# 2009 USQCD Facilities Proposal: Dynamical Anisotropic-Clover Lattice Production for Hadronic Physics


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


(Dated: March 15, 2009)


We propose continuing the production of dynamical anisotropic Clover configurations suitable for spectroscopy and simple hadronic physics. We request using the time under the DOE INCITE 2009 award given to USQCD, and early use 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 anticipated first 6 months of time at the beginning of 2010. While the ORNL Cray is currently being used under INCITE, the gauge production could also use the ANL Bluegene/P. The time requested in this proposal is only for the lattice generation and not for subsequent valence calculations. We request using the equivalent of **70** million Cray XT4 core-hours under the Class *A* proposal guidelines. The online disk storage request is is **2 TB**. The tape storage request is **20 TB**. These two media requests are equivalent to **53880** 6n-node hours.


## I. PHYSICS GOALS

The determination of the excited state spectrum and structure of QCD is a major goal in hadronic physics. The complete combined analysis of available experimental data on the photoproduction of nucleon resonances is the 2009 milestone in Hadronic Physics (HP), and the measurement of the electromagnetic properties of the low-lying baryons is an HP 2012 milestone. The so-called hybrid mesons, in which gluons play an explicit structural role, are equally important; recently, the 12 GeV upgrade at Jefferson Laboratory has received "CD-3" – construction can begin – and the GlueX Collaboration proposal to seek information about exotic mesons is a flagship component of the upgrade, yielding insight into the origins of confinement.

Given the current intense experimental efforts in hadron spectroscopy, the need to predict and understand the hadron spectrum from first principles calculations in QCD is clear. In this proposal, we continue 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, the equivalent of **70** million Cray XT4 core-hours, is justified under the guidelines of a *Class A* proposal.

This proposal focuses on the more computationally demanding part of this program, namely the generation of the lattices themselves which are described in Refs [1, 2]. We believe they are more widely useful besides the accompanying work described here. 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. [3–5], along with the calculation of the exotic $\pi_1$ 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 relevant for JLab's Hall D program, as well as for higher spin states and excited states, at the strange quark mass. This latter part of the project is a continuation of the charmonium spectrum project and transition work published in Ref. [6–8] where the first excited state transition form-factors within charmonium. The baryon spectrum calculation, a continuation of the work in $N_f = 0$ (Ref. [9]) and $N_f = 2$ (Ref. [10]), will use the multiple volumes available to determine the light quark and strange quark octet masses. Both the baryon and meson spectrum determinations will involve new techniques for multi-hadron operators. There is also a proposal, a combination of several from last year, for ground state and excited state transition baryon form-factors (*K. Orginos*) relevant to JLab's experimental program.

In addition, we are also aware that NPLQCD (*M. Savage*) intends to use these lattices for their multi-hadron investigations. Their recent very high statistics work (Ref. [11]) shows clearly how a fine temporal resolution is effective at isolating noisy signals. Also, we are aware that EMC (*B. Tiburzi*) also intend to use these lattices.

The larger lattices that we propose will also be used to study hadron-hadron interactions of two baryons, each containing one very heavy quark that provides an unambiguous center-of-mass. With the anisotropic lattices, we can use two-baryon sources and two-baryon sinks, with a varying displacement $R$ between the two baryons, in order to study the Born-Oppenheimer potential $V(R)$ between the baryons. Our operators for ground-state baryons have shown that baryonic states achieve a plateau quite close to the source and this will be essential in order to determine the interaction potential. The use of baryonic operators with different spin projections will allow us to study the spin-dependence of the static interactions.

## II. COMPUTATIONAL STRATEGY

### A. Actions and Parameters

The use of anisotropic lattices is essential to this project. The dense excited state of baryons extends to around 3 GeV. Using conventional isotropic lattices, with spacings around 2 $\text{GeV}^{-1}$, one can expect significant lattice discretization errors. Using a highly anisotropic lattice, e.g. with temporal lattice spacing of 6 $\text{GeV}^{-1}$, the discretization errors are reduced and better signals are provided for excited states. A well defined positive-definite transfer

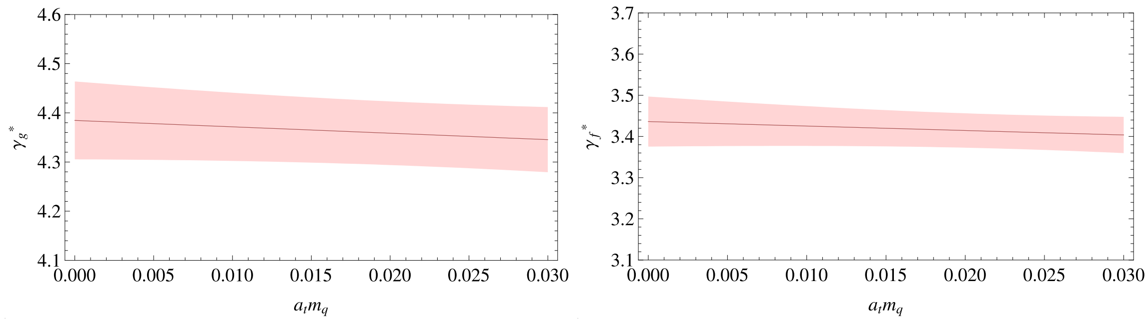


FIG. 1: Tuned gauge and fermion action parameters, $\gamma_g^*$ (left) and $\gamma_f^*$ (right), as functions of the unrenormalized PCAC quark mass, $a_t m_q$. There is negligible mass dependence in the tuned anisotropy parameters. With a renormalized anisotropy of $\xi = 3.5$ and a spatial lattice spacing $a_s = 0.1227(8)$ fm, the maximum extent of the horizontal axis is about 170 MeV.

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 [12, 13] with *Stout-link* smearing [14] of the *spatial* gauge fields in the Clover action. E.g., 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 $1 \times 2$ rectangle in the time direction. This gauge action was used for the determination of the glueball spectrum [15]. First anisotropic $N_f = 2$ stout-link smeared simulations, but with the Hamber-Wu action, have been performed in Ref. [16]. 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 [17]. 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 to be very near to 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]. We show the final result of the tuning in Fig. 1. This plot shows the tuned bare gauge and fermion anisotropies as a function of the unrenormalized PCAC quark mass. These tuned anisotropies ensure that the renormalized gauge and fermion anisotropies are equal to each the desired value $\xi = 3.5$. Very importantly, we find no statistically significant dependence of the bare anisotropies on the mass.

The lack of mass dependence is in contrast to the work of CPPACS who used $N_f = 2$ tadpole-improved Clover but with no stout-smearing, and also the Iwasaki gauge action [18]. They find a strong mass dependence which is presumably related to their different choice of fermion action (no projector property), no stout-link smearing (the choice of clover co-

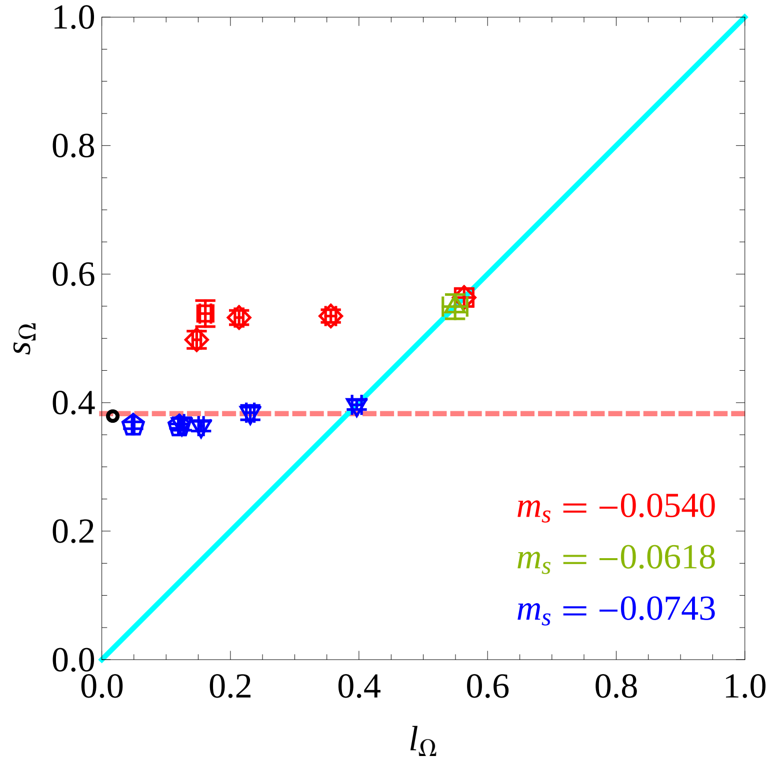


FIG. 2: The "Newport News" plot. The horizontal and vertical axes are ratios of QCD stable hadron masses whose leading order $\chi$PT dependence is on the light quark ($m_l$) and strange quark masses ($m_s$). Shown are 3 different strange quark mass simulations. The physical limit is the black circle. 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 red symbols are $12^3 \times 96$ and $16^3 \times 96$ $N_f = 2+1$ runs varying $m_l$. The red square is a $16^3 \times 96$ simulation at the same mass as the lightest $12^3 \times 96$ point - the latter is presumably finite volume effected. The pink dashed line is a line of slope 0 anchored by the corresponding $N_f = 3$ point. The blue symbols are the current $N_f = 2+1$ simulations. The pentagons are $24^3 \times 128$ runs while the others are $16^3 \times 128$. 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).

efficients becomes important), the different choice of clover coefficients (this is discussed some in Ref. [13] where this form of fermion action arises), and the choice of gauge action. 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 results have appeared in Ref. [2]. The idea is to parameterize observables with dimensionless coordinates involving the light and strange quark mass

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

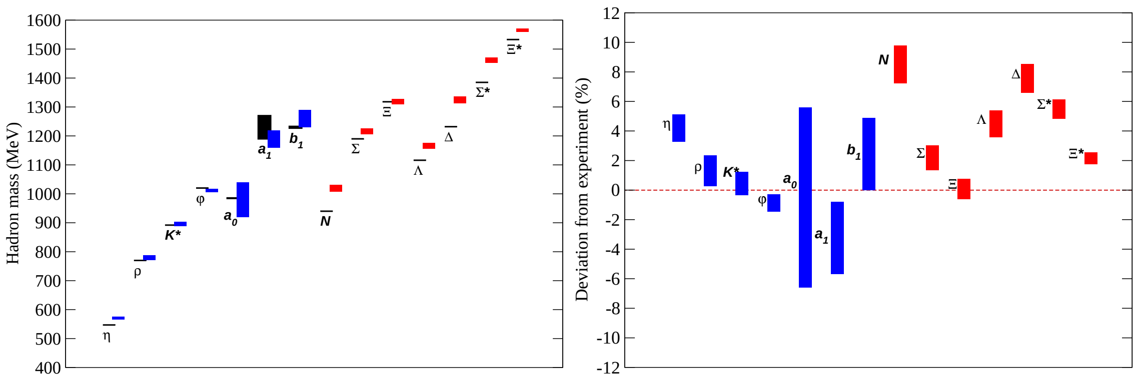


FIG. 3: Summary of the extrapolated hadron masses compared with their experimental values. The extrapolation follows the linear form in Eq. 2 using the coordinates $l_\Omega$ and $s_\Omega$.

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

The strange quark tuning strategy is then to consider the $N_f = 3$ case, tune the bare quark mass till $s_\Omega$ 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 2. We see the strange quark mass choice that projects onto the physical value of $s_\Omega$ is a fairly reasonable choice as the light quark mass has decreased.

We observe that the chiral perturbation theory of observables can be recast using these coordinates. We used this method to extrapolate our hadron masses to the physical limit using a simple linear form

$$\frac{a_t m_H}{a_t m_\Omega} = d_0 + d_l l_\Omega + d_s s_\Omega \quad . \tag{2}$$

The results for our data sets are shown in Figure 3 as published in Ref. [2]. All the data is used (except for the small volume data when a larger volume is available. Also, the 230MeV data set was not available at the time of Ref. [2]). The largest discrepancy with experiment is 10%. We found that if we considered only the data from the final strange quark mass, we obtain very consistent results. Thus, the higher order corrections are apparently small over the mass ranges found here. We note we could also consider higher order forms of this chiral perturbation and reexpress the standard expressions for observables in terms of $l_X$ and $s_X$. The new form can still be considered a mass independent scheme.

We note that the recent BMW calculation in Ref. [19] also used these coordinates but with the reference mass the $\Xi$ baryon – another equally fine choice, but one that doesn't allow as much leverage in determining the strange quark mass. There results have much smaller pion masses than we have available.

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 that determining resonance states from their decay or scattering states as the principal concern in this project. 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 it is only there that we believe we can state with some confidence that we have determined the hadronic spectrum, albeit at some fixed lattice cutoff.

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 $\beta$, but at a fixed $l_\Omega$ and $s_\Omega$.

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.

The current inventory of datasets is shown in Table I. We have finished high statistics $16^3 \times 128$ runs that enable first attempts at characterizing the spectrum and provide small volumes for a multi-volume study. This on-going work is described in the proposal by *D. Richards*. Also finished are $20^3 \times 128$ and $24^3 \times 128$ runs at $m_\pi = 383$MeV. A subset of this data was used for our initial ground state spectrum determination in Ref. [2]. NPLQCD has used the $20^3 \times 128$ at $m_\pi = 383$MeV for their recent high statistics measurements in Ref. [11].

The focus in the short term is finishing the remaining $20^3 \times 128$ runs, and the $24^3 \times 128$ run at $m_\pi = 230$MeV. The DOE time requested will go towards the $32^3 \times 256$ run at this lower mass.

## III. SOFTWARE

Our dynamical anisotropic Clover gauge generation has been using the *Chroma* HMC code with RHMC for the three flavors, and multi-timescale integration. Currently, the Clover action is four-dimensionally even-odd preconditioned. We have compared RHMC to Hasenbüsch style mass preconditioning for the two light flavors. We have recently switched to using a mass preconditioning with a twisted mass version of clover The twisted mass clover term is then handled by RHMC. This technique has stabilized the lower eigenvalues (a problem at small quark mass) which needs to be bounded for RHMC.

The Clover CG inverter runs at about 265 Mflops per node on a QCDOC rack for typical subgrid sizes such as $4^4$. On the 6n cluster, the CG runs at about 1950 Mflops/6n-node while the multi-mass inverter - what is used in RHMC - runs at typically 1400 Mflops/6n-node on both the 6n and 7n clusters up to 512 cores and $16^3 \times 128$ lattice sizes.

Extensive optimizations have gone into the code to support the Bluegene architectures including support for the Bagel utilities. In particular, Clover runs have been made on the San Diego BG/L and the ANL BG/P. We stand ready to incorporate further BG/P improvements in the future.

| $m_\pi$(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 | $12^3 \times 96$ | 9980 |
| | | | $16^3 \times 128$ | 11735 |
| | | | $20^3 \times 128$ | 5100 |
| 560 | -0.0808 | -0.0743 | $16^3 \times 128$ | 11010 |
| | | | $20^3 \times 128$ | 3995 |
| 448 | -0.0830 | -0.0743 | $16^3 \times 128$ | 13470 |
| | | | $20^3 \times 128$ | 2130 |
| 383 | -0.0840 | -0.0743 | $16^3 \times 128$ | 11005 |
| | | | $20^3 \times 128$ | 13090 |
| | | | $24^3 \times 128$ | 12985 |
| | | | $32^3 \times 256$ | *warming* |
| 230 | -0.0860 | -0.0743 | $24^3 \times 128$ | 5000 |
| | | | $32^3 \times 256$ | *warming* |

TABLE I: Current inventory (Mar. 16, 2009) 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.

At ORNL, after the dual core upgrade of the machine to a single socket quad-core Cray XT4, and the "upgrade" of the operating system to Compute Node Linux (CNL), we find a Mflop/core performance resulting in 1400 Mflops/cray-node which is equal to the performance of the 6n cluster. We note that on comparing the performance before and after the operating system change to Linux, the performance has dropped by 25%.

As part of our on-going software development, we have reimplemented the inverter and all of HMC to use version of the Clover action with "temporal preconditioning" (Ref. [20]). Here, the one-dimensional time derivative of the action is pulled out and inverted directly. We find that the condition number (and hence the CG count) of the Clover-Dirac operator to be improved by a factor of 2.6. Unfortunately, we see no real improvement in the inverter times. We need a three-dimensional communication pattern - time is kept local. Thus, for long time extents, the surface area is larger than in a corresponding four dimensional pattern, resulting in more time needed for communications, at least for very large core counts.

We have finished our first implementation of threading within QMT, QDP++ and Chroma. We are using a hybrid model – there is one master thread that handles the communications, and all the cores on the node participate in the computations. This technique coalesces messages into larger ones compared to a non-hybrid version, thus improving scaling.

The ORNL run-time environment forces us to run at very large core counts - typically a

| $L_s$(fm) | 2.45fm | 2.95fm | 3.93fm | 4.91fm |
|---|---|---|---|---|
| $m_\pi$(MeV) | $20^3 \times 128$ | $24^3 \times 128$ | $32^3 \times 256$ | $40^3 \times 256$ |
| 833 | 6k, TACC[1.0M](10) | | | |
| 560 | 7k, TACC[1.5M](6.7) | | | |
| 448 | 8k, TACC[2.1M](5.4) | | | |
| 383 | 13k, done | 13k, done | 11k, Tenn[22M](7.4) | |
| 230 | | 6k, PSC[6M](3.2) | 11k, ORNL[70M](4.2) | |
| 140 | | | | 11k, INCITE[390M](3.4) |

TABLE II: Cost estimates in Millions of cores hours for lattice production. Here, the actual times on the big machines are used. The entries are #traj,location[M-core-hours]($m_\pi L$). The number of traj. is indicated by the "k". The goal is to produce at least 11k trajectories. Some of the production is partially completed, and the remaining production cost is quoted. Configurations are saved every 5th trajectory. The current NSF production is indicated in blue. The USQCD (red) entries are what are proposed for the 2009 allocation period. The ORNL run ($32^3 \times 256$ at 230MeV) is underway now. The current production could be moved to another system. The physical limit calculation (green) run was proposed for early use time and could take advantage of other resources. The cost for the physical limit calculation is assumed to scale like Eq. 3, namely like $m_\pi^{-2}$.

minimum of 8k cores. 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 – 16k cores. 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.

Finally, we are investigating a new deflation technique suitable for HMC. There is a potential for further decreases in our times per trajectory.

## IV. REQUIRED RESOURCES

To determine our computational needs, we must take into account how the ratio of spatial to temporal lattice spacings $a_s/a_t = \xi$ affects scaling. We use the scaling ansatz

$$\mathrm{Cost}_{\mathrm{traj}} = C\xi^{1.25}\left(\frac{\mathrm{fm}}{a_s}\right)^6 \cdot \left[\left(\frac{L_s}{\mathrm{fm}}\right)^3\left(\frac{L_t}{\mathrm{fm}}\right)\right]^{5/4} \cdot \left(\frac{135\mathrm{MeV}}{m_\pi}\right)^2 . \tag{3}$$

This is the cost of producing a fixed number of trajectories, and does not take into account any critical slowing down in the quark mass, hence the scaling like $1/a_s^6$. We note that there is an overall factor of $\xi^{1.25}$ arising from the volume$^{5/4}$. The upshot is that we have taken one direction close to the continuum limit, so we must pay for it. We normalize the scaling with the timing on $24^3 \times 128$ at $m_\pi = 230$MeV.

The current and proposed production is shown in Table II The main focus in this allocation year will be production on $32^3 \times 256$ at $m_\pi = 230$MeV using INCITE resources. For 11k trajectories, the total time estimated is about $70M$ Cray XT4 core hours. We are

| volume | Existing TB | Proposed TB | Total TB |
|---|---|---|---|
| $12^3 \times 96$ | 0.4 | | 0.4 |
| $16^3 \times 128$ | 1.3 | | 1.4 |
| $20^3 \times 128$ | 1.3 | 1.1 | 2.4 |
| $24^3 \times 128$ | 1.7 | 0.6 | 2.3 |
| $32^3 \times 256$ | 0.2 | 9.7 | 9.9 |
| total | | | 17 |

TABLE III: Storage requirements for gauge configurations.

using NSF resources to finish $20^3 \times 128$ runs at the heavier masses, and NSF resources to finish the $24^3 \times 128$ run at $m_\pi = 230$MeV, and the start of a $32^3 \times 256$ run to complete the $m_\pi = 383$MeV mass set. We anticipate having the physical limit calculation to finish the $\beta = 1.5$ datasets.

As more resources become available, our first choice would to finish and/or extend the $32^3 \times 256$ run at $m_\pi = 383$MeV. We would also like to take an initial start at the physical limit calculation. While the time need is large, valuable experience would be gained by making measurements on a limited statistics dataset.

## V. USE FOR ADDITIONAL RESOURCES

We note that the INCITE allocation is at ORNL, but the production could be moved to the BG/P at ANL or other resources. The HMC code has been extensively tested and optimized for the BG/P. The physical limit calculation on $40^3 \times 256$ was proposed for early use time at ORNL, but was not approved. It also could be moved to other available resources.

## 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

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 III. It is expected that **17 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

| media | Storage TB | 6n-equiv |
|---|---|---|
| disk | 2 | 26940 |
| tape | 3+17 | 26940 |
| total | | 54880 |

TABLE IV: Storage requirements turned into 6n-equivalent node hours

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 **2 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 IV. The 6n-equivalent cost is **53880** 6n-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, and D. G. Richards, Phys. Rev. **D73**, 074507 (2006), hep-ph/0601137.

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

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

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

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

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

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

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

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

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

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

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

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

[19] S. Durr et al., Science **322**, 1224 (2008).
[20] R. Edwards, B. Joó, and M. Peardon, in preparation (2008).