2011 USQCD Facilities Proposal: Dynamical Anisotropic-Clover Lattice Production for Hadronic and Nuclear Physics

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
International:
Nilmani Mathur
Tata Institute, India
Michael Peardon, Sinead Ryan
Trinity College, Dublin
(Dated: March 21, 2011)

We propose continuing the production of dynamical anisotropic Clover configurations suitable for studies in spectroscopy, hadronic structure and nuclear interactions. The current focus of production is the start of a new $40^3 \times 256$ dataset 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. This project falls under Class A proposal guidelines as the configurations will be used for multiple projects meeting USQCD scientific objectives. We request using the DOE INCITE 2011 time of **30 M** ORNL Cray core-hours as well as the first half year of INCITE 2012 ORNL time. In addition, we request using discretionary time when it becomes available in this next allocation year at both ANL and ORNL. The first half of year of 2011 ORNL time has already been allocated. Assuming the ORNL allocation to be the same in 2012, we expect the first half year to amount to 15 M ORNL Cray hours. The combined request for time starting July 1, 2011 to June 30, 2012 is **30 M** ORNL Cray hours. The ORNL time will result in the production of 1500 trajectories. The online disk storage request is is **6 TB**. The tape storage request is **23 TB**. These two media requests are equivalent to **223 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 and nuclear interactions.

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. Construction for the 12 GeV upgrade at Jefferson Laboratory 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.

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, **30** million ORNL Cray core-hours and discretionary time at both ANL BG/P and ORNL, is justified under the guidelines of a *Class A* proposal.

This time will allow for the thermalization and generation of 1500 trajectories of $40^3 \times 256$ lattices at a pion mass of 230MeV. However, based on experience with a smaller lattice size at this pion mass, we anticipate needing roughly 7000 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. [3–9], 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. [10–12] where excited state transition form-factors within charmonium were computed for the first time within lattice QCD. The baryon spectrum calculation, a continuation of the work in $N_f = 0$ (Ref. [13]), $N_f = 2$ (Ref. [14]), and $N_f = 2+1$ (Ref. [15]) will use the multiple volumes available to determine the light quark and strange quark baryon masses. Both the baryon and meson spectrum determinations will use new techniques for multi-hadron operators.

A significant advance was made with the development of the "distillation" algorithm for hadron operator and correlator construction [16]. A recent letter [6] and paper [7] demonstrates the efficacy (and power) of the method by extracting high lying exotic meson states as well as even spin-4 states. This excited state technology has also been extended to the determination of phase shifts [8] - a crucial ingredients in the determination of hadron resonance states.

The distillation technique provides a convenient framework for the computation of all-to-all propagators necessary for the calculation of correlators with annihilation diagrams. A recent calculation [9] has, for the first time, computed the highly excited spectrum of (single-particle) isoscalar mesons, including states of high spin. Hidden-flavor mixing angles are extracted and compared with existing phenomenological estimates. A notable feature of this work is the calculation of light *exotic* J^{PC} isoscalar meson states appearing at a mass scale comparable to their corresponding isovector partners. Thus, multiple exotic states maybe accessible to JLab's Hall D project.

Work is on-going using lighter masses and larger volumes [17, 18]. 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.

We note that the Hadron Spectrum calculation of excited spectra, as well as the recent successes with GPUs, featured as *Selected FY10 Accomplishments in Nuclear Theory in the FY12 Congressional Budget Request [19].*

The 2010 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.

The NPLQCD project (*M. Savage*) intends to use these lattices for their multi-hadron investigations. Their recent very high statistics work (Ref. [20–23]) shows clearly how a fine temporal resolution is important for extracting signals in three-nucleon systems. The recent result in Ref. [23] demonstrates a bound state in the $\Lambda\Lambda$ channel (the H-Dibaryon). The NPLQCD proposal anticipates using the lattices that are the subject of this proposal, and are under generation now.

The EMC project (*B. Tiburzi*) 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 gauge and fermion actions that are used in this project have been established for some time. Here, we recount some of the basic ideas, and refer to Refs. [1, 2] for more details.

The use of anisotropic lattices has been, and continues to be essential to this project. A project goal is to resolve the dense excited state spectrum of baryons up 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. For this reason, a chiral fermion action, such as the overlap action, is *not* preferred since the lack of positivity results in unphysical

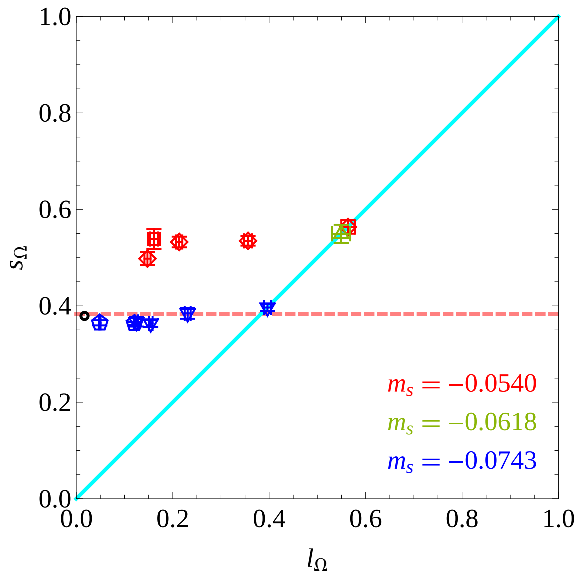

FIG. 1: The "Newport News" plot. The horizontal and vertical axes are ratios of QCD stable hadron masses whose leading order χPT dependence is on the light quark (m_l) and strange quark masses (m_s). Shown are 3 different strange quark mass simulations. The physical limit is the black circle. The cyan 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 $40^3 \times 256$ lattices at the lowest pion mass of 230MeV.

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 described in detail in Refs. [1, 2]. The spatial gauge fields in the Clover action are Stout-link smeared. 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 [24]. 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 [25]. 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 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 based [1].

As we are working with 2 light and 1 strange quark mass, we must disentangle the deter-

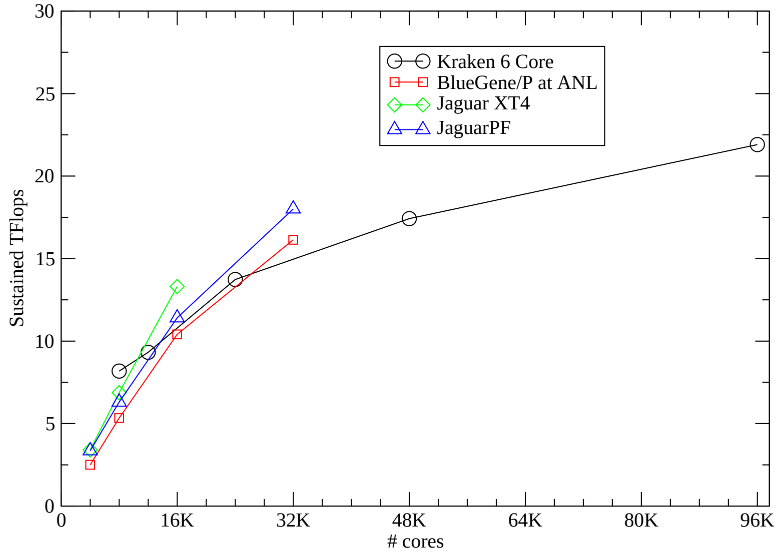

FIG. 2: Hard scaling of the performance of the Mixed Precision BiCGStab solver on a variety of resources (including Kraken Hex-core). The lattice size is $32^2 \times 256$. Measurements were taken prior to the Jaguar Hex-core upgrade, so Jaguar at this point was a dual socket quad core XT5. Jaguar XT4 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.

mination of the strange quark mass from a determination of the lattice spacing. The idea we adopted [2] involves parametrizing 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.

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.

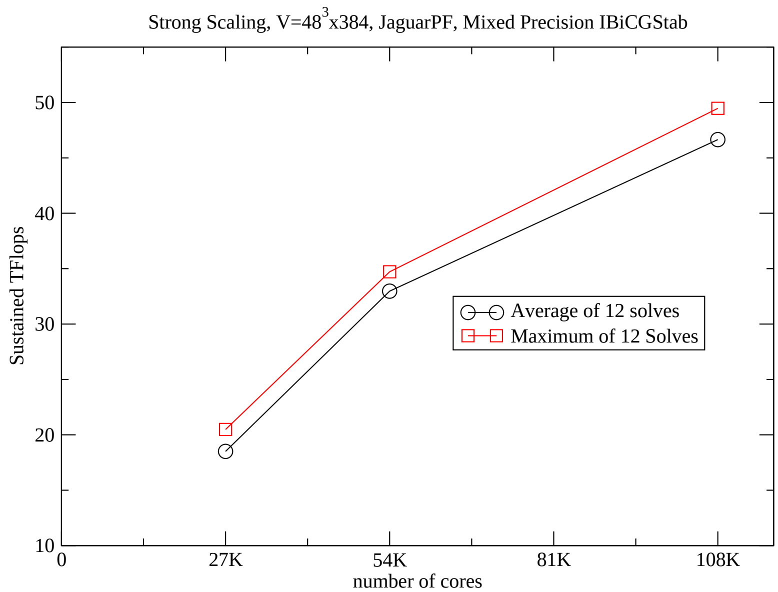

FIG. 3: This graph shows the hard scaling of the performance of the Mixed Precision BiCGStab solver on JaguarPF using a $48^2 \times 384$ lattice. core. This lattice size would be used for a physical limit calculation were sufficient computing time available. The proposed lattice size is smaller - namely, $40^3 \times 256$. Current production is using 19K cores.

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. The 2010 INCITE allocation at ORNL, plus extra time found through other resources, have enabled us to reach (very close) to an initial target of 7000 usable trajectories on the $32^3 \times 256$ dataset at 230MeV. In addition, we have also extended the 383MeV dataset at $32^3 \times 256$. Based upon our experience with valence measurements on these lattices and on smaller ($24^3 \times 128$) lattices which have 10k trajectories, we request extending both $32^3 \times 256$ datasets further to around 10k trajectories. Priority is given to the 230MeV ensemble. We request running under discretionary time at ANL for such a purpose.

The focus of this proposal is the start of a new dataset, a $40^3 \times 256$ ensemble at the lowest pion mass of 230MeV. The first half-year part of the 2011 INCITE allocation at ORNL is being used. We currently are using 19k core partitions as this size gives a reasonable trade-off on efficiency plus gives sufficient priority in the queues. At this level, we estimate the 30M core hour allocation plus 15M core hours from the first half of 2012 will delivery 1500 trajectories, plus some for thermalization. Clearly, a multi-year program is needed for this dataset (such was the case for the previous $32^3 \times 256$ dataset). We expect to utilize the NSF Blue Waters machine later this year, but in any case, we expect this $40^3 \times 256$ dataset to require additional resources to reach adequate statistics.

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
			$32^3 \times 256$	4116
230	-0.0860	-0.0743	$24^3 \times 128$	12760
			$32^3 \times 256$	8108
			$40^3 \times 256$	warming

TABLE I: Inventory as of March 10, 2011. 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 (blue), 1120 traj. in total should be dropped for thermalization. The proposed dataset $40^3 \times 256$ (red) is thermalizing at the time of writing.

A few times, the SPC has asked for clarifications about when a switch to a finer lattice spacing might be planned. We answer that the goal, first and foremost, is to do a decent job at one lattice spacing before moving on. This entails using high statistics - order 10k trajectories at multiple pion masses and volumes - for valence measurements. We have detailed observations that (at least) rotational breaking effects from lattice discretizations are negligible in our measurements. However, high statistics are needed to reduce gauge noise in isoscalar meson determinations, as well as to reduce systematic uncertainties in determining phase shifts and scattering lengths. We feel that moving to a second lattice spacing before completing the proposed $40^3 \times 256$ dataset would be premature.

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, 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 essential for the Cray machines, this threading technique is not necessary on BG/Ps. We could take advantage of those systems using the conventional QMP model.

While good scaling of the dslash operator is essential, ultimately it is the performance of the solvers that matters. Recent improvements in our solver technology has significantly improved performance. Firstly, 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 [26] with the Improved BiCGStab approach of [27]. 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.

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 are currently using 19k core partitions on JaguarPF (the only system now available at ORNL) for our $40^3 \times 256$ ensemble to achieve a reasonable performance as well as queue priority. We note that hard scaling is considerably improved on the new Cray XE6 at NERSC, but since NERSC is not part of this allocation request, we will not report on it here.

As a point of reference, we show in Figure 3 the hard scaling of the BiCGStab solver on a $48^3 \times 256$ lattice. This lattice size is what we use for a physical limit calculation if sufficient computing time were available (which there is not).

We envisage a significant software effort in the near future that will be needed to efficiently utilize new architectures coming online. Firstly, there is the NSF Blue Waters system in which we have been actively involved, as well as the BG/Q system at ANL. A high priority concern is the efficient support of heterogeneous systems. Namely, the successor to the current JaguarPF system at ORNL will be a Cray based machine using Nvidia GPUs. Of course, there is also the JLab and (soon) FNAL gpu clusters.

A new SciDAC effort at JLab is aimed at reimplementing QDP++ with a back end library system that initially supports GPUs and also CPUs in a new way. By reimplementing QDP++ while keeping the same interface, we can move large code blocks and data onto

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
$32^3 \times 256$	6.0		6.0
$40^3 \times 256$	0.0	7.5	7.5
total	12.5		20

TABLE II: Storage requirements for gauge configurations.

the GPUs thus enabling improved performance for HMC as well as correlator contractions which are becoming an increasingly time consuming part in Nuclear Physics calculations. Further performance improvements come from incorporating Level 3 libraries, and to this end we have made an initial port of the Chroma HMC system to GPUs using the QUDA level-3 library. The goal is to have such a hybrid software system in place so as to exploit the recently announced Titan system at ORNL, as well as potentially using the JLab GPU cluster for gauge generation.

IV. REQUIRED RESOURCES

Based upon current performance on JaguarPF at ORNL, we conservatively estimate needing 45M Cray core hours to produce 1500 trajectories on the $40^3 \times 256$ ensemble at a 230MeV pion mass. We will use 19K cores so as to maintain a reasonable performance as well as priority in the queues.

V. USE FOR ADDITIONAL RESOURCES

Once equilibrated, we will split off a second stream of the gauge generation, and would like to move it to ANL. However, based upon previous experience, we suspect that running a $40^3 \times 256$ dataset under discretionary time at ANL will be problematic. Instead, our request is to extend the previous dataset ($32^3 \times 256$) at ANL.

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.

media	Storage TB	Jpsi-equiv
disk	6	161K
tape	3+20	62K
total		223K

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

VII. DATA STORAGE

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. A working space of **6 TB** is expected.

Under the storage requirements of USQCD, the online disk space and off-line tape storage have Jpsi-equivalent costs, and are summarized below in Table III. The Jpsi-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, M. J. Peardon, D. G. Richards, and C. E. Thomas, Phys. Rev. **D82**, 034508 (2010), 1004.4930.

[8] J. J. Dudek, R. G. Edwards, M. J. Peardon, D. G. Richards, and C. E. Thomas (2010), 1011.6352.

[9] J. J. Dudek et al. (2011), 1102.4299.

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

[11] J. J. Dudek, R. G. Edwards, N. Mathur, and D. G. Richards, Phys. Rev. **D77**, 034501 (2008), 0707.4162.

[12] J. J. Dudek, R. Edwards, and C. E. Thomas, Phys. Rev. **D79**, 094504 (2009), 0902.2241.

[13] S. Basak et al., Phys. Rev. **D76**, 074504 (2007), 0709.0008.

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

[15] J. Bulava et al., Phys. Rev. **D82**, 014507 (2010), 1004.5072.

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

[17] C. Morningstar et al., AIP Conf. Proc. **1257**, 779 (2010), 1002.0818.

[18] J. Foley et al., AIP Conf. Proc. **1257**, 789 (2010), 1003.2154.

[19] Department of Energy FY 2012 Congressional Budget Request, `http://www.cfo.doe.gov/budget/12budget/Content/Volume4.pdf`.

[20] S. Beane, PoS **CD09**, 028 (2009), 0912.5388.

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

[22] S. R. Beane et al. (NPLQCD), Phys. Rev. **D81**, 054505 (2010), 0912.4243.

[23] S. R. Beane et al. (NPLQCD) (2010), 1012.3812.

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

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

[26] M. A. Clark, R. Babich, K. Barros, R. C. Brower, and C. Rebbi, Comput. Phys. Commun. **181**, 1517 (2010), 0911.3191.

[27] 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.