2010 USQCD Facilities Proposal: Dynamical Anisotropic-Clover Lattice Production for Hadronic and Nuclear Physics

Justin Foley, Colin Morningstar
Carnegie Mellon University
Kostas Orginos
College of William and Mary
Jozef Dudek, Robert Edwards, Bálint Joó,
David Richards, Christopher Thomas
Jefferson Lab
Stephen Wallace
University of Maryland
Huey-Wen Lin
University of Washington
International:
Nilmani Mathur
Tata Institute, India
Michael Peardon, Sinead Ryan
Trinity College, Dublin

(Dated: February 27, 2010)

We propose continuing the production of dynamical anisotropic Clover configurations suitable for studies in spectroscopy, hadronic structure, nuclear interactions, as well as searches for physics beyond the Standard Model. The current focus in production is the generation of 10000 trajectories of $32^3 \times 256$ lattices at a pion mass of 230 MeV. The time requested in this proposal is only for the lattice generation and not for subsequent valence calculations. We request using the time under the DOE INCITE 2010 award given to USQCD, and discretionary time when it becomes available in this next allocation year. Since the INCITE time is a multi-year award that terminates at the end of 2010, we ask for the second 6 months of time at the end of 2010 bringing the total time for 2010 to 40 million Cray hours. A new INCITE proposal should be submitted shortly, and we expect to request time for the period starting in 2011. Currently, a second stream at the target parameters are in production at ANL using INCITE time. We request using the INCITE time of **40 M** ORNL Cray core-hours and **20 M** ANL BG/P core-hours under the Class A proposal guidelines. This combined request will result in the production of 4000 trajectories. The online disk storage request is is **4 TB**. The tape storage request is **20 TB**. These two media requests are equivalent to **170 K** Jpsi-core hours.

I. PHYSICS GOALS

The central goal of this proposal is to provide the lattices necessary for calculations within nuclear and high-energy physics that are based on the QCD action. These calculations can be broadly characterized as under hadronic spectroscopy, hadronic structure, nuclear interactions, and physics beyond the Standard Model.

The determination of the excited state spectrum and structure of QCD is a major goal in hadronic physics. The measurement of the electromagnetic properties of the low-lying baryons, including the determination of nucleon resonances, is an HP 2012 milestone. The 12 GeV upgrade at Jefferson Laboratory has received "CD-3" and construction is underway. A flagship component of the upgrade is the new Hall D which has a major focus on spectroscopy. In particular, the GlueX Collaboration proposal seeks information on exotic meson states, where gluons might play an explicit structural role in the formation of the so-called hybrid mesons. In addition, there are new experimental programs in spectroscopy at BES-III (China) and GSI/Panda (Germany).

Hadronic structure calculations aim to build a three-dimensional picture of hadrons with the determination of the spin, flavor, and spatial distribution of quarks and gluons. The determination of hadronic form-factors is a major component of the 12 GeV upgrade at JLab.

Furthermore, the precise determination of the hadronic matrix elements of quark bilinears can provide important theoretical input for low energy experiments searching for physics beyond the standard model. The proposed precision measurement of neutron decay parameters at the Ultra Cold Neutron (UCN) source at LANL is such an experiment.

The calculation of nuclear processes directly from QCD can provide a new understanding of stellar evolution. The calculations of hyperon-hyperon scattering as well as properties of three and four nucleon systems are a major goal within nuclear physics.

Given the current intense experimental efforts in these areas, the need for generating lattices suitable for such calculations is clear. In this proposal, we continue the generation of such lattices that were begun initially to enable the comprehensive study of the spectrum of both baryons and mesons, including exotics, in full QCD.

The relevance of this work to DOE's mission is illustrated in the Executive Summary of the SciDAC-2 proposal, namely that this work intends to "calculate the masses of strongly interacting particles and obtain a quantitative understanding of their internal structure". In addition, the configurations that are generated are of use to the wider QCD community. Thus, the time requested, **40** million ORNL Cray core-hours and **20** million ANL BG/P core-hours, is justified under the guidelines of a *Class A* proposal.

This time will allow for the generation of 4000 trajectories of $32^3 \times 256$ lattices at a pion mass of 230MeV. However, based on experience at larger pion masses, we anticipate needing roughly 10000 trajectories. So, additional time, should it be available at ORNL, ANL, or other facilities, will also be needed.

This proposal focuses on the generation of the lattices which are described in Refs [1, 2]. There are additional critical components to the overall effort.

The project by *D. Richards* focuses on the determination of the baryon spectrum and meson spectrum following techniques developed in Refs. [3–6], along with the calculation of the exotic π_1 photo-coupling amplitude which is relevant for JLab's Hall D program. This project will use these $N_f = 2 + 1$ configurations for a calculation of the transition form-factors in the hybrid channel, as well as for higher spin states and excited states. This latter part of the project is a continuation of the charmonium spectrum project and transition work published in Ref. [7–9] where the first excited state transition form-factors within charmonium were computed. The baryon spectrum calculation, a continuation of the work in $N_f = 0$ (Ref. [10]) and $N_f = 2$ (Ref. [11]), will use the multiple volumes available to determine the light quark and strange quark octet masses. Both the baryon and meson spectrum determinations will use new techniques for multi-hadron operators.

A significant advance was made recently with the development of the "distillation" algorithm for hadron operator and correlator construction [12]. A recent letter [9] demonstrates the efficacy (and power) of the method by extracting high lying exotic meson states as well as even spin-4 states. Work is on-going using lighter masses and larger volumes. The use of anisotropic lattices (fine temporal lattice spacing), extended (non-local) operators, and variational technologies are crucial in these efforts. The new technique is well suited for studies of decays using multi-hadron operators.

The FormFactor proposal (*K. Orginos*) is using these lattices for the computation of ground state and excited state transition baryon form-factors relevant to the experimental programs at JLab and Mainz. In addition, this proposal also intends to compute, at $Q^2 = 0$, the 5 possible quark bilinear combinations in the nucleon form-factor which will provide information of neutron decay parameters relevant for the BSM searches.

In addition, the NPLQCD projects (*M. Savage*) intends to use these lattices for their multi-hadron investigations. Their recent very high statistics work (Ref. [13–15]) shows clearly how a fine temporal resolution is important for extracting signals in three-nucleon systems. The NPLQCD proposal anticipates using the lattices that are the subject of this proposal, and are under generation now.

The EMC project (*B. Tiburzi*) also intends to use these lattices for calculations of hadronic polarizabilities. Also, the DISCO project (*J. Osborn*) is using these lattices for the computation of the disconnected insertions within baryon form-factors.

II. COMPUTATIONAL STRATEGY

A. Actions and Parameters

The use of anisotropic lattices is essential to this project. The dense excited state spectrum of baryons extends to around 3 GeV. Using conventional isotropic lattices, with spacings around 2 GeV^{-1}, one can expect significant lattice discretization errors. Using a highly anisotropic lattice, e.g. with temporal lattice spacing of 6 GeV^{-1}, the discretization errors are reduced and better signals are provided for excited states. A well defined positive-definite transfer matrix is crucial to the determination of the excited state hadron spectrum (as well as the glueball spectrum). For this reason, a chiral fermion action, such as the overlap action, is *not* preferred since the lack of positivity results in unphysical oscillations in correlators. Finally, it is important to use fermion and gauge actions with small scaling violations.

For this project, we have chosen to use the $N_f = 2 + 1$ anisotropic Clover fermion action [16, 17] with *Stout-link* smearing [18] of the *spatial* gauge fields in the Clover action. The gauge fields entering the fermion action are not smeared in the time direction, thus preserving the transfer matrix. In addition, we use a tree-level tadpole-improved Symanzik gauge action with no 1×2 rectangle in the time direction. This gauge action was used for the determination of the glueball spectrum [19]. For our project, only two iterations of Stout smearing are used with a weight of the staples of $\rho = 0.14$ which is below some critical (but classical) upper bound where the glue fields get inordinately amplified [21]. The principal reason for using a Stout-smeared fermion action is to make the fermion matrix more stable at small quark masses. Also, different lattice spacing scales are involved. The time direction is very fine, and after the spatial Stout smearing we find the tadpole factors for the temporal and spatial links to be very nearly 1. This means that one might expect the tadpole corrected

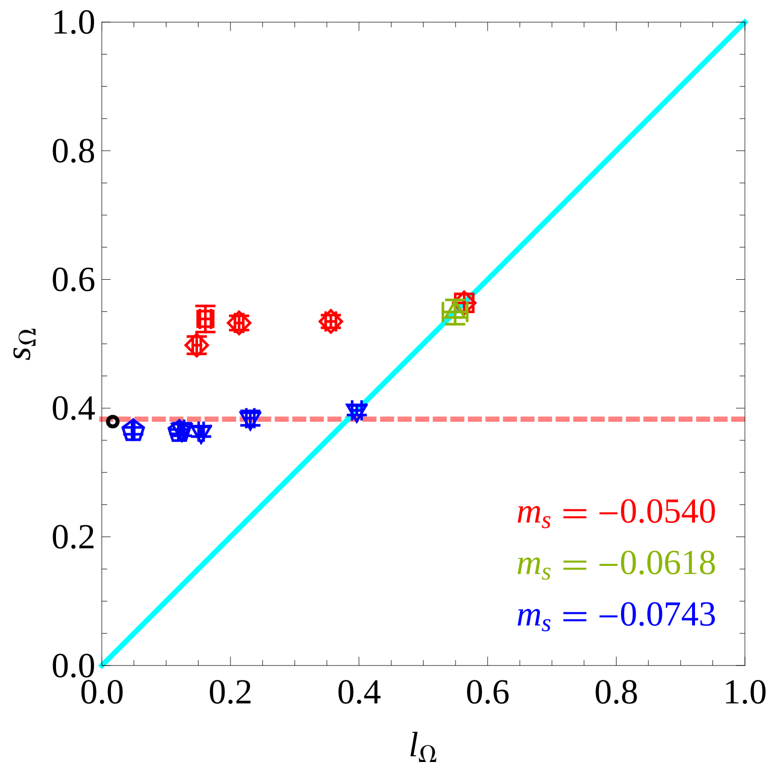

FIG. 1: The "Newport News" plot. The horizontal and vertical axes are ratios of QCD stable hadron masses whose leading order χPT dependence is on the light quark (m_l) and strange quark masses (m_s). Shown are 3 different strange quark mass simulations. The physical limit is the black circle. The cyan diagonal line is the $N_f = 3$ line. The normalizations of the axes is chosen to give 1 in the infinite quark mass limit. The pink dashed line is a line of slope 0 anchored by the physical point. The blue symbols are the current $N_f = 2 + 1$ simulations. The pentagons, the lightest two pion masses (blue), are $24^3 \times 128$ runs while the others are $16^3 \times 128$. Results for $20^3 \times 128$ lattices are not shown. We see that the current simulations lie close to this line indicating the strange quark mass is close to the physical value (at this lattice spacing). The current work is generating $32^3 \times 256$ lattices at the lowest pion mass of 230MeV.

tree-level Clover coefficients might be very near their non-perturbative estimates, and this is in fact what we observe. Namely, while we did not attempt to non-perturbatively tune the Clover coefficients; nevertheless, they are in fact compatible with being non-perturbatively tuned.

We have employed a new method for the tuning of the anisotropies based on the Schrödinger functional method at zero renormalized quark mass [1]. The tuned anisotropies ensure that the renormalized gauge and fermion anisotropies are equal to the desired value $\xi = 3.5$. Very importantly, we find no statistically significant dependence of the bare anisotropies on the mass. The fact that we see no significant mass dependence allows us to fix the anisotropies for all masses, thus greatly simplifying our efforts.

The combination of the anisotropy tuning, controlling the lattice spacing, and tuning the strange quark mass is rather involved; however, in the end a simple method was developed to separately determine the strange quark mass and lattice spacing. A final version of the

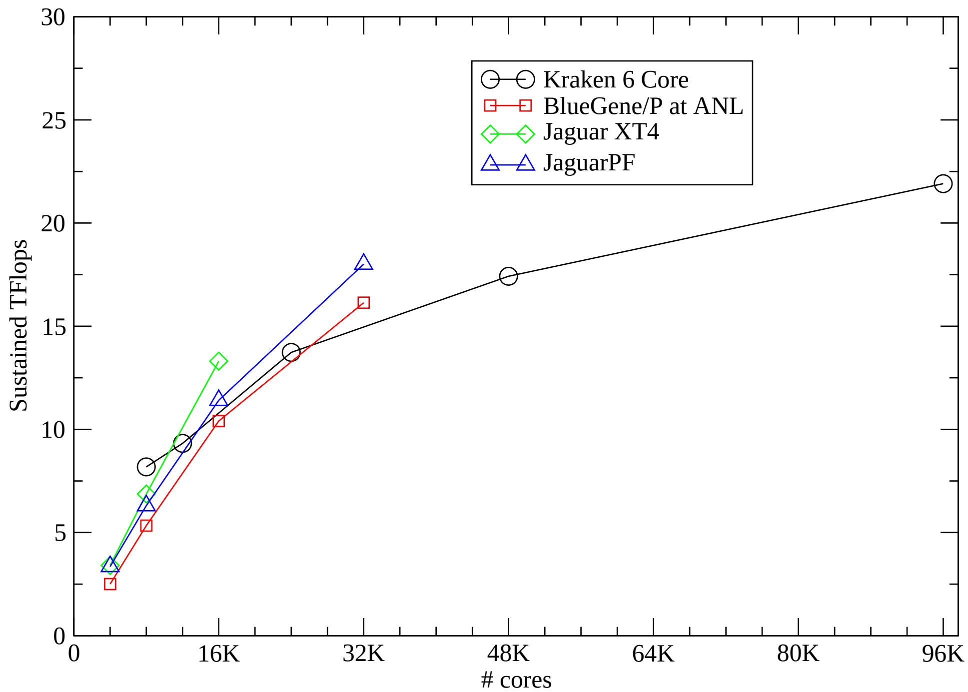

FIG. 2: Hard scaling of the performance of the Mixed Precision BiCGStab solver on a variety of resources (including Kraken Hex-core). Measurements were taken prior to the Jaguar Hex-core upgrade, so Jaguar at this point was a dual socket quad core XT5. Jaguar XT4 was a single socket quad-core XT4. Error bars are fairly large for the performances on the Cray systems, with significant downward fluctuations. Performances on the BG/P system are fairly stable.

results have appeared in Ref. [2]. The idea is to parameterize observables with dimensionless coordinates involving the light and strange quark mass

$$l_X = \frac{9m_\pi^2}{4m_X^2}, \qquad s_X = \frac{9(2m_K^2 - m_\pi^2)}{4m_X^2} \tag{1}$$

where m_X is some reference scale, preferably a stable particle under the strong interactions. We chose the Ω baryon. The normalization is chosen so that in the infinite quark mass limit the coordinates l_Ω and s_Ω tend to 1.

The strange quark tuning strategy is then to consider the $N_f = 3$ case, tune the bare quark mass till s_Ω takes its physical value. One can approach the physical limit in this plane of coordinates; however, it is simpler to fix the strange quark mass and vary only the light quark mass. Choosing ratios cancels the explicit lattice spacing and removes the need for a scale determination. The results for our lattices are shown in Figure 1. We see the strange quark mass choice that projects onto the physical value of s_Ω is a fairly reasonable choice as the light quark mass has decreased.

We note that the recent BMW calculation in Ref. [23] also used these coordinates but with the reference mass the Ξ baryon – another equally fine choice, but one that doesn't allow as much leverage in determining the strange quark mass.

Finally, we find that the integrated autocorrelation times for our datasets is moderate. For the pion correlator, $\tau_{int} \approx 25(10)$ for the lightest pion mass dataset.

B. Production Plans

Our multi-year plan is to generate a large number of configurations, at multiple volumes and masses, to enable high statistics measurements – all at initially one lattice spacing. We view determining resonance states from their decay or scattering states as the principal objective of the spectrum component of this project. In general, determining the properties of hadronic and nuclear physics observables at a light pion mass is a primary objective. The quark mass dependence is an important issue we need to address. While we will attempt chiral extrapolations, the goal is to get to the physical limit where we believe we can state with some confidence that we have determined the hadronic spectrum, albeit at some fixed lattice cutoff. In the interim, calculations at the lowest available pion masses are essential.

The current inventory of datasets is shown in Table I. We have finished high statistics $16^3 \times 128$, $20^3 \times 128$, and $24^3 \times 128$ runs that enable first determinations of the hadron spectrum with multiple volumes down to pion masses of 383MeV. A limited statistics $32^3 \times 256$ ensemble at this 383MeV pion mass is also available. In addition, a high statistics ensemble on $24^3 \times 128$ at 230MeV pion masses is also completed.

The focus of this proposal is the continuation of current runs for the generation of a larger volume $32^3 \times 256$ ensemble at the lowest pion mass of 230MeV. The first half-year part of the 2010 INCITE allocation at ORNL is being used. In addition, some time at ANL under INCITE and discretionary time are being used for the generation of a second stream for this dataset.

The issue of discretization effects is, of course, an important one. To attempt to have some measure of cutoff effects, we will at some point move to a finer lattice, but at some moderate pion mass that is close to one of our existing values, and compare the spectrum we determine at two (nearly) equivalent pion masses, but different cutoffs. In the language of the coordinates used above, we determine the spectrum at different gauge couplings β, but at a fixed l_Ω and s_Ω.

At some point, once we feel we have adequately resolved the spectrum at the initial coupling of $\beta = 1.5$, we would move to a finer lattice. However, as far as cost, it might be more effective to move to a fine isotropic lattice.

Last year, the SPC asked for clarifications about when a switch to a finer lattice spacing might be planned. We note that during the 2009 year of running at ORNL, we found that to maintain reversibility in the algorithm, we needed to switch to double precision. Significant code changes were needed to maintain reasonable computational cost. In the near term, we are focusing available resources on generating a large statistics $32^3 \times 256$ ensembles at the lowest pion masses. Should some time become available, we could puruse more aggressively the characterization of discretization effects outlined above.

III. SOFTWARE

Our dynamical anisotropic Clover gauge generation has been using the *Chroma* HMC code with RHMC for the strange quark, and Hasenbüsch style mass preconditioning for the two light quarks, along with multi-timescale integration. Currently, the Clover action is four-dimensionally even-odd preconditioned.

To amortize the cost of overhead from communications, we have implemented a threaded version of the Wilson-Dirac operator, as well as the rest of QDP++ and Chroma, using the light-weight threading package (called QMT) developed under SciDAC. In this model,

m_π(MeV)	m_l	m_s	volume	# traj.
1563	-0.0540	-0.0540	$12^3 \times 96$	4670
1120	-0.0699	-0.0540	$12^3 \times 96$	5550
783	-0.0794	-0.0540	$12^3 \times 96$	4770
635	-0.0826	-0.0540	$12^3 \times 96$	3885
635	-0.0826	-0.0540	$16^3 \times 96$	1800
1305	-0.0618	-0.0618	$12^3 \times 96$	2000
833	-0.0743	-0.0743	$12^3 \times 96$	10575
			$16^3 \times 128$	11735
			$20^3 \times 128$	11090
560	-0.0808	-0.0743	$16^3 \times 128$	11010
			$20^3 \times 128$	10245
448	-0.0830	-0.0743	$16^3 \times 128$	13470
			$20^3 \times 128$	10280
383	-0.0840	-0.0743	$16^3 \times 128$	11005
			$20^3 \times 128$	13090
			$24^3 \times 128$	12985
			$32^3 \times 256$	4000
230	-0.0860	-0.0743	$24^3 \times 128$	12760
			$32^3 \times 256$	3774

TABLE I: Inventory as of Feb. 27, 2010. Total number of trajectories (including those from thermalization) of $N_f = 2+1$ anisotropic Clover runs at $\beta = 1.5$, $a_s = 0.1227(8)$fm, $\xi = 3.5$, $a_t^{-1} = 5.62(4)$GeV. The smaller strange quark mass $m_s = -0.0743$ is the final target value. Configurations have been saved every 5th trajectory. The $32^3 \times 256$ runs are archived every 2nd trajectory. Typically, 1000 traj. are dropped for thermalization. For the lightest mass $32^3 \times 256$ dataset, 1480 traj. in total should dropped for thermalization from the two streams.

only one core on a node is responsible for communications, while the linear algebra part of the work is threaded over all the cores. This technique coalesces messages into larger ones compared to a non-hybrid version. The effect is to significantly improve scalability on the Cray systems.

While good scaling of the dslash operator is essential, ultimately it is the performance of the solvers that matters. During the last year we have made several improvements in our solver technology. It was found that to to maintain reversibility in the molecular dynamics part of our Hybrid Monte Carlo algorithm, with such light quark masses and large lattices that single precision solutions were no longer adequate and one had to move to double precision. To counter the factor of 2 slowdown from this move and to improve our scaling to larger partitions, we have combined the multiple precision approach of [24] with the Improved BiCGStab approach of [25]. The Improved BiCGStab algorithm reduces the 4 separate global reduction stages in conventional BiCGStab algorithm to just 1 global reduction and the multiple precision technique allows one to run the calculation primarily in single precision, and yet still achieve double precision accuracy.

The ORNL run-time environment forces us to run at very large core counts - a minimum of 9100 cores on the old XT4 system and 135000 cores on the new XT5 system to run at a "normal" priority. The problem sizes per core are quite tiny, and we are latency bound. We are successfully using the threading implementation of Chroma's HMC to allow scaling to very large core counts. While important for the Cray machines and clusters, this threading technique is not necessary on BG/Ps. Thus, we could take advantage of those systems using the conventional QMP model.

In figure 2 we show the hard scaling of our mixed precision BiCGStab solver on a variety of resources, including Kraken after its hex-core update. In this plot, we have hard scaled our problem size up to 96K cores on Kraken. An interesting feature of the plot is that at 'an order of magnitude level' it shows similar performances on the XT4, XT5 and BlueGene hardware up to about 24–32K cores. We would envisage running on Jaguar at the 24K core mark, as one can see noticeable strong scaling degradation at 48K and 96K cores.

However, we note that currently the old XT4 system at ORNL is operational, and we have been using that system at 16K cores. There is less user demand for this system, scaling is better, it is fairly stable and actually has a higher throughput and is more cost effective than the newer XT5 system for our size of jobs. We will continue to use the XT4 while it is operational.

We intend to continue improving the performance of our application through further development throughout the allocation period. We are actively investigating using the Force Gradient integrator [26]. This method can potentially decrease the number of integration steps required for a given acceptance rate, hence decrease the number of calls to the inverters.

IV. REQUIRED RESOURCES

At this point in our software development and in our production running, the cost of the gauge generation at the lightest $32^3 \times 256$ ensemble at a 230MeV pion mass is understood, up to fluctuations in timings and machine instabilities observed at ORNL. The time requirments on both the ORNL Crays and the ANL BG/P are basically the same when run at comparable number of cores (16k to 32k). For 1000 trajectories, 15 M-core-hours are needed. We currently have 3774 trajectories in total, resulting in 2250 trajectories available after thermalization. The request of 40M (ORNL) and 20M (ANL) core-hours will generate another 4000 trajectories bringing the total number of trajectories (after thermalization) to 6250. We need another 3800 trajectories to reach the target of 10000 trajectories.

We anticipate having a physical limit calculation to finish the $\beta = 1.5$ datasets.

V. USE FOR ADDITIONAL RESOURCES

As noted before, currently we are using INCITE resources at ORNL and ANL. As more resources become available, we intend to finish the $32^3 \times 256$ run at $m_\pi = 230$MeV. The remaining 3800 trajectories require another 57 M-core-hours using current Cray and BG/P machines.

We would also like to take an initial start at a physical limit calculation. While the time needed is large, valuable experience would be gained by making measurements on a limited statistics dataset. However, expected algorithmic improvements from the force-gradient integrator should lower the cost of such a calculation since the method becomes more effective at lower pion masses.

volume	Existing TB	Proposed TB	Total TB
$12^3 \times 96$	0.4		0.4
$16^3 \times 128$	1.4		1.4
$20^3 \times 128$	2.2		2.2
$24^3 \times 128$	2.5		2.5
$32^3 \times 256$	8.8	4.6	13.4
total			20

TABLE II: Storage requirements for gauge configurations.

VI. DATA SHARING

The initial $N_f = 2$ anisotropic Wilson lattices have been archived on the ILDG in single precision. They also have been stored at JLab on tape. The Clover lattices are already on tape at JLab, and are available to the whole collaboration as they are generated. However, given that the current ILDG implementation at FNAL makes all configurations public, the Clover lattices will not be placed there immediately. Interested parties should request access to the $20^3 \times 128$ lattices and the $m_\pi = 230$Mev $24^3 \times 128$ lattices which were generated using NSF Teragrid resources.

VII. DATA STORAGE

media	Storage TB	Jpsi-equiv
disk	4	108K
tape	3+20	62K
total		170K

TABLE III: Storage requirements turned into Jpsi-equivalent node hours

The $N_f = 2$ Wilson lattices have been stored on tape at JLab in single precision, taking roughly 3TB of tape to hold.

The storage requirements for the $N_f = 2 + 1$ configurations is shown in Table II. It is expected that **23 TB** will be needed for off-line storage for all the configurations proposed and the ones that currently exist. At any one time, one mass set would be needed for analysis. The worst case expected in this allocation year is for the $32^3 \times 256$. Given the number that can be produced in this year, a working space of **4 TB** is expected.

Under the storage requirements of USQCD, the online disk space and off-line tape storage have 6n-equivalent costs, and are summarized below in Table III. The 6n-equivalent cost is **170K** Jpsi-node hours.

VIII. EXCLUSIVITY

The follow-on computation of the excited resonance spectrum for both baryons and mesons in the light-quark sector are considered exclusive elements of this proposal, reserved for the Hadron Spectrum Collaboration. Otherwise, there are additional proposals requesting specific calculations.

---

[1] R. G. Edwards, B. Joo, and H.-W. Lin, Phys. Rev. **D78**, 054501 (2008), 0803.3960.

[2] H.-W. Lin et al. (Hadron Spectrum), Phys. Rev. **D79**, 034502 (2009), 0810.3588.

[3] S. Basak et al., Phys. Rev. **D72**, 094506 (2005), hep-lat/0506029.

[4] S. Basak et al. (Lattice Hadron Physics (LHPC)), Phys. Rev. **D72**, 074501 (2005), hep-lat/0508018.

[5] A. C. Lichtl (2006), hep-lat/0609019.

[6] J. J. Dudek, R. G. Edwards, M. J. Peardon, D. G. Richards, and C. E. Thomas, Phys. Rev. Lett. **103**, 262001 (2009), 0909.0200.

[7] J. J. Dudek, R. G. Edwards, and D. G. Richards, Phys. Rev. **D73**, 074507 (2006), hep-ph/0601137.

[8] J. J. Dudek, R. G. Edwards, N. Mathur, and D. G. Richards, Phys. Rev. **D77**, 034501 (2008), arXiv:0707.4162 [hep-lat].

[9] J. Dudek, R. Edwards, and C. Thomas (2009), 0902.2241.

[10] S. Basak et al., Phys. Rev. **D76**, 074504 (2007), arXiv:0709.0008 [hep-lat].

[11] J. M. Bulava et al., Phys. Rev. **D79**, 034505 (2009), 0901.0027.

[12] M. Peardon et al. (Hadron Spectrum), Phys. Rev. **D80**, 054506 (2009), 0905.2160.

[13] S. R. Beane et al. (2009), 0903.2990.

[14] S. R. Beane et al., Phys. Rev. **D80**, 074501 (2009), 0905.0466.

[15] S. R. Beane et al. (2009), 0912.4243.

[16] M. Okamoto et al. (CP-PACS), Phys. Rev. **D65**, 094508 (2002), hep-lat/0112020.

[17] P. Chen, Phys. Rev. **D64**, 034509 (2001), hep-lat/0006019.

[18] C. Morningstar and M. J. Peardon, Phys. Rev. **D69**, 054501 (2004), hep-lat/0311018.

[19] C. J. Morningstar and M. J. Peardon, Phys. Rev. **D60**, 034509 (1999), hep-lat/9901004.

[20] R. Morrin, A. O. Cais, M. Peardon, S. M. Ryan, and J.-I. Skullerud, Phys. Rev. **D74**, 014505 (2006), hep-lat/0604021.

[21] S. Capitani, S. Durr, and C. Hoelbling, JHEP **11**, 028 (2006), hep-lat/0607006.

[22] T. Umeda et al. (CP-PACS), Phys. Rev. **D68**, 034503 (2003), hep-lat/0302024.

[23] S. Durr et al., Science **322**, 1224 (2008).

[24] M. A. Clark, R. Babich, K. Barros, R. C. Brower, and C. Rebbi (2009), 0911.3191.

[25] in *ICA3PP '02: Proceedings of the Fifth International Conference on Algorithms and Architectures for Parallel Processing* (IEEE Computer Society, Washington, DC, USA, 2002), p. 324, ISBN 0-7695-1512-6.

[26] A. D. Kennedy, M. A. Clark, and P. J. Silva (2009), 0910.2950.