Office of High Energy Physics

Report on the

LQCD-ext III Science Review

July 9-10, 2019

Executive Summary

The Science Review of the USQCD collaboration's plans to extend their mid-scale computing research program LQCD-ext II to a third five year period, 2020-2024, was held on July 9-10, 2019, at the Cambria Hotel in Rockville, MD. The purpose of the review was to assess the science goals and the computing strategy presented in the proposal "Computational Resources for Lattice QCD: 2019-2024" (FWP FNAL 19-11). In particular, the LQCD-ext II team was instructed to address four charge points:

1. 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 the Office of High Energy Physics (HEP) for the period 2020-2024?
2. 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.
3. Why is an extended project needed if the Office of Advanced Scientific Computing Research (ASCR) is providing the lattice community access to Leadership Class machines? In particular, is mid-scale hardware, such as Computer Processor Units (CPU) or Graphical Processor Units (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?
4. What are the plans at Fermilab and Brookhaven National Laboratory (BNL) for LQCD Institutional Cluster computing? How are these plans incorporated into your proposal for the LQCD research program in 2020-2024?

Seven expert reviewers from the high energy physics and computer science communities heard presentations by the LQCD-ext II team that addressed the charge points 1-4. In addition, the review team was asked to either endorse the present funding scenario of $2M/year (with annual escalations tied to inflation) or to modify this level based on additional needs or deficiencies of the LQCD effort. The review team was very favorably impressed by the LQCD-ext III proposal and the presentations at the July 9-10 review. The review team unanimously endorsed the proposed $2M/year plan but directed HEP to consider additional funds for storage needs at an approximate level of $200,000/year. This recommendation was based on information provided by USQCD at the review which demonstrated the value to the world wide lattice gauge theory community of sharing lattice configurations in an open fashion. Although the physics goals of the next five years will be considerably different from those of the past five years, the LQCD-ext III team convinced the reviewers that their plans presented in their proposal were well aligned with the HEP experimental program and would, in fact, be unique and vital to those interests. The computational strategy of the team was strongly endorsed by the reviewers and the essential need of mid-scale computing resources to analyze the gauge configurations produced by simulations on Leadership Class machines was validated.

Introduction and Background

The Department of Energy Offices of ASCR, HEP and Nuclear Physics (NP) have been involved with the National Lattice Quantum Chromodynamics Collaboration (USQCD) in hardware acquisition and software development since 2001. The original Lattice Quantum Chromodynamics (LQCD) information technology (IT) hardware acquisition project LQCD started in 2006 and ran through 2009. The scientific goal of this project was to use dedicated computer hardware to produce predictions of QCD for quantities relevant to the experimental programs supported by HEP and NP with sufficient accuracy so they could be used to support the analyses of those experiments in heavy quark physics, heavy ion collisions and light quark spectroscopy. This project was very successful and was extended another five years, 2010-2014, and was renamed LQCD-ext. The lattice gauge theory field matured enormously over this second period and its further extension, LQCD-ext II, 2015-2019, is producing predictions for strongly interacting phenomena that rival the results of several ongoing HEP experiments. Since lattice simulations provide the only systematic and potentially exact theoretical calculations in the field, its results are unique and influential. These points were validated annually at the project’s annual progress reviews. These progress reviews have strengthened the project in several ways:

1. The allocation process that USQCD uses to assign science projects to the hardware purchased by the LQCD team has become more transparent and successful at supporting projects that are well aligned with the experimental programs of HEP and NP.
2. The governing bodies of USQCD have become more transparent and democratic in support of the entire lattice community of theorists.
3. The procurement process of LQCD has been validated annually and was expanded to consider GPU based clusters and other “disruptive” technologies which have led to the project’s continued success in reaching and surpassing its milestones by wide margins.
4. The computer strategy of USQCD where they produce large ensembles of gauge configurations on Leadership Class supercomputers and then analyze those configurations on the special purpose dedicated clusters constructed under the IT hardware project was reviewed and validated annually.
5. The influence and acceptance of lattice methods have grown considerably through the initiation of workshops with other communities and through the development of “speaker bureaus” etc.
6. The recent change in the project, directed by HEP, to develop and use Institutional Clusters at Fermilab and BNL has had an impressive start.

Details of the annual progress made by LQCD can be found in their annual progress review reports and will not be reproduced here aside from the points listed above. Those reports also contain information on the administrative structure of USQCD, the procurement strategy of the hardware project, etc., which is essential to understanding the issues being reviewed here.

The LQCD-ext II team has now submitted a proposal for another five year extension, "LQCD-ext III: Computational Resources for Lattice QCD: 2020-2024" (FWP FNAL 19-11), which is the subject of this scientific review. Since the scientific landscape in HEP will be very different in the coming five year period, one of the central goals of the review was to validate that the LQCD-ext II team's plans will continue to be aligned with the experiments planned by HEP for this period. The proposal was supplemented with seven additional white papers where individual lattice research areas presented their five year plans. These white papers were available to the review team several weeks before the July 9-10 review and provided depth and detail to the group's proposal for continuing capacity computing. These white papers have also appeared on the arXiv and several publication services and journals have expressed interest in devoting entire issues to their dissemination. It has become a tradition for USQCD to present their detailed five year plans in this manner and the reception by the research and publication communities has been very positive.

Over the next five years HEP's domestic experimental program will focus on the Intensity Frontier which will consist of the muon g-2 experiment, the muon-to-electron (Mu2e) conversion experiment and a program of neutrino oscillation experiments leading up to the Long Baseline Neutrino Facility (LBNF/DUNE). In flavor physics, the US will collaborate with Japan's upgrade of its KEKb accelerator and detector BELLE II programs which plan to extend b-quark physics measurements of the earlier BaBar program at SLAC and the BELLE program at KEKb by a 40-fold increase in luminosity. At the Energy Frontier, the US will continue in their ATLAS and CMS collaborations at the LHC where run 3, scheduled to begin after a two year shutdown to upgrade both detectors, should extend the search for new states in the TeV energy range very significantly. In addition, the LCHb detector at the LHC will study properties of heavy quarks to high precision and will search for new composite states of QCD.

Another central goal of the review was the validation of the computational model of the LQCD research program. LQCD simulations are typically done by generating gauge configurations on Leadership Class machines ("capability" computing) and then calculating matrix elements on clusters ("capacity" computing). Over the last two years, the computing model for mid-scale hardware available to the project has changed. HEP has directed both BNL and Fermilab to build and operate Institutional Clusters to serve its experiments and theorists. The LQCD project is now expected to work within that model and help design and construct Institutional Clusters that should run lattice codes efficiently but also serve the wider experimental community as well. The purpose of Institutional Clusters is to economize, serve a broad class of physicists and their experimental and theoretical needs, and still deliver the mid-scale computing needs of the very intensive computing challenges of the lattice program. The computing model for this USQCD-led program is: (1) gain access to large allocations on Leadership Class machines to produce ensembles of gauge configurations, and (2) use these ensembles for a large number of calculations of matrix elements and scattering processes which involve current operators, quark propagators, etc., to make predictions relevant to the HEP experimental program. The calculations of the second step are well matched to the power and architectures of Institutional Clusters. Those clusters can be designed in various ways, including GPUs, CPUs with GPU accelerators, etc. These matric element calculations are done most cost effectively on such clusters where groups of processes can be assigned to different projects and the huge gauge configurations can be parsed through the cluster's memory using their considerable I/O bandwidth. Since the cluster hardware evolves rapidly, dependent on the commercial market of PC's, tablets, smart phones and gaming consoles, the computational model of the LQCD projects which have relied on a balance of Leadership Class machines and commodity clusters must be re-validated and possibly updated on a regular basis.

The charge points presented to the project were:

1. What is the scientific case for continuing simulations of 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?
2. 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.
3. Why is an extended project needed if ASCR 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?
4. What are the plans at Fermilab and BNL for LQCD Institutional Cluster computing? How are these plans incorporated into your proposal for the LQCD research program in 2020-2024?

The review began when Andreas Kronfeld, the spokesperson for USQCD, and Richard Edwards, his deputy, gave overviews of USQCD, and presented its organizational and governance structures. USQCD acts as a federation of lattice collaborations and organizes the community and its proposals for computing resources, so that all allocation processes and the resulting research efforts are fair, balanced, productive and efficient. These items are subject to annual LQCD reviews and over the years HEP review teams have influenced all aspects of USQCD, its governance and allocation processes as well as the particulars of the LQCD mid-scale computing effort. This review team considered and assessed these topics in detail. The presentations by Dr. Kronfeld and Edwards were followed by four science talks given by Ruth Van de Water (Flavor Physics), Andreas Kronfeld (Neutrinos), Zohreh Davoudi (Fundamental Symmetries) and Ethan Neil (Beyond the Standard Model). These four talks were the heart of the review since they argued that capacity computing will be necessary for USQCD to execute calculations that will guide and challenge the ongoing HEP experimental program over the next five years. The next three talks turned to management and project issues. Bill Boroski, the contract project manager of the mid-scale computing research effort, presented the developments of FY 2019 operations and discussed the transition to the Institutional Cluster model. Andreas Kronfeld then presented the computational requirements and milestones of the research plan. The plan showed the computation resources needed to reduce the uncertainties in particular lattice calculations in order to impact several experimental measurements ongoing or planned at the HEP laboratories. Bill Boroski then followed with a presentation on the LQCD effort to employ the new Institutional Cluster model for mid-scale computing. He explained how the research effort interacts with the computing divisions at Fermilab and BNL to design clusters which can satisfy the needs of the lattice community while still being usable by a wider set of researchers. Andreas Kronfeld then completed the presentations with a summary and general outlook for lattice gauge theory over the next five years.

The following sections of this report summarize the reviewers' comments on each of the four charge points.

Charge 1

Findings

The collaboration presented four science talks describing the anticipated program over the next 5 years. The talks covered quark and lepton flavor physics, neutrino-nucleus scattering, fundamental symmetries, and beyond the standard model physics.

The collaboration presented the connections of their science priorities to the P5 priorities.

Comments

The scientific case for continuing lattice simulations of QCD remains compelling past 2019. The review panel commended the collaboration on developing a balanced scientific portfolio that is well aligned with the P5 priorities for the next five years.

Lattice QCD plays an increasingly important role in theoretical developments at both the intensity and energy frontiers, as it currently presents a first-principle systematically-improvable approach to computation of a wide class of observables. USQCD laid out a logical approach to addressing the problems related to interpretation of experimental data: from the most immediate (EW+Higgs physics, flavor physics, and (g-2)-related computations), to intermediate (computations in neutrino and lepton-flavor violating physics, PDFs, and Dark Matter detection), to future (BSM lattice computations) problems.

The collaboration presented an interesting mix of projects some of them being more conventional and others being rather speculative (e.g., stealth dark matter). The goals outlined in the whitepapers align with the experimental and theoretical physics goals of HEP, ranging from calculations of hadronic quantities that are directly needed for the interpretation of experiment, and thus are indispensable (strong coupling, quark masses, flavor physics, muon g-2 and b physics), to more indirect calculations that constrain models currently used in the interpretation of experiments (neutrinos), to theoretical investigations where lattice techniques complement other theory investigations (BSM).

Recommendations

None.

Charge 2

Findings

The collaboration presented several cases showing where their program will impact the experimental programs.

The collaboration prepared 7 white papers outlining the underlying science, 4 of which are directly connected with the HEP experimental program (with the other 3 categorized under nuclear physics).

Comments

The whitepapers are effective in describing the interconnections between the proposed LQCD research program and specific experimental and theoretical investigations outside that program.

Review and revision of these whitepapers on a regular basis, such as every five years, is valuable both inside and outside the LQCD community.

The muon g-2 task is of direct relevance to the experimental muon g-2 program at Fermilab, because without a firm theoretical prediction the measured value in itself has little to teach us. LQCD will probably produce the best calculation of the hadronic light-by-light contribution to g-2 on the relevant time scale. It is less clear that it will also deliver a competitive computation of the Hadronic Vacuum Polarization (HVP) contribution before experimental results become available. However, even a 0.5-1% computation would already constitute an important validation of the dispersive result. Over the next few years, the precision of the lattice calculation of the lowest-order HVP contribution is expected to reach that of the dispersive approach. Also, members of USQCD have invented a method for combining lattice and dispersive techniques that combine the strengths of the two methods.

The extraction of alpha_s and heavy-quark masses (m_b, m_c) from lattice calculations with a percent level accuracy is of great relevance to precision global tests of the Standard Model (SM), which include both electroweak and Higgs observables, and provide one of the most stringent constraints on physics beyond the SM (BSM). All observables are calculated including several orders of QCD and EW corrections, and the precision needed on input parameters is at least at the percent level. Recommendations from the Flavor Lattice Averaging Group (FLAG) are nowadays regularly included in all SM global fits. The precision reached by lattice calculations on light-quark masses is also impressive.

The exploratory studies of parton distribution functions (PDFs) on the lattice are very important. Lattice PDFs are complementary to existing PDF global fits and are based on a completely different approach. Their combination can provide PDFs over the entire range needed by LHC measurements and will improve the overall accuracy of predictions.

The flavor physics task is of direct relevance to the international effort in beauty and charm physics, in particular LHCb and Belle 2; without a precise quantitative understanding of the SM expectation it is impossible to discover new physics. For quark-flavor physics, USQCD is carefully tracking what matrix elements are needed to interpret experimental results, and where precision needs to be improved and can be improved, as in the calculations related to CKM matrix element extractions and rare decays.

Because of the increase of computational power, more sophisticated lattice methods, such as those based on Luscher's finite-volume approach, now allow applications of lattice QCD simulations to problems with multiple hadrons in in- or out-states that were not suited for such. In particular, calculations of kaon decays on the lattice lead to the resolution of the delta I = ½ problem, and will allow a computation of epsilon-prime with controlled errors. Members of USQCD proposed to further develop these methods to apply them to multichannel problems such as D-decays. One of the long-term goals is to fully compute D-mixing observables that could be affected by new physics.

Within the nucleon matrix element task, the computation of the axial form factor is of direct relevance to the neutrino program, including NOvA, SBN and DUNE. Having a good understanding of this form factor from theory will not reduce cross section uncertainties per se. The neutrino nucleus cross section problem is complex and has many moving pieces: nucleon-nucleon correlations, initial state distributions, final state interactions, etc. Knowing the axial form factor will stop one piece from moving and thus facilitates addressing the other moving pieces. Determination of nucleon transition form factors to the delta and other low mass baryon resonances would be similarly valuable.

A computation of the inclusive nucleon hadronic tensor in the transition region between quasi-elastic and Deep Inelastic Scattering (DIS), even with large, ~10% uncertainties, would be highly valued by the neutrino community. Efforts to compute nuclear matrix elements for light (A≤4) nuclei are not directly relevant to neutrino experiments unless they can be extended to carbon, oxygen, and argon, or provide insight into the validity of nuclear effective theories.

Precise lattice calculations relevant to studies of violations of fundamental symmetries (such as baryon and lepton number) and neutrinoless double beta decay will eventually become relevant. USQCD proposed a gradual approach to this problem, where precision does not play a leading role, as in flavor physics calculations.

The BSM effort provides valuable insights into strongly-coupled Quantum Field Theory (QFTs). Combined with a bottom-up approach to a more general Effective Field Theory (EFT) extension of the SM Lagrangian, the information on strongly-coupled QFT could validate the EFT approach and narrow the spectrum of possible UV completions. The investigation of both direct and indirect effects of new physics within the SM EFT framework has been growing in recent years and is now an integral and very active component of the LHC physics program. Several analyses have already appeared where SM EFT interactions are being constrained by LHC Run 2 data. More theoretical insights gathered through lattice studies could provide very valuable guidance in selecting which measurements could be more relevant.

Historically, USQCD has a successful track record in keeping a healthy balance between applying lattice methods to the calculation of fundamental quantities and physical observables that could have a direct input on experiments, and developing new methods and ideas in lattice gauge theory. We support maintaining this balance.

Recommendations

None.

Charge 3

Findings

Several reasons for the needs of an extended project were presented, including exploration of new methods and new science ideas, opportunities for junior researchers to start new projects (good examples are the recent progress in the computation of the hadronic light-by-light contribution to muon g-2, and a new project on Lambda_b decays, both initiated by junior USQCD members), fast turn-around time often not available at the Leadership Class Facilities (LCFs), and projects with workflows that are not well suited to high-capacity computing.

For 2019, 75% of the computing resources were obtained from INCITE projects on the LCFs while 25% were provided by the mid-scale cluster resources. The balance between both has shifted considerably over the years.

Comments

This project addresses a variety of physics topics with various computing needs. Those topics include projects that are at different stages of development. Some topics, such as generation of various gauge configurations, are both computationally intensive and mature enough to run on LCFs. The majority of projects, especially at the earlier stages of development, require more frequent “interactions” with physicists. Once the projects mature, some could be moved to LCFs. In addition, professional development of junior lattice theorists is easier with Institutional Clusters (ICs), where postdocs and junior faculty can be responsible for their own projects.

The collaboration made a convincing argument why the LCFs alone cannot fulfill their needs. Currently about 25% of all computing power available to the collaboration resides in the clusters. Much of the post-processing and data analysis requires a human in the loop which is not compatible with LCF workflows. Also, much of the work is exploratory in nature, that is, failure is to be expected and codes are very bleeding edge and thus this work is not suitable for execution on LCFs at all.

The cluster resources provide important opportunities for developing new approaches on small scales that, once ready, can then be ported to the larger production resources. The LCFs in general do not provide resources for such development work (only very minimal, like development queues or smaller machines with very short queues).

The balance between the usage of the LCFs (via INCITE allocations) vs the mid-scale cluster usage has shifted considerably over the years and will probably continue to do so, given the anticipated arrival of two exa-scale machines in the US in the time frame of 2021/2022. The collaboration made a convincing case that a moderately sized cluster environment will be important for their science in the future as well. While the size of this resource will not grow as quickly as the leadership computing resources, the balance between them seems to be adequate even in the coming years.

Recommendations

The mid-scale computing clusters should continue to compose a significant portion of the resources available for USQCD research.

Charge 4

Findings

The project documents, specifically the project execution plan, describe a management plan that incorporates both the existing dedicated clusters, and support for design, procurement, and operations within the institutional cluster model.

The USQCD collaboration, as part of its scientific operations, makes data from many of its computations on LCF machines available to all researchers inside and outside the collaborations. This contributes to the utility of these LCF computations, but also leads to significant data storage requirements.

Comments

The cost quoted by BNL and FNAL for additional storage of 4PB of data is approximately 10% of the five year LQCD-ext III budget.

The integration of the cluster resource management into the IC resources has started. Overall, the collaboration seems satisfied with how this new arrangement is working. Further monitoring and reporting about how this new arrangement works compared to the dedicated resources would be very valuable, also for other communities who might want to make a case for resources within the ICs at Fermilab and BNL.

Overall, the USQCD approach to mid-scale computing could serve as a model for other HEP communities as well.

Recommendations

The current project execution plan should be reviewed and revised after more experience is gained with the design, procurement, and operations phases of running on the institutional clusters.

The USQCD collaboration should review its additional storage needs and strategy and propose a specific plan at a future annual review.

The LQCD-ext III risk register should add the risk that IC technical solutions supported by FNAL or BNL become, in the future, not optimal in performance per unit cost for LQCD computing needs.

For DOE: We urge DOE to continue monitoring the usage and availability of the ICs at BNL and Fermilab to ensure that the USQCD collaboration obtains the full support and resources they need from the ICs.

For Fermilab/BNL: We urge both laboratories to keep ensuring that USQCD has sufficient input into the design decisions for the systems so that they will continue to provide the best possible computing environment for their projects.

For DOE: We recommend funding at the requested level for the proposed duration of the project.

Given the exploratory nature of some of the presented activities (a promising example is PDF calculations), and the fact that theory can be fast evolving in unexpected directions over a span of several years, the DOE should consider the possibility of increasing the USQCD budget for specific new research efforts during this grant period.

APPENDIX A

Charge Letter to the LQCD-ext III Team

Dr. Andreas Kronfeld
High Energy Theory Group
Wilson Hall
Fermi National Laboratory

P.O. Box 500
Batavia, IL 60510-0500

Dear Dr. Kronfeld,

The High Energy Physics (HEP) division of the Office of Science of the Department of Energy will conduct a review of your proposal for the extension of the LQCD research program to the next five year period, 2020-2024, at the Cambria Rockville Hotel at 1 Helen Heneghan way, Rockville, MD 20850 (https://www.cambriasuitesrockville.com/) on July 9-10, 2019. A review panel consisting of computational scientists and high energy theoretical and experimental physicists will evaluate both the scientific and computing plans that you have presented to us in your recent whitepapers.

This review will focus on the scientific justification of the proposal and its implementation on mid-scale Institutional Cluster hardware operating at Fermilab and Brookhaven National Laboratory.

The critical issues to be examined in the July 9-10 review include:

1. 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?
2. 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.
3. 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?
4. 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?

The review will begin with a closed executive session at 8:30AM on July 9, followed by presentations by you and your team that address the four charge points. The second half of the review will consist of additional executive sessions, preliminary report writing and a close-out where the review team will give you immediate feedback on your plans and presentations. You should work with John Kogut, the Federal Project Manager, and Bill Kilgore, Program Manager for Theoretical Physics, to generate an agenda for the review.

Each panel member will be asked to review those aspects of the review presentations that are within their scope of expertise. Each will write an individual report on his/her findings. These reports will be due at the DOE two weeks after completion of the review. John Kogut will accumulate the reports and produce a final summary report based on the information in the letters. That report will have recommendations for your consideration that you and your team should respond to in a timely fashion.

If you have additional questions, please contact John Kogut in HEP and/or Bill Kilgore.

We look forward to an informative and stimulating review.

Sincerely,

James Siegrist
Associate Director
Office of High Energy Physics

APPENDIX B

Reviewers for 2019 LQCD-ext II Science Review

1. **Maarten Golterman, Cal State SF (analytic LQCD and other)**
   maarten@stars.sfsu.edu or maarten@sfsu.edu

2. **Alexey Petrov, Wayne State (flavor physics)**
   apetrov@wayne.edu

3. **Katrin Heitmann, ANL (computational cosmology)**
   heitmann@anl.gov

4. **P. Huber (Va. Tech) (Neutrino theory)**
   pahuber@vt.edu

5. **Lee Roberts (Boston U.) (Muon g-2 experimentalist)**

roberts@bu.edu

6. **Kevin McFarland (T2K, neutrino exp)**
kevin@rochester.edu

7. **Laura Reina, Florida State (pQCD, collider pheno)**
reina@hep.fsu.edu

APPENDIX C

Review Agenda

July 9

08:30 Executive session (30 min)

09:00 Logistics and Introductions (10 min) – *John Kogut & Bill Boroski*

09:10 Overview of Scientific Program (15 min) – *Andreas Kronfeld*

09:25 USQCD Governance and NP-HEP Cooperation (30 min)– *Robert Edwards*

09:55 Science Talk 1: Quark and Lepton Flavor Physics (40 min) – *Ruth Van de Water*

10:35 Break (20 min)

10:55 Science Talk 2: Neutrino-Nucleus Scattering (40 min) – *Andreas Kronfeld*

11:35 Science Talk 3: Fundamental Symmetries (40 min) – *Zohreh Davoudi*

12:15 Working Lunch

1:15 Science Talk 4: Beyond the Standard Model (40 min) – *Ethan Neil*

1:55 LQCD-ext II: 2019 Accomplishments and Performance (40 min) - *Bill Boroski*

2:35 LQCD-III: Computational Requirements and Milestones (40 min) – *Andreas Kronfeld*

3:15 Break (20 min)

3:35 LQCD-ext III: Institutional Cluster Computing & Operations Model (40 min) – *Bill Boroski*

4:15 Summary (15 min) – *Andreas Kronfeld*

4:30 Executive Session (60 min)

5:30 Committee request for additional information – ***John Kogut / Proposal Leadership***

5:45 Adjourn

July 10

8:30 Response to committee questions and discussion (90 min)

10:00 Break (10 min)

10:10 Executive Session / Preliminary Report Writing

12:00 Working Lunch

1:00 Closeout

2:00 Adjourn