2019 USQCD Facilities Proposal: Meson Resonances and their Couplings from Anisotropic Clover Lattices

The Hadron Spectrum Collaboration

—

Raúl Briceño
Old Dominion University

Robert Edwards, Bálint Joó, Luka Leskovec, Frank Winter
Jefferson Lab

Jozef Dudek, *Christopher Johnson, Archana Radhakrishnan*
College of William and Mary

Bipasha Chakraborty, Christopher Thomas, *Antoni Woss*
University of Cambridge, Cambridge, UK

Maxwell Hansen
CERN, Switzerland

Nicolas Lang, Michael Peardon, *David Tims*, Sinéad Ryan, David Wilson
Trinity College, Dublin, Ireland

Nilmani Mathur
Tata Institute, Mumbai, India

—

(students in italics)

We propose to advance our calculations of the meson resonance spectrum using $32^3 \times 256$ anisotropic-clover lattices at a pion mass of 230 MeV, as well as new ensembles of $24^3 \times 256$ at pion masses of 275 MeV and 325 MeV. These investigations are aimed at tackling the challenges that arise in scattering amplitude studies as the quark masses are progressively decreased, notably those due to inelastic and multi-channel decays. We will determine resonance properties, including mass, widths and coupling to decay channels as the pion mass changes, using the rigorous finite-volume approach to coupled-channel scattering. Notably, these calculations will, for the first time, attempt to understand the contribution of three-body channels to scattering amplitudes.

We request 2.0M K80-equivalent hours on the GPU cluster at JLab, and 102.4M KNL core-hours on the KNL system at JLab. Finally, we request 100 TByte of new tape storage, equivalent to 0.13M Sky core-hours, and 500 Tbyte online disk storage equivalent to 3.75M Sky core-hours.

I. INTRODUCTION

In the last decade we have witnessed tremendous amount of progress towards the determination and understanding of the excited state spectrum of hadrons within Quantum Chromodynamics (QCD). The progress has been partly driven by the various worldwide experimental efforts, within these is included the GlueX experiment of the new Hall D at Thomas Jefferson National Accelerator Facility (Jefferson Lab). The GlueX experiment is one of the flagship experiments of the $338 million upgrade of Jefferson Lab, and its primary purpose is to provide a detailed description of "exotic" mesons. These are states that are expected to provide the most direct evidence of new configuration of matter, where gluonic degrees of freedom are manifest, known as "hybrids".

The proposed program is the continuation of a long term effort to systematically study these states directly from the QCD using the only rigorous means of non-perturbatively determining QCD observables, namely lattice QCD (LQCD). Phenomenologically, all exotic candidates manifest themselves in the QCD spectrum as resonances that quickly cascade to multiple channels containing two as well as three QCD-stable hadrons. In other words, they are understood as dynamical enhancements in the scattering amplitudes of byproducts. Consequently, it is not too surprising that in order to study *exotica* using LQCD one must determine the scattering amplitudes with their quantum numbers. This is a necessary theoretical effort running in parallel to experimental searches. As concisely expressed by the most recent report by the DOE Nuclear Science Advisory Committee for the 2015 Long Range Plan: "Underscoring this huge progress, LQCD plays an essential role in guiding experimental work. GlueX at JLab, one of the flagship experiments of the 12-GeV Upgrade, is designed to search for exotic particles where the glue is in an energetically excited state. Initial LQCD calculations motivated the experiment and guided its design. Recent LQCD results confirm the mass range of the predicted particles. And in the future, LQCD calculations of hadron dynamics will play a critical role in the analysis of the data."

The goal of this program is to serve as a guide, confirm, and ultimately complement experiment. This is being done by determining the prominent decay channel, provide constraints in the dynamical coupling describing the production and decay mechanism, and ultimately presenting a structural decomposition of these states which will be able to give insight to their true nature beyond what is experimentally feasible.

In the last 10 years we have developed most of the necessary numerical and formal tools to carry out a program to determine the scattering amplitude of a wide range channels, including those with exotic quantum numbers . As is reviewed below, during the developmental stages our attention has primarily been focused on unphysically heavy values of the quark masses, where these calculation were tractable for a variety of reasons discussed. At this stage, we are able to start carrying out the calculation of exotica using the physical values of the quark masses.

Given the set of challenges associated with the study of resonant scattering amplitudes using the physical quark masses, we have outlined a multistep procedure towards the study of exotica. In the first stage, which is the first milestone of this project, we will study the spectrum of the byproducts of the exotic mesons. This is a necessary first step.

The Titan and Summit systems have been crucial for this project. The *Chroma* code and the *QUDA* library have been well optimized for these systems. A major development has been the optimization of the HMC gauge generation for these systems. The integration of the highly optimized multi-grid solver for GPUs reduces the communications and lowers wall clock times significantly compared to previous years. The integration of the multigrid solver into the gauge generation has reduced the computational resources going from Titan to Summit by **73x**. This tremendous advance has accelerated our project, and under Summit Early Science Program, we have generated enough of these physical limit gauge fields that we can make first calculations of the spectrum of QCD that is directly comparable to experiment and thus guide and influence new experimental searches.

The current USQCD project will capitalize on these developments. We will use recently generated gauge configurations, with size $48^3 \times 512$, using values of the quark masses in the physical limit of QCD. On the GPU systems at JLab, we will compute propagators on the lattices. Using the KNL clusters, we will carry out first studies of low-energy scattering observables on these lattices, and on existing lattices at heavier than physical values.

II. CURRENT STATUS

To this day there are no published results of first principle studies of resonant scattering amplitudes in channels with exotic quantum numbers. Our project is pursuing these calculations now using unphysically heavy values of the light quark masses. We will review the obstacles that we have had to overcome to reach the present status, as well as those that we are actively addressing to achieve the ultimate goal of performing studies of exotica at the physical point of QCD.

Given that lattice QCD is the only systematically improvable non-perturbative method available to obtained QCD observables, our program and the proposed project focuses on this. Lattice QCD is necessarily defined in a discretized and finite volume. For the observables we are considering discretization effects are typically small and can be systematically removed.

On the other hand, finite volume effects leads to one of the major challenges in the study of scattering amplitudes and consequently resonances. This is related to the fact that there are no asymptotic states in a finite volume, which implies that there is no direct mechanism for extracting scattering amplitudes from finite-volume observables. Fortunately, we can make use of non-perturbative relations, that we have partly developed, to map the LQCD-determined finite-volume spectra onto the infinite-volume amplitudes. These relations become increasingly complicated to implement is the presence of a large number of thresholds and where the decays involve states with more than two particles. This is indeed the case for exotic mesons. For example, the lowest lying π_1 resonance has the quantum numbers of $\pi\eta, \pi\eta', \pi(\rho \to 2\pi), \ldots$. As we discuss below, we have made significant progress on studying resonances that couple to multiple two-body channels. We developed the necessary formalism, techniques, and to this date are still the only collaboration to carry out studies of coupled-channels systems. This has culminated in the first study of the π_1 for unphysically heavy quark masses, where it only couples to two-body channels. Our collaboration is also leading efforts to extension of this framework for systems where three-particle channels can go on-shell, which will allow for future studies of the π_1 are lighter and eventually physical values of the quark masses.

As we have made steady progress towards this goal, we are now at a stage where we can begin the generation of quark propagators using physical values of the quark masses on recently generated gauge configurations. The first physics outcomes of this new project will have to be determination of the low-lying two-particle scattering amplitudes and resonances. This is a necessary first step towards studying exotics, as we have proven that in order to

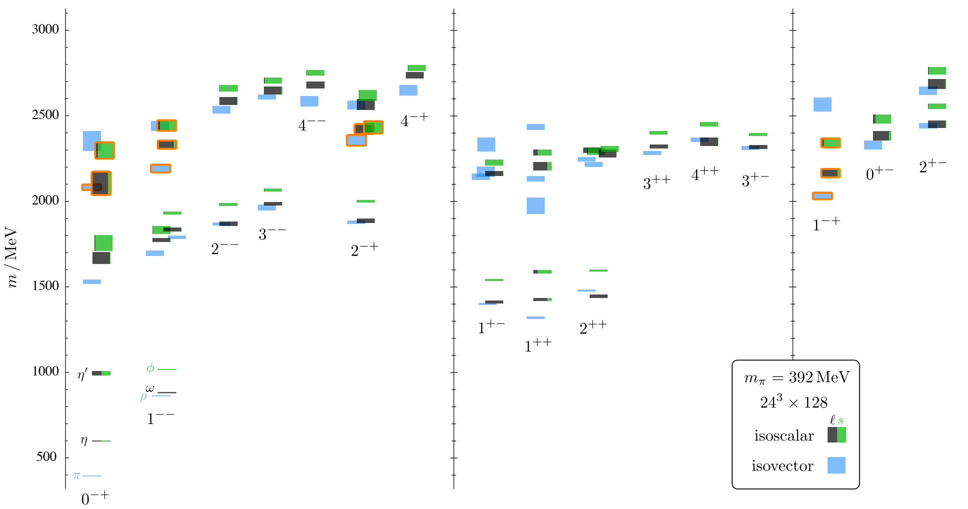

FIG. 1: From Ref. [2]: Isovector and isoscalar meson spectrum. The vertical height of each box indicates the statistical uncertainty on the mass determination. The existence of the exotic states in the right panel is a major focus of the GlueX experiment at Jefferson Lab. The states outlined in orange are the lowest-lying states having dominant overlap with operators featuring a chromomagnetic construction – suggesting their interpretation as the lightest hybrid meson supermultiplet.

study three-particle decays one must determine the scattering amplitude of all possible two-particle subsystems [1]. Finally, this set of physical ensembles, in conjunction with previously generated ensembles at a range of values of the quark masses, will give us a unique insight into the nature of the QCD spectrum.

A. Spectroscopy and Amplitude Analysis within Lattice QCD

As reported in a series of papers, the highly-excited spectrum has been computed both for mesons and for baryons, and we summarize here some of the salient results for this proposal.

A major motivation for previous ALCC awards came from the first calculation of the isovector and isoscalar meson spectrum on our dynamical anisotropic lattices with one strange quark and two light-quark masses corresponding to a pion mass of 392 MeV. Phenomenological insights from this calculation include the observation of a detailed spectrum of states, shown in Fig. 1, having a large overlap with operators having essential gluonic content that had been suggested [2–4] as good candidates to be hybrid mesons. The results suggest that the masses of exotic mesons are near 2 GeV which is within the experimental reach of JLab's GlueX experiment.

The next major step, which has been a major focus of this program in recent years and the focus of our previous milestones, is the determination of the decay modes of the states. The observed spectra that have been shown in Figures 1 do not allow us to extract this information as the expected multi-hadron levels do not appear. We address this by adding into the operator basis a set resembling multi-mesons [5]. Once these operators are included, we are then in a position to extract resonance information.

Having obtained the discrete energies from finite-volume lattice QCD calculations, we can then relate these to the infinite-volume scattering amplitudes. The required formalism for

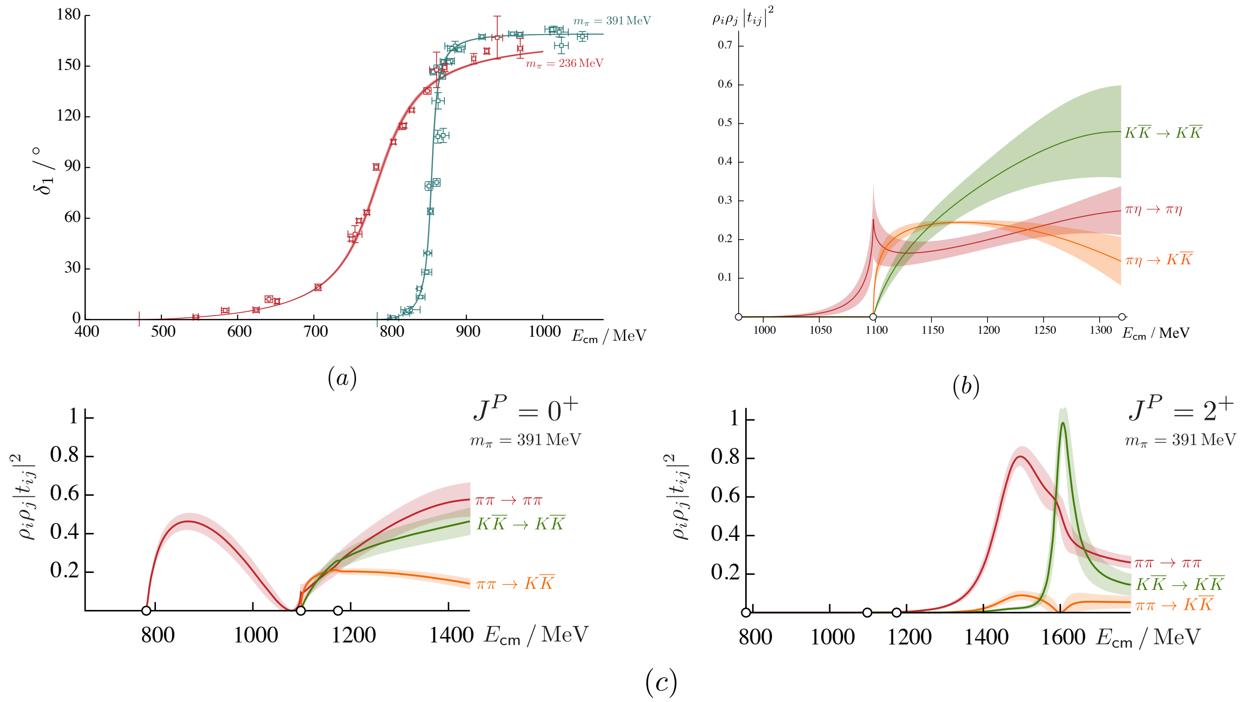

FIG. 2: (a) Shown is the resonant $\pi\pi$ scattering phase shift as a function of energy obtained using two values of the $m_\pi = 236, 391$ [5, 6]. (b) Shown are the values of the components of the scattering amplitude for the isotriplet $\pi\eta - K\overline{K}$ coupled system determined using $m_\pi = 391$ MeV [7]. The resonant behavior is due to the presence of the a_0 resonance. (b) Shown are the values of the components of the scattering amplitude for the S- and D-wave isoscalar $\pi\pi - K\overline{K}$ coupled system determined using $m_\pi = 391$ MeV [8]. The resonant behavior is due to the presence of the σ, f_0, f_2 and f_2' resonances.

relativistic elastic scattering was first developed by Lüscher [9, 10] for the case of a system in its rest-frame, and was extended to moving frames in [11–13]. In the last few years, we had to generalized these ideas to arbitrary two-body system [14] in order to be able to first the determination of a exotic resonance.

We have used this methodology now to study several important channels. The first resonant implementation was in the isovector $\pi\pi$ channel [5, 6], shown in Fig. 2(a) for two values of the quark mass corresponding to $m_\pi = 236, 391$ MeV. Most recently, these same ensembles were used to study the isoscalar $\pi\pi$ channel, where the enigmatic σ resonance resides. This is a controversial state whose existence was a source of debate for over half a century. By investigating the energy dependence of the scattering amplitude of this channel, shown in Fig. 3(a), we were able find the manifestation of this state for the two values of m_π studied. As seen in Fig. 3(b), for the heavier ensemble, it is a bound states, while for the ensemble closer to the physical point, it is a broad resonance resembling the experimental one. Being the first QCD evidence for this state, the results were published in Phys. Rev. Lett. [15] and article was showcased in the cover of Phys. Rev. Lett. for the January 13th edition, Fig. 3(c).

Higher energy resonances, in general, involve more than one decay channel, and hence involve coupled-channel scattering amplitudes. The finite-volume formalism has been ex-

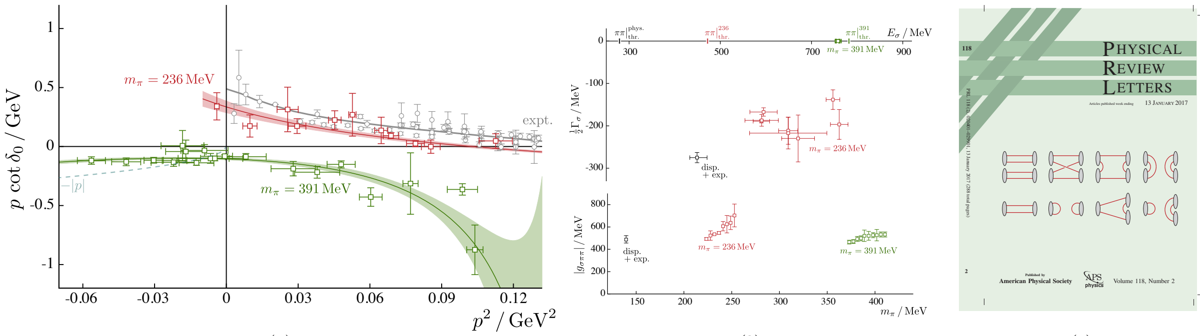

(a) (b) (c)

FIG. 3: Results from Phys. Rev. Lett. [15]. (a) Shown are the scattering amplitudes of the scalar/isoscalar elastic $\pi\pi$ mesons for two values of the light quark masses corresponding to $m_\pi = 236, 391$ MeV. The red and green square point correspond to the values constrained directly from the lattice QCD spectrum. This result is compared with previous experimental values shown in light grey. The energy dependence of these scattering amplitudes is used to determine the properties of the enigmatic σ resonance. (b) The top panel shows the values of the mass and width of σ as a function of the quark masses. The bottom panel shows the $\sigma \to \pi\pi$ coupling. (c) Shown is the cover of the Phys. Rev. Lett. edition of January 13th, 2017, where this calculation was showcased.

tended for such cases in Refs. [14, 16–19]. We have been able to successfully implement the algorithmic and formal technology to extract scattering amplitudes for a few important examples where multiple channels are kinematically open. This was first demonstrated for the scattering of $K\pi$ and $K\eta$ which resulted in two publications Phys. Rev. Lett. [20], and a longer Phys. Rev. D [21]. This channel turned out to be a nearly decoupled system. We later studied the $\pi\eta, K\overline{K}$ scattering amplitude in the isovector channel. As in experiment, this was demonstrated to be a strongly coupled system, and it was the first of its kind to be studied via lattice QCD. The three components of the scattering amplitude, $\pi\eta \to \pi\eta$, $\pi\eta \to K\overline{K}$, and $K\overline{K} \to K\overline{K}$, are depicted in Fig. 2(b). The fact that this is indeed strongly coupled is evident from the fact that the $\pi\eta \to K\overline{K}$ component of the scattering amplitude is of the same magnitude as the other two for the kinematic regime constrained from the finite-volume spectrum. Also evident is the strong resonant behavior near the $K\overline{K}$ threshold, which is shown to be due to the presence of the a_0 resonance.

We have most recently extended the results presented published in Phys. Rev. Lett. [15], by considering energies above the $K\overline{K}$ and $\eta\eta$ thresholds [8]. This exploratory study was carried out using the ensembles corresponding to $m_\pi = 236, 391$ MeV. The results are summarized in Fig. 2(c). The 0^+ amplitudes show strong mixing between the $\pi\pi$ and $K\overline{K}$ scattering states. In addition to the previously found σ, we found the first observation of the would be $f_0(980)$ from QCD. The 2^+ partial wave amplitudes show clear evidences of the low-lying f_2 and f_2' resonances. This study culminated the first complete investigation of the low-lying scalar and tensor nonet via lattice QCD for given set of quark masses.

To date, there are no calculations of a hadronic resonance using physical values of m_π. There has been several historical limitations for this, all of which we have been systematically addressing. In this project, we propose to carry out the first exploration of the a resonant amplitude at the physical point. The first calculation that will be carried out is that of the isovector $\pi\pi$ amplitude. Beyond providing a perfect testing ground for these ideas, as

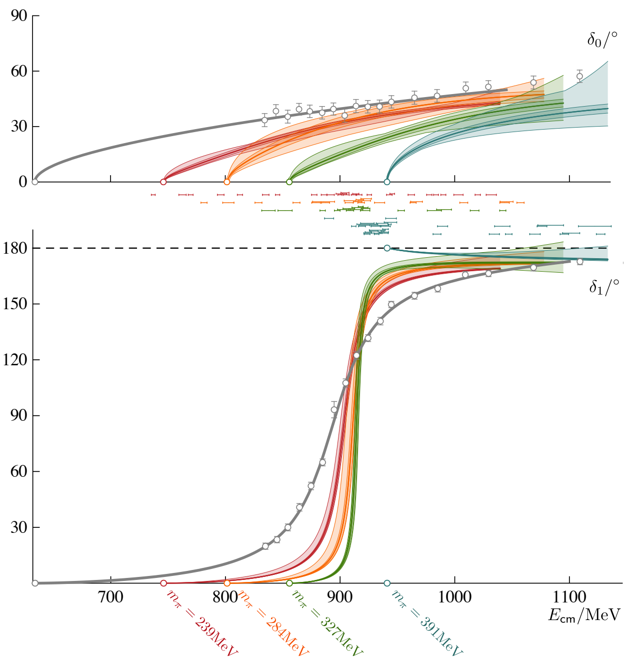

FIG. 4: Preliminary results for the quark mass dependence of isospin-$\frac{1}{2}$ elastic $K\pi$ scattering amplitudes determined under ALCC. Shown are S-wave (top) and P-wave (bottom) phase shifts, determined at four different pion masses, indicated by their color and legend at the bottom of the plot. The central line indicates the mean. The colored points in the middle are the energy levels determined in rest and moving frames. The outer bands include the uncertainty over parameterizations, mass and anisotropy variations. The grey points are experimental data from the LASS experiment, and the grey curves are a guide from simple parameterizations. The results suggest the K^* appears as a resonance with a coupling that is nearly constant across quark masses, while the origin of the putative κ meson is less clear. Establishing the resonant behavior of these amplitudes is important for disentangling the decays of higher mass states, such as exotic mesons.

illustrated by Fig. 2(a), the resonant isovector $\pi\pi$ state plays an important role in a variety of interesting phenomenology, which we discuss below.

The low energy scattering of kaons and pions provides another window into the most basic properties of QCD. Such as $\pi\pi$ scattering described above for the σ and the ρ, there is the possibility of a strange quark mass version of such states. The origin of such states is interesting in its own right since the analytic properties of the scattering amplitudes is expected to be more complicated that the case of $\pi\pi$. Shown in Fig. 4 is a preliminary study of the quark mass evolution of S-wave and P-wave phase shifts in isospin-$\frac{1}{2}$ elastic $K\pi$ scattering. Understanding how resonances can appear as thresholds are changed with the light quark masses is crucial as these low-energy states will feature prominently in decays of higher energy states.

At higher energies hadrons, including exotic mesons, can decay into multi-meson final states through an intermediate states featuring resonances of non-zero intrinsic spin. To make such studies possible, we first extended the underlying formalism relating finite-volume spectra to amplitudes to incorporate the possibility of spinning particles [14]. In Ref. [22] we reported on the first calculation of the energy dependence of partial-wave scattering amplitudes for $\rho\pi$ in isospin-2, including the coupled S and D-wave system with $J^P = 1^+$. In this exploratory study, we used quarks with pion mass of 700 MeV, so the ρ becomes a stable hadron lying some way below the $\pi\pi$ threshold. We find clear evidence for dynamical mixing between the coupled S and D-wave channels. This is also the first non-zero determination of an observable that is dependent on the tensor hadronic force using LQCD. Establishing this so called 'spinning-hadron' analytic formalism was a necessary prerequisite before tackling exotic meson systems.

Finally, we highlight a technical development that is of relevance when studying exotica. While we have investigated meson-meson systems, we recognized the significance of multi-hadron operators in accessing the finite-volume spectra. This leads one to naturally ask if other operators may be of significance. Particularly, it is clear that dynamical degrees of freedom could include tetraquark-like constructions as suggested to explain recent observations in the charmonium sector. We have carried out an initial investigation in charmonium including tetraquark as well as meson-meson operator constructions in Ref. [23]. This is not just a technological achievement, but it also points to no strong indication for any bound state or narrow resonance in these channels, including the $Z_c(3900)$, which is interesting on its own right.

B. Three-body scattering

As we decrease the quark mass toward the physical limit, three-body thresholds can become open for resonant decays, and it is crucial that a consistent formalism be developed to accomodate such decay modes. There has been steady progress towards these goals recently reported in Ref. [1]. A relativistic, model-independent three-particle quantization condition has been derived for the case of identical scalar particles with a G-parity symmetry and where the two-particle K matrix has a pole in the kinematic regime of interest. The result involves intermediate infinite-volume quantities with no direct physical interpretation. It was shown how these quantities are related to the physical three-to-three scattering amplitude by integral equations. This work opens the door to study processes such as $a_2 \to \rho\pi \to \pi\pi\pi$, in which the ρ is rigorously treated as a resonance state.

III. PROPOSED PHYSICS PLAN

We propose two major threads of calculations using the existing anisotropic Clover ensembles, and the new datasets on $32^3 \times 256$ and $24^3 \times 256$ with $m_\pi = 275$ MeV, and $24^3 \times 256$ with $m_\pi = 325$ MeV we have generated using external resources. We will compute two-point functions using the $32^3 \times 256$ lattices with $m_\pi = 236$ MeV and $m_\pi = 275$ MeV, and the two new $24^3 \times 256$ ensembles. We will use these new lattices to compute the quark mass dependence of resonance couplings.

We will compute propagators on the new $48^3 \times 512$ lattices with $m_\pi = 170$ MeV. Once these are ready, we will use them in two-point function contractions to compute scattering amplitudes.

A. Meson-Meson scattering

A principal goal of our spectroscopy program is establishing the branching fractions for decays of hadrons, including putative exotic mesons. It is these decay couplings that can inform and confront experimental analyses, such as those within GlueX.

Higher energy states are expected to decay into multiple channels with multiple partial waves including S-waves. Thus, a particularly important question concerns the quark-mass dependence of hadronic couplings of resonances. For example, as seen in Fig. 3(b) the $\sigma \to \pi\pi$ is shown to be consistent for the two computed values of the quark mass and the physical point [15]. If this is observed to be generally true, this would add confidence that calculations of exotic states at heavier quark masses can be related to the physical limit.

As mentioned earlier, we have completed first studies of the isovector and isoscalar scalar sector to energies above the $K\overline{K}$ threshold for the heavier $m_\pi = 391$ MeV ensemble [7, 8]. We plan to extend our previous studies of the scalar and vector resonances to lighter values of the quark masses. In particular, we plan to obtain the spectrum of low-lying resonances at $m_\pi \sim 236, 275, 325$ MeV. In total, these resonances include the a_0, a_1, a_2, and the b_1. These calculations will extend the first studies of the $f_0(980)$, $a_0(980)$, and $a_2(1320)$ resonances which we found to be near the $K\overline{K}$ threshold and strongly coupled to $\pi\pi/K\overline{K}$, and $\pi\eta/K\overline{K}$, resp.

With these analyses in place, we can establish the resonance behavior of the light scalar and vector meson sectors of QCD. This is a necessary step before tackling higher energy sectors such as the exotics, and an interesting investigation in its own right.

B. Three-meson scattering

As discussed above, the quantization condition and formalism for three-body scattering [1] has been sufficiently developed that we can now proceed to numerical tests of amplitude determinations. We will use the the range of spatial volumes available for $m_\pi \sim 391$ MeV. We will carry out our calculations for the $\pi\pi\pi$ system first in isospin-3, and then proceed to lower isospins. In isospin-2, the two-particle system corresponding to the ρ will appear resonant in $\pi\pi$. In isospin-1, there are also expected to be resonances in the three-body system. A case in point is the a_1 with scattering amplitudes in channels featuring $\pi\pi\pi$ and $K\bar{K}\pi$. We have already begun testing in three-body systems. These investigations are necessary first steps towards tackling higher energy resonances.

IV. COMPUTATIONAL PLAN

The first part of our request is time for the solution of the light and strange quark Dirac equations to construct the annihilation quark lines needed to construct many-body correlation functions on the physical limit lattices of $48^3 \times 512$. In the language of *distillation*, these are the *perambulators* where we invert on each time-slice of the lattice and on each source vector. We are using $N = 640$ eigenvectors of the three-dimensional gauge-covariant Laplacian as sources. We contract the solution vector with the eigenvectors on those same time-slices. These propagator objects – the "annihilation lines" – can be reused, post-facto, for constructing multi-hadron correlation functions. We need the annihilation lines for both light and strange quarks. The solution vectors will be computed on the GPUs using the QUDA *AMG* code.

The second part of the proposal will use the KNL-s for calculation of the hadron two-point correlation functions on the existing ensembles of propagators on the light quark lattices with $m_\pi \sim 236, 275, 325$ MeV. These calculations will address our program of mapping the quark mass dependence of the low mass resonances of QCD. We have already begun investigating more vigorously the contribution of three-body processes in scattering amplitudes, including those with resonant two-body decays.

The two phases of calculations are independent of the each other. All the perambulators and operator elemental components for the lattices with heavy pion masses already exist.

A. Codes and Libraries

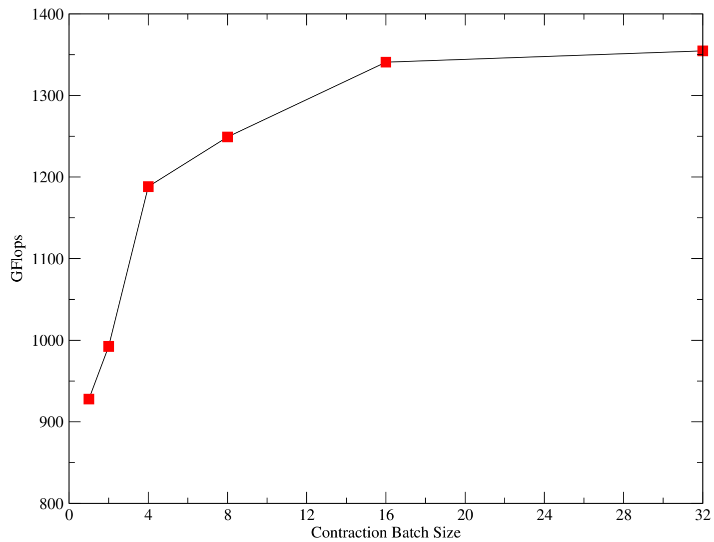

FIG. 5: Performance results of an exterior contraction (*zgemm*) of two mesons with various batch sizes. The performance approaches the peak available for *zgemms* on the JLab KNL system.

We are using in production now optimized codes, called from Chroma, for construction of the hadron correlation functions (contractions) on the KNL systems at JLab, and optimized propagator codes for production on these same systems as well.

We will use the Chroma software system in conjunction with the QUDA library to calculate light and strange quark solution vectors on GPUs. For this purpose, we will utilize the recently developed implementation of Adaptive-Smoothed Aggregation multigrid (AMG) technique implemented in the QUDA library for lattice QCD on GPUs [24–27]. AMG is well suited to the demands of a large number of multiple right-hand sides as the multigrid projection subspace is computed once, and reused for the large number of inversions that are required.

To construct the correlation functions from the perambulators and the corresponding operator elements, we have developed a package, *Redstar*, that automates operator construction and the many Wick contractions that their resulting correlation functions entail [4, 20, 21, 28, 29]. *Redstar* is effectively a work flow engine that coordinates several stages of the calculations. From a list of target correlation functions, the code performs Wick contractions symbolically and produces graphs of the contrations where hadron operators appear as vertices in the graph. Under SciDAC and ECP, the code has been extended to find the minimum number of temporaries – a contraction of the vertices which are tensors in distillation space – required to evaluate the graphs. In general, tensor contractions can be recast as matrix multiplication. In addition to finding the minimum number of temporaries, the code also orders the calculation of those intermediate nodes (temporaries) so as to minimize the memory high-water mark - the largest number of temporaries held in memory at any one time. The code is optimized for dense matrix multiplies on KNL systems using a combination of the Intel MKL library and our own library developed under the Exascale project. The double-precision complex matrix multiplies (*zgemm*) that we must compute per time-slice, necessary for the evaluation of the graphs, is the dominant part of the contraction cost. A technique we use to increase performance is "batching" of the *zgemms* – grouping together into one long stream several of these multiplies. In Figure 5 we show the performance approaching the peak performance of the JLab KNL machines. With these recent improvements, the computations in two and three-body systems are quite tractable.

B. Computational Cost estimates for GPUs

We will make use of the QUDA Multigrid to compute the light and strange annihilation quark-lines on 64 time-slices of the $48^3 \times 512$ lattices with $N = 640$ distillation vectors. The choice of the number of vectors is based on the need to resolve mesons at high momentum. More vectors improves the operator overlap onto the state at momentum expected to be important, namely $|\vec{p}|^2 = 6$ in lattice units. Early tests with $N = 256$ were not sufficient to achieve the desired precision. In addition, the derivative-based single-particle operator constructions are more effective with a larger number of vectors. On $32^3 \times 256$ lattices, we used $N = 256$ vectors.

We will use 32 time-slices for each time-source, this means we will have two time-sources of statistics to average when we eventually compute correlation functions. We want to compute these annihilation lines on the 100 existing configurations, but as it will be clear, there is not enough GPU time at JLab. We have accurate timings of the AMG inverter on the Summit V100 system, and now also have preliminary timings on the JLab RTX 2080 system for both light and strange quark masses. We note that at these small pion masses, the number of iterations of the inverter is only weakly dependent on the quark mass.

One useful tool at our disposal is the geometry of how we divide our lattice layout over the GPUs. As we have lattices with a very long time extent, we can choose to have all the parallelism in the time-direction. This configuration is well matched to the communication network on the 8-node GPU boxes. Performance varies (considerably) with the number of GPUs and how many dimensions in which they communicate. The "only-time" parallelism is the optimal configuration.

For the proposed $48^3 \times 512$ lattices, we find the timings are 3.34 sec for the light quark inversion on 128 GPUs, and 1.55 sec for the strange quark. We will use $N = 640$ distillation source vectors and use this source vector for each Dirac component of the source, thus we are performing 2560 inversions per time-slice for each annihilation quark line. For 64 time-slices, the total time for the light quarks is 2.0M RTX hours and 0.9M for the strange quarks and on 100 configurations. This is clearly more than the available time at JLab. We will scale the number of configurations we can use to fit the available machine time. Thus, 70 configurations is feasible including using all of the new RTX system with a total cost of **2.0M** RTX2080 gpu-hours, or equivalently, 2.0M K80 gpu-hours. Below 50 configurations, the statistics are quite limited, although we have two time-sources which effectively doubles our statistics.

We comment on the choice of the number of time-slices. For isoscalar observables, we have Wick contractions that feature the product of traces (disconnected bubbles) of the meson operator. On all of our other ensembles, we have all the annihilation lines available. We increase statistics by averaging the product of the traces around the lattice. We then use these (time-averaged) Wick contractions within a correlation function that may feature connected graphs over, say, 32 time-slices. Our choice to use 64 time-slices on the large lattice is a pragmatic one, allowing for some averaging of disconnected graphs.

C. Computational Cost estimates for KNLs

A scientific goal we will pursue this year is the calculation of correlation functions that allow us to extract resonance properties of low lying mesons. In particular, we find the costs largely driven by correlation functions featuring three-body decays. We use the variational method to evaluate a matrix of correlation functions featuring a large basis of operators including single and multi-meson operators. After Wick contractions, these correlation functions will require the evaluation of a large number of graphs. In general, the cost of evaluating correlation function is determined by the number of intermediate nodes that are needed to evaluate the graphs. As discussed above, a key part of our Exascale development has been developing algorithms to evaluate a large number of graphs with the minimum number of such intermediate (temporary) nodes.

We will compute correlation functions over 40 time-slices. For the $32^3 \times 256$ lattices at $m_\pi = 236$ MeV, we will average over 4 time-sources to increase statistical accuracy over 485 configurations. For the $24^3 \times 256$ lattices at $m_\pi = 275$ and 325 MeV, we will average of 4 time-sources and use 400 and 312 configurations, resp. We have found this number of time-sources to give us resonable statistics in previous studies.

The basic contraction of two nodes involves the contraction over both spin as well as distillation indices. The basic operation is of the form $\sum_{\beta=1}^{4}\sum_{j=1}^{N} M^{ij}_{\alpha\beta}P^{jk}_{\beta\gamma}$, for some number of distillation indices. Based on our NESAP and Exascale development effort, we have optimized the Intel Knights Landing (KNL) implementation to have a performance twice that of the Nvidia K20 GPU implementation for the largest matrix size of $N = 256$. As the new JLab GPU system will feature gamer cards without ECC, will not consider using that system for our contraction work. For the current proposal, we will use the timing of *0.00214 secs* per hadron multiply for $N = 256$.

We will carry out contractions in the isospin= 1, 2 and 3 channels with energies above $\pi\pi\pi$ and $\bar{K}K\pi$ thresholds. Again, we will consider multi-particle operator constructions using projected operators for the three light quark mass ensembles with pion masses at 236, 275 and 325 MeV, resp. The smallest pion mass calculations will use $N = 256$ distillation vectors while the others will use $N = 162$ vectors. We note that in some sectors at light pion masses, we must consider projected operator constructions involving the η and η' up to $|\vec{p}|^2 = 6$. These constructions are all in place.

As an example, considering $\pi\pi\pi$ and $K\bar{K}\pi$ in isospin = 1 and $\Lambda^{PG} = T_1^{+-}$ where the a_1 resonance can be found, we find that up to $E_{cm} = 1580$ MeV on the $L = 32$ lattice with $m_\pi = 236$ MeV, there are 196443K graphs that must be evaluated in a basis of 37 operators. There are 40×216463 intermediate nodes that must be computed to evaluate the graphs for 40 time-slices.

We find that considering total momentum up to $|\vec{p}|^2 = 4$ (and relative momentum up to $|\vec{p}_{rel}|^2 = 6$) for all the possible helicities and three isospins, we arrive at 21.6M number of temporaries per configuration with 4 time-sources, on the lightest pion mass ensemble ($m_\pi = 236$ MeV). For 485 configurations, this requirement is 73**M KNL core-hours** using the conversion of 0.563 KNL / Sky core.

For the $24^3 \times 256$ ensembles at the heavier pion masses, there are less two-particle operators with large relative momentum needed to reach the target center-of-mass energies. With total momentum of the system up to $|\vec{p}|^2 = 4$ (and relative momentum up to $|\vec{p}_{rel}|^2 = 4$), we need 8.6M temporaries for 4 time-sources per configuration on the 275 and 325 MeV ensembles. Having a lower internal momentum for the two particle constructions significantly reduces the number of temporaries. For each of the datasets, this requirement is **29.4 KNL core-hours**. Thus, the total cost of all the contractions is **102.4M KNL core-hours**.

V. STORAGE

For the contractions, we must hold on disk the light and strange quark "annihilation" propagators (perambulators), the "forward" propagators, and the meson 'elementals' which are used to construct the meson operators.

For 485 configurations of the $32^3 \times 256$ lattices, the propagator and meson elemental objects are 1026 GB in size per configuration. Thus, we need about 500 TB of hot disk-space for this ensemble. No new tape space is needed for the proposed contractions.

We do request new tape storage at JLab for data produced under an LCF calculation that will be needed for a future proposal. We estimate 100 TB of new tape will be needed for these anticipated calculations.

Thus, our request for new tape storage is **100 TBytes**, equivalent to **0.13M Sky-equivalent core-hours**. The request for disk storage is **500 TBytes** which is equivalent to **3.75M Sky-equivalent core-hours**.

VI. READINESS AND RUN SCHEDULE

As described in detail in Sec. IV A, the propagator generating codes in *Chroma* and *QUDA*, and the *Redstar* contraction codes are in production now. They are well supported, and in use by several different project members at JLab, as well as other sites outside of USQCD. Regarding the anticipated run schedule, we have/can/will run continually throughout the year. The codes have run efficiently enough, in fact, that we have nearly burned up our allocation on the KNL-s for this year. All proposed work is "comfortably" parallel, relying on existing configurations and input data. All the jobs can run independently, as long as the input data is present. Thus, we rely on a high disk footprint to maintain high throughput in the queues. Our scripts are designed so that we can submit hundreds of jobs and let them filter through the queues according to the fairshare. This method has proven to be highly effective at JLab. Our anticipated work schedule is such that we are computing bound - waiting on the current data-set to finish before launching the next data-set.

VII. DATA SHARING & EXCLUSIVITY

We have are in the process of generating *DOI-s'* (Document Object Indicators) for our gauge ensembles and propagators. Users will need an account at JLab to access the data.

The calculation of the hadron resonance spectrum for states composed of the u, d, s, c quarks using the proposed data sets is an exclusive part of this proposal.

---

[1] R. A. Briceño, M. T. Hansen, and S. R. Sharpe, Phys. Rev. **D99**, 014516 (2019), 1810.01429.

[2] J. J. Dudek, R. G. Edwards, P. Guo, and C. E. Thomas (Hadron Spectrum), Phys.Rev. **D88**, 094505 (2013), 1309.2608.

[3] J. J. Dudek, Phys.Rev. **D84**, 074023 (2011), 1106.5515.

[4] J. J. Dudek, R. G. Edwards, B. Joo, M. J. Peardon, D. G. Richards, et al., Phys.Rev. **D83**, 111502 (2011), 1102.4299.

[5] J. J. Dudek, R. G. Edwards, and C. E. Thomas (Hadron Spectrum Collaboration), Phys.Rev. **D87**, 034505 (2013), 1212.0830.

[6] D. J. Wilson, R. A. Briceno, J. J. Dudek, R. G. Edwards, and C. E. Thomas, Phys. Rev. **D92**, 094502 (2015), 1507.02599.

[7] J. J. Dudek, R. G. Edwards, and D. J. Wilson (Hadron Spectrum), Phys. Rev. **D93**, 094506 (2016), 1602.05122.

[8] R. A. Briceno, J. J. Dudek, R. G. Edwards, and D. J. Wilson (2017), 1708.06667.

[9] M. Luscher and U. Wolff, Nucl. Phys. **B339**, 222 (1990).

[10] M. Luscher, Nucl. Phys. **B364**, 237 (1991).

[11] K. Rummukainen and S. A. Gottlieb, Nucl.Phys. **B450**, 397 (1995), hep-lat/9503028.

[12] C. Kim, C. Sachrajda, and S. R. Sharpe, Nucl.Phys. **B727**, 218 (2005), hep-lat/0507006.

[13] N. H. Christ, C. Kim, and T. Yamazaki, Phys.Rev. **D72**, 114506 (2005), hep-lat/0507009.

[14] R. A. Briceno, Phys. Rev. **D89**, 074507 (2014), 1401.3312.

[15] R. A. Briceño, J. J. Dudek, R. G. Edwards, and D. J. Wilson, Phys. Rev. Lett. **118**, 022002 (2017), 1607.05900.

[16] S. He, X. Feng, and C. Liu, JHEP **0507**, 011 (2005), hep-lat/0504019.

[17] R. A. Briceno and Z. Davoudi (2012), 1204.1110.

[18] M. T. Hansen and S. R. Sharpe, Phys.Rev. **D86**, 016007 (2012), 1204.0826.

[19] P. Guo, J. Dudek, R. Edwards, and A. P. Szczepaniak, Phys.Rev. **D88**, 014501 (2013), 1211.0929.

[20] D. J. Wilson, J. J. Dudek, R. G. Edwards, and C. E. Thomas, Phys. Rev. **D91**, 054008 (2015), 1411.2004.

[21] J. J. Dudek, R. G. Edwards, C. E. Thomas, and D. J. Wilson (Hadron Spectrum), Phys. Rev. Lett. **113**, 182001 (2014), 1406.4158.

[22] A. J. Woss, C. E. Thomas, J. J. Dudek, R. G. Edwards, and D. J. Wilson (2018), 1802.05580.

[23] G. K. C. Cheung, C. E. Thomas, J. J. Dudek, and R. G. Edwards (Hadron Spectrum), JHEP **11**, 033 (2017), 1709.01417.

[24] M. A. Clark et al., Comput. Phys. Commun. **181**, 1517 (2010), 0911.3191.

[25] R. Babich, M. A. Clark, and B. Joo, ACM/IEEE Int. Conf. High Performance Computing, Networking, Storage and Analysis, New Orleans (2010), 1011.0024.

[26] M. Clark and R. Babich, *QUDA: A library for QCD on GPUs*, `http://lattice.github.io/quda/`.

[27] M. A. Clark, R. Brower, and M. Cheng, Proceedings of the 2014 GPU Technology Conference (2014).

[28] C. E. Thomas, R. G. Edwards, and J. J. Dudek, Phys. Rev. **D85**, 014507 (2012), 1107.1930.

[29] L. Liu et al., JHEP **1207**, 126 (2012), 1204.5425.