yet.
- NP Office understands the centrality of lattice QCD to fulfill its mission.
- The NP hardware project is new. To reiterate some things—
  - the Offices want to keep the science unified under USQCD auspices;
  - the science of USQCD is *de facto* reviewed with the hardware;
  - thinking about coordinating hardware reviews has begun, but details need to be understood;
  - two distinct models (IC & dedicated hardware) in one review vs two reviews.

- A challenge is the desirable goal of the same level of funding from NP as from HEP:
  - without lowering HEP funding.
- NSAC has laid out milestones that require lattice QCD.
- Are there persuasive arguments to go beyond satisfying these milestones?
- HEPAP hasn't laid out such milestones; the argument then is phrased as "the interpretation of an approved requires lattice-QCD calculations".
- Can any of this be adapted for the NP Office?
- An advantage of the new leadership setup is that the EC (and hence the collaboration) will be led by two people with connections to both Offices.