Project Description

1 Introduction

Funds are requested from the National Science Foundation's Major Research Instrumentation (MRI) program to develop a supercomputing cluster to support the research programs at the University of Richmond in nuclear physics and astrophysics. The instrument will also be available to senior personnel at other institutions for work on closely related projects. The research groups at Richmond typically support 5-6 undergraduates each summer and during the academic year (the University of Richmond is a primarily undergraduate institution). These students routinely go on to careers in science and engineering. Both of the Richmond programs have external support from the US Department of Energy (DOE) (Gilfoyle in nuclear physics) and the National Science Foundation (NSF) (Bunn in astrophysics).

The nuclear physics research is centered on unraveling the structure of the nucleon and nature of quark confinement. Additional senior personnel in nuclear physics will use the instrument at Ohio University, Virginia Tech, and Union College. They are members of the CLAS Collaboration (with Gilfoyle) that is responsible for the operation of a large particle detector (CLAS) at the Thomas Jefferson National Accelerator Facility (JLab) in Newport News, VA. The astrophysics research will focus on simulations of a new cosmic microwave background polarimeter being built by researchers at Brown University and the University of Wisconsin in collaboration with Co-PI Bunn. Senior personnel at both Brown and Wisconsin will be heavily involved in the project.

We have considerable experience with supercomputing. One of us (Gilfoyle) was a co-principal investigator on a project that developed a computing cluster in 2001 with support from the NSF at Richmond. He has been the manager of that project since then. That existing system is near the end of its useful life, but the infrastructure to support it is still sound and the proposed system will benefit from that investment.

2 Nuclear Physics

The research effort in nuclear physics is part of the program at JLab. The primary mission of JLab is to reveal the quark and gluon structure of nucleons and nuclei and to deepen our understanding of matter and the confinement of quarks. Nucleons are comprised of three valence quarks held together by the exchange of gluons, along with a quark-gluon sea of particles. Quantum Chromodynamics (QCD), the fundamental theory of particle physics, is a highly successful description of quarks at very high 4-momentum transfers or Q^2 [1], but at the energies where the nucleons exist (the non-perturbative region), it has proved to be a daunting challenge to solve [2]. At low $Q^2 < 0.5$ $(\mathrm{GeV/c})^2$ the "hadronic" picture of nuclei (*i.e.*, nuclei made of protons and neutrons) has been successful at reproducing a wide range of measurements [8]. However, the transition region between these extremes is poorly understood in the range $Q^2 \approx 0.5 - 10.0$ $(\mathrm{GeV/c})^2$ and mapping the geography of this transition is an essential goal of nuclear physics (see Long-Range Plan of the Nuclear Science Advisory Committee (NSAC)[1] [3]).

The central instrument at JLab is a superconducting electron accelerator with a maximum energy of 4-6 GeV, a 100% duty cycle, and a maximum current of 200 μA. Our research is done

[1]NSAC is an advisory committee that provides official advice to the Department of Energy (DOE) and the National Science Foundation (NSF) on the national program for basic nuclear science research.

in Hall B with the CEBAF Large Acceptance Spectrometer (CLAS). This device is a large (45-ton), toroidal multi-gap magnetic spectrometer with nearly full solid angle coverage. A toroidal magnetic field is generated by six iron-free superconducting coils. The particle detection system consists of drift chambers [4], Cerenkov detectors [5], scintillation counters [6] for time-of-flight measurements, and electromagnetic calorimeters [7]. Together there are about 33,000 detecting elements capable of acquiring approximately 1 terabyte of data per day. The Richmond group has been part of the CLAS Collaboration which built and operates the detector since its inception.

The analysis of these large data sets requires significant computing resources. First pass analysis is done on the JLab computing farm, but final results require additional analysis. Demand is high for the computing resources at JLab, and it can routinely take a day for a submitted batch job to start. We have developed our own local computing cluster so we can analyze our data in a timely fashion. We also simulate the response of CLAS to separate real physics effects from artifacts of the detector. This stage requires large disk space to store the Monte Carlo events and, more importantly, considerable computing power. Our simulation generates Monte Carlo events at only about 2-3 events per second.

2.1 Out-of-Plane Structure Functions of the Deuteron

The hadronic model of nuclear physics has been successful at low Q^2, but it is not well-developed in the GeV region. There have been few measurements to challenge theory in this region until recently [8, 9]. We must provide a baseline for the hadronic model so deviations at higher Q^2 can be attributed to quark-gluon effects with greater confidence [3, 10]. To this end, we are investigating the out-of-plane structure functions of the deuteron, the simplest nucleus, using the reaction $D(\vec{e}, e'p)n$ with CLAS. The cross section for the reaction is

$$\frac{d\sigma^5}{d\omega d\Omega_e d\Omega_{pq}} = C\left(\rho_l f_l + \rho_t f_t + \rho_{TT} f_{TT} \cos\phi_{pq} + \rho_{LT} f_{LT} \cos 2\phi_{pq} + h\rho'_{LT} f'_{LT} \sin\phi_{pq}\right) \tag{1}$$

where C and the ρ_i are functions of the known electron parameters, h is the helicity of the electron beam, and ϕ_{pq} is the azimuthal angle of the ejected proton relative to the 3-momentum transfer $\vec{q}$ and the 3-momentum of the beam. The unique, nearly-4π solid angle of CLAS creates an inviting opportunity to study the ϕ_{pq}-dependent structure functions f'_{LT}, f_{LT}, and f_{TT}. They represent a model-independent measurement of a little-studied part of the deuteron cross section.

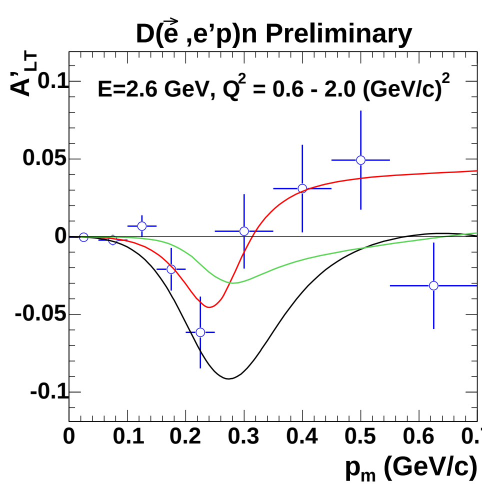

Figure 1: Preliminary results for the asymmetry A'_{LT} proportional to f'_{LT}.

These three structure functions are extracted by measuring different moments of the out-of-plane production in CLAS. Each of these moments is related to a different asymmetry which is, in turn, proportional to a particular structure function. The data cover the 4-momentum transfer range $Q^2 = 0.2 - 5.0$ (GeV/c)^2. We are studying the reaction in quasi-elastic kinematics first and later will investigate higher energy transfers. Our preliminary results for the asymmetry A'_{LT} (proportional to f'_{LT}) in Figure 1 at a beam energy of 2.6 GeV show significant structure which is reproduced by a calculation from Jeschonnek (red curve) [11]. Two other calculations by Arenhoevel (black curve) [12] and Laget (green curve) [13] do not reproduce the data as well. The extraction and analysis of the other two structure functions (f_{LT} and f_{TT}) and investigations of different kinematic regimes for all three structure functions are ongoing. This work is part of a CLAS Approved Analysis (CAA) entitled "Out-of-Plane Measurements of the Structure Functions of the Deuteron." A CLAS Collaboration member can write a proposal to analyze an existing data set, the proposal is reviewed by a committee of Collaboration members, and defended before the full Collaboration who then vote to approve the project. Gilfoyle is the spokesperson on this CAA [14]. The analysis of the data and the simulations of CLAS are computationally intensive and would use the proposed cluster.

2.2 Magnetic Form Factor of the Neutron

The elastic electromagnetic form factors are the most basic observables that describe the internal structure of the proton and neutron. The differential cross section for elastic electron-nucleon scattering can then be calculated in the laboratory frame as [15]

$$\frac{d\sigma}{d\Omega} = \sigma_{Mott}\left(G_E^2 + \frac{\tau}{\epsilon}G_M^2\right)\left(\frac{1}{1+\tau}\right) \tag{2}$$

where σ_{Mott} is the cross section for scattering from a point particle, G_E is the electric form factor, G_M is the magnetic form factor, $\tau = Q^2/4M^2$, ϵ is the polarization of the virtual photon, and M is the nucleon mass. There are four elastic form factors (electric and magnetic ones for each nucleon) and their evolution with Q^2 characterizes the distributions of charge and magnetization within the proton and neutron. They are stringent tests of non-perturbative QCD and are connected to generalized parton distributions (GPDs). Conventional parton distributions describe the longitudinal momenta of the nucleon constituents, but integrate over the transverse structure losing, for example, information about the orbital angular momentum of the partons. With exclusive measurements and the GPD formalism one can determine the longitudinal momenta and transverse position of the partons inside the nucleon, their orbital angular momentum, and quantum interference effects [16]. The elastic form factors are also important challenges for lattice QCD to meet. Lattice QCD is one of the more promising avenues for solving non-perturbative QCD, and one of its important tests will be the accuracy with which it can reproduce the elastic form factors in this Q^2 range [17].

We are part of a broad assault on the four elastic nucleon form factors at JLab [18]. Our focus is on G_M^n, the magnetic form factor of the neutron. To measure G_M^n we use the ratio of elastic $e-n$ to elastic $e-p$ scattering on deuterium. The ratio method is less vulnerable to uncertainties than previous methods. We have completed data collection and most of the analysis for a measurement of G_M^n in the range $\mathrm{Q}^2 = 0.2-5.0$ $(\mathrm{GeV/c})^2$ [19, 20]. A report describing the analysis of two out of the three sets of running conditions in the experiment is under Collaboration review. Our group at Richmond has taken on the analysis of the third remaining data set. We note the G_M^n data sets are the same ones used in the deuteron structure function analysis described in Section 2.1. We have also submitted a Letter-of-Intent to the JLab Program Advisory Committee (PAC) to make the same measurements at higher Q^2 as part of the JLab 12-GeV Upgrade.[2] A Letter-of-Intent is a preliminary proposal for beam time at JLab, and it provides feedback from the PAC before the researchers make the large investment of time and effort required to produce a full proposal. The Letter-of-Intent was approved by the PAC in August 2006 [22]. We are now developing a full beam time proposal. Gilfoyle is the spokesperson and contact person for the Letter-of-Intent. The work of completing the existing experiment analysis and preparing the full proposal will take advantage of the computing cluster proposed here.

[2] The DOE plans to upgrade the accelerator at JLab from a beam energy of 6 GeV to 12 GeV. The upgrade will require extensive changes to the accelerator and to CLAS to take advantage of the new physics opportunities. The JLab, 12-GeV-Upgrade is the fourth highest priority of the DOE office of Science in the next 20 years [21].

2.3 Quark Propagation and Hadron Formation

The confinement of quarks inside hadrons is perhaps the most remarkable feature of QCD and solving its mysteries is an essential goal of nuclear physics [3]. In one picture a struck quark stretches the color string until $q\bar{q}$ pairs tunnel up from the vacuum and break the string. The full picture with full QCD is more complicated.

We have proposed a broad program of measurements and analyses to determine the mechanisms of confinement in forming systems. We use the nucleus as a ''detector" to probe the hadronization formation length and the time scale on which a pre-hadron (such as a bare $q\bar{q}$ pair) becomes dressed with its own gluonic field. The response of the hadron to the presence of the nucleus depends on the time scale on which hadronization takes place inside the nucleus. The ratio of hadrons produced relative to the production from deuterium (hadronic multiplicity ratio R^h_M) and transverse momentum broadening Δp_T^2 are the two primary observables. Using a wide range of nuclear targets one can measure the quark production time and hadron formation times with different hadrons produced in the reaction. The production time is the lifetime of a deconfined quark, and it will be determined by analyzing the kinematic dependence of Δp_T^2. The formation times are the time intervals required to form the color field of hadrons, and these will be determined from the kinematic dependence of the hadronic multiplicity ratio R^h_M. A proposal to do this experiment at high Q^2 as part of the 12-GeV Upgrade was approved by the JLab PAC August, 2006 [23]. Gilfoyle is a co-spokesperson on that proposal and will be responsible for analyzing the π^0, η, and η' production. We will use the proposed instrument to simulate the physics and the upgraded CLAS detector response to prepare for the 12-GeV Upgrade.

2.4 Technical Projects

The measurements of the nuclear reactions described above are subject to radiative corrections. The code DEEP_EXCLURAD developed by one of us (Gilfoyle) can be used for exclusive reactions using electrons [24]. Radiative corrections are usually calculated using the formalism developed by Schwinger and Mo and Tsai [25, 26], but that method is valid only for inclusive electron scattering and not for exclusive ones studied here. These new calculations are based on the method developed by Afanasev, *et al.* [27] and use the program DEEP to calculate the deuteron response functions [28]. We use an adaptive step size to perform some of the integrals so the run time can vary from a few tens of seconds to several hours for a single kinematic point. We need hundreds of such points for our analysis. This makes it essential to have access to the supercomputing cluster.

We are also committed to development projects for the JLab 12-GeV Upgrade to double the beam energy of the electron accelerator and enhance the experimental equipment in Hall B [29]. We will be responsible for design, prototyping, development, and testing of software for event simulation and reconstruction. The improved CLAS detector (called CLAS12) will have prodigious software requirements. Event simulation is an essential aspect of the design of CLAS12 and the eventual precision of the detector. For many experiments, the quality of the results will be limited by systematic uncertainties instead of statistical ones. The work will make significant use of the proposed cluster during its lifetime.

2.5 Role of Senior Personnel in Nuclear Physics

Faculty from institutions besides the University of Richmond are part of the nuclear physics portion of this project. All are members of the CLAS Collaboration with Gilfoyle and have been users of the existing cluster. Their participation here will raise their own research productivity and make better use of the proposed instrument. Just as important, the number of students, undergraduate and graduate, that will learn sophisticated data analysis and data mining methods will increase. Below we describe the research of the other senior personnel in nuclear physics.

Dr. M.F. Vineyard is a professor at Union College in Schenectady, NY. Union is a primarily undergraduate institution like Richmond. He is working with Gilfoyle on the analysis of the G_M^n data (see Section 2.2) and studying the effects of the nuclear medium on excited states of the nucleon using photoproduction of the η meson as a more selective probe for some resonances. Dr. Vineyard is a co-spokesperson on those two measurements and on a Letter-of-Intent approved in 2006 to extend the G_M^n measurements to higher Q^2[19, 22, 30]. Dr. Vineyard is also involved in technical projects developing software for the JLab 12-GeV Upgrade.

Dr. K.H. Hicks is a professor at Ohio University in Athens, OH. Ohio University is a research institution and Dr. Hicks typically has a postdoctoral fellow and 1-3 doctoral students in his group. His research is focused on strangeness electro- and photo-production as a way to probe the structure of the kaon itself and as a means to identify new hadronic states and to unravel the structure of known ones. He is co-spokesperson on three approved proposals at JLab including one with Gilfoyle in 2006 to study quark propagation and hadron formation as part of the JLab 12-GeV Upgrade (see Section 2.3) [23, 31, 32]. He is also the spokesperson on three CLAS Approved Analyses to investigate similar physics.

Dr. D. Jenkins is an emeritus professor at Virginia Tech in Blacksburg, VA. Virginia Tech is a research institution. Dr. Jenkins is a spokesperson on JLab experiment "The Photoproduction of Pions." Analysis is now being completed for this experiment. He is part of another experiment to measure pion, eta and kaon photoproduction from a polarized target scheduled to begin March 2007. He is also a member of a group searching for exotic hybrids through coherent production on helium in preparation for the 12-GeV upgrade at JLab.

3 Astrophysics

The astrophysics research we propose concerns analysis of observations of the cosmic microwave background (CMB) radiation. For over ten years now, CMB observations have been among the main contributors to the extraordinary advances in precision cosmology. The maps made by the COBE [33] and WMAP [34] satellites have led the way, along with a large number of suborbital experiments. With continued advances in observing technology in recent years, CMB observations are expected to be in the forefront of further great advances in cosmology in the coming decades. In particular, we are now at the beginning of an era of CMB polarimetry.

CMB polarization has already been detected [35, 36, 37, 38, 39, 40]. Numerous instruments are currently attempting to further characterize the polarization of the CMB, and more are under development. In the near future, maps of CMB polarization are expected to refine estimates of cosmological parameters (e.g., [41]), probe the ionization history of the Universe [42] and the details of recombination [43], and measure gravitational lensing due to large-scale structure [44]. Most exciting of all, polarization maps may provide a direct probe of an inflationary epoch in the extremely early Universe [45, 46].

CMB polarimetry is an extremely high priority in the astrophysics community. As a joint DoE/NASA/NSF Task Force noted, "the accurate measurement of CMB polarization is the next critical step in extending our knowledge of both the early Universe and fundamental physics at the highest energies ... As our highest priority, we recommend a phased program to measure the large-scale CMB polarization signal expected from inflation" [47]. Similar opinions are expressed in the National Research Council's decadal survey of astronomy and astrophysics [48] and their report *Connecting Quarks with the Cosmos* [49].

NASA's *Beyond Einstein* road map for future astrophysics programs includes a dedicated CMB polarimeter known as the Einstein Inflation Probe (EIP) [50]. Co-PI Bunn is one of the leaders, along with Peter Timbie at Wisconsin and Gregory Tucker at Brown, of a NASA-funded Mission

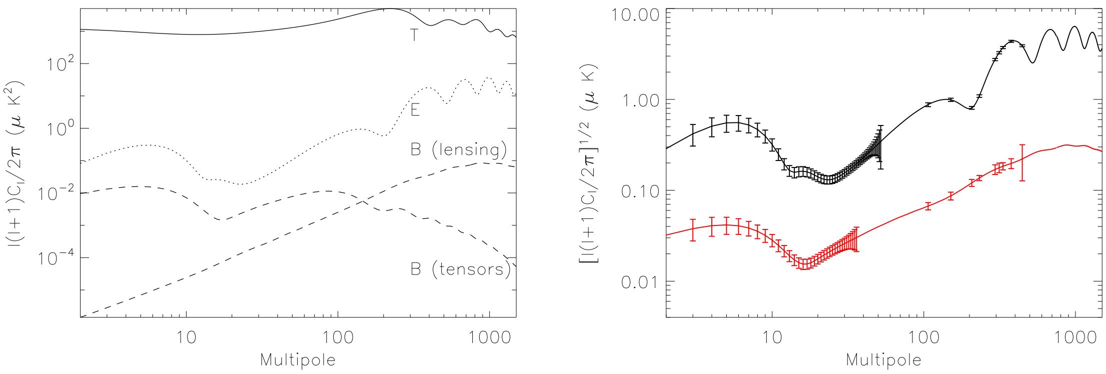

Figure 2: The left panel shows CMB angular power spectra for temperature anisotropy, E polarization, and B polarization. For the B spectrum, the tensor and lensing contributions are shown separately. These spectra are based on the best-fit model from WMAP [34], with the tensor-to-scalar ratio taken to be 0.1. The right panel shows forecasted errors on the E (upper) and B (lower) signals for EPIC [51]. In addition to characterizing the E spectrum with unprecedented precision, EPIC will be able to measure both of the major B signals: the lensing contribution at high multipoles (small angular scales) and the tensor contribution at low multipoles. The latter provides a unique test of inflation.

Concept Study for the EIP, the Einstein Polarization Interferometer for Cosmology (EPIC). If eventually selected, this instrument would be an interferometer using bolometric detectors, the highest-sensitivity detectors at millimeter wavelengths. The members of the EPIC collaboration are currently constructing and deploying a ground-based prototype four-element millimeter-wave bolometric interferometer (MBI-4). We plan in the next few years to extend this to a 16-element balloon-borne instrument (MBI-16). Recent status reports on MBI and EPIC may be found in [51, 52]. The work proposed herein is to develop data analysis and simulation tools in support of MBI and EPIC. Although we will focus primarily on the MBI/EPIC instrument design, we expect our results to be applicable to CMB interferometers more generally. Timbie and Tucker are senior personnel on this proposal and will be heavily involved in this research.

The key to understanding the science that can be derived from a CMB polarization map is the fact that any such map contains two components — a scalar E component and a pseudoscalar B component [45, 46], which probe different physical phenomena. In particular, ordinary density (scalar) perturbations produce only E-type polarization (to linear order). As a result, the B component is predicted to be smaller than E by an order of magnitude or more over all angular scales (see Figure 2). However, the very fact that density perturbations do not produce B-type polarization makes detection of the B component more valuable: the B channel is a clean probe of other types of perturbations. By far the most exciting prospect is the use of B modes to detect primordial gravitational waves (tensor perturbations) produced during an inflationary epoch. If this tensor B component is detected, we will have a direct probe of the Universe at far earlier times than any other method can provide.

Other than the signature of primordial tensor perturbations, the dominant source of B-type polarization in the CMB is expected to be gravitational lensing of E modes by large-scale structure. These two predicted sources of B modes probe very different epochs: the tensor contribution is imprinted on the CMB at the time of last scattering but is a relic of the extremely early Universe; the lensing contribution is produced at much later times.

All CMB polarization detections to date have been of the E component; the chief goal of the EIP, as well as future suborbital missions, will be to detect the B component. The B-type polarization signal is extremely weak; however, by combining thousands of low-noise detectors with long integration times, the next generation of experiments will have the raw sensitivity to detect it. Figure 2 shows forecasts of E and B error bars for the EPIC mission concept. The greater challenges are removal of non-cosmological foreground signals and control of systematic errors such as pointing errors, imperfectly known beam patterns, crosstalk among detectors, etc. The primary science goal of the astrophysics research proposed herein is to perform simulations of interferometric CMB polarization observations to assess the effects of systematic errors.

In Section 3.1, we outline some differences between interferometric CMB observations and traditional imaging systems, and in Section 3.2 we describe the simulations to be performed.

3.1 Interferometric CMB polarimetry

Traditionally, measurements of the CMB have used two different approaches: direct imaging and interferometry, and both technologies are being seriously examined as candidates for future B-mode experiments such as the EIP. The goal of MBI and EPIC is to combine the advantages of interferometry with bolometric detectors, the lowest-noise detector technology at millimeter wavelengths. There are several advantages to this approach:

Cost and Simplicity. For measurements that require high angular resolution, large single dishes are often impractical for a number of reasons including mass, deformations due to gravity and cost. Interferometers effectively enable high angular resolution by reproducing the resolution performance of a large dish; the trade-off is a reduction in collection area if the interferometer area is not filled.

Clean Optics. Interferometers have simple, reflection-free optics, removing various sources of spurious polarization. The microwave signal enters the instrument via corrugated horn antennas, which have extremely low sidelobes and easily calculable, symmetric beam patterns.

No chopping and scanning. Traditional imaging systems for CMB observations typically use some form of chopping, either by nutating a secondary mirror or by steering the entire instrument at a rate faster than the $1/f$ noise in the atmosphere and detectors. With no need for rapid chopping, the time constants of the detectors can be relatively long.

Better Resolution for Equivalent Size. An interferometer has angular resolution roughly twice as good as a monolithic dish of the same size. The reason is that the signal in a filled dish is dominated by spacings that are much smaller than the aperture diameter. This angular resolution factor is important because the size of the aperture is a cost-driver for the EIP. Angular resolution is important for CMB polarization measurements in two ways. First, imperfections in the shape and pointing of beams couple the CMB temperature anisotropy into false polarization signals. These problems can be reduced significantly if the CMB is smooth on the scale of the beam size, which happens for beams smaller than $\sim 10'$ [53]. Second, removing contamination of the tensor B-mode signal by B-modes from weak lensing requires maps of the lensing at higher angular resolution than the scale at which the tensor B-modes peak [54].

Clean separation of E and B modes. In any incomplete sky map, there is some “leakage” between E and B components [55, 56, 57]. One simple way to understand this is to note that the E-B separation can be done trivially mode by mode in Fourier space. With incomplete sky coverage, individual Fourier modes cannot be measured. The data from an interferometer consists of visibilities, which have narrow window functions in Fourier space. As a result, there is less E-B leakage in interferometric data than in imaging data [58].

Direct Measurement of Stokes Parameters. Interferometers can measure the linear polarization Stokes parameters Q and U directly, without differencing signals from different detectors. As a result, some sources of systematic error that cause large spurious signals in a traditional experiment are greatly mitigated. For example, effects such as differential pointing errors (“squint”) cause the unpolarized temperature anisotropy to contaminate the E and B polarization channels in a traditional imaging experiment [53] but do not in an interferometer [59]. Since the unpolarized signal is orders of magnitude larger than the polarized signal (Fig. 1), this makes a big difference.

3.2 Systematic error simulations

It is clearly essential to have a detailed, quantitative understanding of the effects of systematic errors and foreground contamination on data from both interferometric and imaging systems. The state of the art is far more developed for imaging systems than for interferometers. Our goal is to close that gap, so that the two technologies can be compared on an equal footing. We have established a theoretical framework for analysis of systematic errors in CMB interferometric polarimetry [59], but this work needs to be supplemented with detailed simulations. We propose to perform such simulations of CMB interferometric polarimetry in order to assess the effects of various systematic errors on MBI/EPIC in particular and on CMB interferometer in general.

Along with the construction of the prototype interferometer MBI-4, the MBI/EPIC research group has been developing a data analysis pipeline. Many steps in the pipeline can be performed with adaptations of publicly-available parallelized code (e.g., [60]); some stages require home-grown code which is currently under development. We will adapt this pipeline to the planned MBI-16 instrument, but simulation of this larger instrument will require more computing power than we presently have. The equipment to be purchased under this grant will enable us to simulate analysis of the MBI-16 data. Our overall goal is to simulate the propagation of a known signal through the instrument and then analyze it in the same manner as we will with the real data. We will determine in precise detail the error properties of both the recovered Fourier-space power spectrum and the recovered image.

We now outline the key steps in the simulation of MBI data. We will assess the computational requirements in section 4.3.

Simulation of time-ordered data (TOD). Given an underlying "true" sky map, a model of the instrument as well as its attitude as a function of time, and a noise model, we need to compute the simulated output time streams from each of the detectors. This can be done with a scattering-matrix model for each of the instrument components. This step is not computationally intensive, and code is already largely developed.

TOD → Visibility-space "map." The raw data from an interferometer is a set of visibilities, which are essentially samples of the Fourier transform of the map, convolved with the primary beam.[3] Because the data are contaminated by correlated noise, the optimal recovery of a Fourier-domain visibility map from the TOD is nontrivial; however, efficient parallelized algorithms such as MADMap [60], originally developed for traditional imaging systems, can be adapted for this purpose.

Power Spectrum Estimation. We wish to determine the maximum-likelihood power spectrum for a given visibility data vector. Once again, standard codes for imaging systems, which have been parallelized and made publicly available, can be adapted to apply to visibility data.

Visibility data → Image. The primary science goal of a CMB experiment is the power spectrum, which can be computed entirely in the visibility domain, without ever constructing a real-space image of the observed map. However, in order to check for errors or foreground contamination in the data, we will surely want to produce actual images from the visibility data. In addition, some CMB studies search for signals beyond merely the power spectrum and so

[3]In principle we should work with spherical harmonic transforms rather than Fourier transforms, but for the relatively small fields of view considered by MBI Fourier transforms are adequate, even when many fields of view are mosaicked together [61].

require real images.

In principle, one can go from visibility space to a real-space image simply by performing an inverse Fourier transform. This is of course computationally trivial and is likely to be useful for quick-look diagnoses of the data. However, because the Fourier plane is not generally fully sampled, and because of the complicated noise properties of the data, the images recovered in this way are unlikely to be adequate for a final analysis.

In traditional radio astronomy, the most common image recovery technique is the CLEAN algorithm [62]. Unfortunately, this algorithm is well-suited to sources with sharp features, not to the diffuse nearly-Gaussian structure of CMB maps. We plan instead to use maximum-entropy image reconstruction, which has been well-developed in the CMB context [63]. Note that maximum-entropy is a nonlinear method, so the error properties of the resulting maps may be non-Gaussian and complicated. It is therefore extremely important to simulate this step in the analysis.

3.3 Component separation

The simulations described in the preceding section will be our primary initial focus. Over a longer time scale, we plan to develop code to test other aspects of the MBI data analysis and to address other problems in CMB data analysis. Chief among these is the problem of component separation.

As MBI-16 and eventually EPIC attempt to characterize B polarization, the issue of component separation (i.e., removal of foregrounds) will be crucial. Both blind techniques (e.g., independent component analysis) and those based on fitting to foreground templates have been proposed for CMB component separation, but few have been adapted to the case of interferometric data. An extremely interesting question is whether these techniques are best applied in visibility space or in a real-space image produced by, e.g., maximum-entropy reconstruction. We plan to develop algorithms to address these questions. Because this work will require development of code from scratch (as opposed to adapting existing code), we anticipate seeking funding for a full-time postdoctoral researcher to work on this project.

3.4 Role of Senior Personnel in Astrophysics

Co-PI Bunn will lead the computational astrophysics research, with extensive support from senior personnel Timbie and Tucker. The Richmond, Wisconsin, and Brown research groups have collaborated closely for the past several years in the development of analysis and simulation code for MBI/EPIC. Much of the development has been done by graduate students in the Wisconsin and Brown research groups. If we acquire a new computing cluster at Richmond, the primary role of Timbie and Tucker will be to supervise members of the Wisconsin and Brown groups, who will work closely with Bunn to adapt their code to the new cluster and to design and perform the required simulations.

In the past, the close collaboration among the three groups has been facilitated by weekly teleconferences and regular face-to-face meetings (supported by NASA and NSF funds). We expect this pattern to continue in the future.

As mentioned above, we plan to seek funding in the next few years for a postdoc to work full-time on analysis and simulation issues. This postdoc would be based primarily in Richmond but would work closely with the other groups and spend large amounts of time at Brown and Wisconsin.

4 Research Instrumentation Needs

We request in this proposal funds for the purchase of a cluster of fifty, dual-CPU computers supported by 5.0 terabyte of disk storage and associated hardware and software to increase the productivity of our research efforts at the University of Richmond and to train our undergraduates in modern analysis methods. We describe the current resources available to our group and our computational needs and present a detailed rationale for the proposed system.

4.1 Current Computing Facilities

The current computer system in the Nuclear and Particle Physics Group includes a computing cluster developed in 2001 with NSF and University funds plus an array of computers for software development and non-CPU-intensive calculations and analysis. The system now consists of 34 machines running the Linux operating system and 3 TByte of RAID storage (most nodes are 1.4 GHz machines). Each machine has 18 GByte of disk space and 256 MByte of memory. The entire system resides on its own subnet, and another machine handles all incoming network traffic and security. It resides in a laboratory with a 50-ton, 60,000-BTU air conditioner, an upgraded electrical panel, and a connection to the building's backup power. Nearby rooms provide space for workstations and our students.

The system we now have is near the end of its useful life. We started out with 49 remote nodes and one master and added five new ones over the years. Only 29 of the original remote nodes still work. The usual failure modes are a dead disk drive or a burned out power supply. We have replaced some of these components to resurrect nodes and swapped parts to keep them going. We have also replaced the master node, fileserver, and several power supplies in the RAIDs. It is worth considering the experience at JLab with a very large computing farm with several hundred nodes. They have found that if a node has problems after 2-1/2 years it is not cost effective to fix it and they replace it. All nodes on the JLab farm are replaced after four years of use. The Richmond cluster is now almost six years old.

4.2 Nuclear Physics Computing Needs

The nuclear physics projects described in this proposal all have considerable computing demands. These demands involve the simulation of the CLAS detector to generate publication-quality acceptance functions and adequate disk space and CPU power to perform "second-pass" analysis of the data. To estimate the CPU demands for simulating the CLAS consider the recent experience with the analysis of deuteron structure functions described in Section 2.1. One simulation required 40 million Monte Carlo events for a single beam energy and toroidal magnetic field setting of CLAS. The typical event simulation rate in the package GSIM (the CLAS simulation software) is about 2 Hz on each remote node. The current cluster will take about 6 cluster-days[4] to complete this simulation; the calendar time can be longer because of competition from other users. The proposed cluster will reduce that time down to about one day. The JLab facilities are heavily subscribed, and our existing cluster is aging and falling short. To close this gap additional computing power is necessary. The cluster proposed here will reduce the demand on the Jefferson Lab cluster, speed the calculation of the CLAS acceptance, and complete the analysis of the CLAS data.

We have learned several lessons from our previous experience. The major bottleneck in our data analysis is the speed of the current switch (about 100 Mbps). When we start an analysis run, the data are copied from the RAID to the remote nodes. We have to slow down the analysis so the data rate through the switch does not get too high because so many nodes are trying to

[4] A cluster-day is 24 hours of time on the existing cluster with no competing calculations being performed.

copy files from the RAID out to the cluster. This bottleneck makes it difficult for multiple users to take full advantage of the system. The switch in the proposed system will be 10 times faster.

The disk needs are large for general storage and on the remote nodes. We currently use 2.5 TBytes of storage on the existing RAID out of the 3.0 TBytes available. If we add the astrophysics users and continue to perform extensive simulations those storage demands will only increase. On the remote nodes, when we analyze real data or Monte Carlo data it is more efficient to temporarily store the data on those nodes if the analysis requires repeated runs through the same data set. We save the time to copy data from the RAID onto the individual nodes. This requires adequate storage on the remote nodes. To summarize, we need more and faster remote nodes, a faster switch to reduce the bottleneck of moving data to the remote nodes, and adequate long-term storage and disk space on the remote nodes.

4.3 Astrophysics Computing Needs

We now assess the computational requirements to perform the proposed astrophysics simulations. For the estimates below, we will assume a month-long flight of a balloon-borne MBI-16, observing 1000 square degrees of sky with a 5° field of view in each pointing. With these parameters, the total number of samples in the time-ordered data (TOD) is $N_{tod} \simeq 3 \times 10^8$. The number of independent visibilities, which is also the approximate number of independent pixels in the final map, is $N_v \simeq N_{pix} \simeq 1 \times 10^4$. As described in detail below, the proposed equipment has 100 2.2-GHz CPUs with 2.2 GB of RAM at each node.

The memory requirements for the algorithms below will be satisfied by the proposed cluster. The time-ordered data is always computed and read in relatively small chunks. The largest matrices are $N_v \times N_v$, which can be stored in the 2 GB of RAM in a single node of our proposed equipment; however, even these matrices are all sparse and almost never need to be stored in memory at a single node. We therefore focus on time requirements, not memory.

We now consider the most computationally intensive steps in the proposed simulations.

TOD $\rightarrow$ visibilities. The scaling properties of standard map-making algorithms are non-trivial. We can estimate the time required for this step by comparing with the 2003 flight of the BOOMERanG telescope [64], which had similar values of N_{tod} and N_{pix} to our benchmark parameters. (This experiment is not interferometric, but the algorithm for producing an interferometric visibility-space map is similar to that for a real-space map in an imaging experiment.) They required 20 min to produce a single map with 128 450 MHz processors. Scaling to our equipment specifications, we should be able to generate one visibility map in 5 min.

Power spectrum estimation. Naive power spectrum algorithms require $O(N_v^3)$ time to compute a single likelihood, which makes searching a many-dimensional likelihood parameter space seem prohibitively expensive. Fortunately, a few important insights in the past decade have improved on that estimate. First, the visibility-space covariance matrix is quite sparse: only a fraction $f_s \simeq 0.02$ of the entries are nonzero, leading to a savings of $f_s^{-3/2} \simeq 350$ in computing time [65]. Second, Monte Carlo Markov chains (MCMC) can be used to replace a search of the entire likelihood parameter space for a maximum [66]. Typically, only a few thousand likelihoods need to be evaluated for each MCMC.

Hobson & Maisinger [65] implemented an MCMC algorithm for interferometric data and performed benchmarks. For $N_v = 5000$, they found that a single likelihood evaluation required 30 s of 1-GHz processor time. Scaling to our parameters, we estimate that we can evaluate a single likelihood in about 1 second. A typical MCMC will therefore take roughly one hour.

Maximum Entropy Image Reconstruction. The scaling properties of CMB maximum-entropy algorithms to the large data sets considered here are not well-known. Each calculation of the quantity to be maximized includes a calculation of the χ^2. Since the noise covariance matrix is not diagonal (but is quite sparse), the time required for this step is nontrivial. As it happens, though, the matrix operations required for this operation are quite similar to those required for a single likelihood evaluation of the power spectrum (see above). We expect that the number of entropy evaluations to find the maximum is of order 10^4, meaning that a maximum-entropy map should take a time of the same order of magnitude as a MCMC likelihood evaluation.

The key goal of these simulations is to assess the statistics of the errors in these analysis methods. This requires a significant number of simulations to be performed. Since the slowest steps we wish to perform have time scales of hours, we can perform hundreds of simulations in a reasonable time scale with the proposed equipment.

4.4 Proposed System

We now describe the proposed system that will satisfy our computational needs. The components are listed in Table 1. A detailed quote for items 1-6 is in the supplementary documents from the vendor LinuxLabs. Below we discuss our reasoning behind the choice of the different components.

Item	Number	Description	Price($)
1	1	Dual Opteron master node, 2.2 GHz, 4 GByte RAM, 5 TByte RAID	17,100
2	49	Remote nodes, 2.2 GHz, 2 GByte RAM, 160 GByte storage	192,521
3	1	HP Procurve switch	16,000
4	5	UPS - 5 minutes	34,635
5	5	cabinets	6,039
6	1	Nimbus OS license, installation, and warranty	21,942
7	-	Hardware items that cost less than $500	1,100
		Total Cost	289,337

Table 1: Proposed computer cluster description and cost (see quote in supplementary documents for more details).

The dual-Opteron processors (item 1) were chosen because of their excellent cost-to-benefit ratio. Their clock speed is about 50% faster than the speed of most of the current remote nodes, and architectural improvements make them 4-5 times faster than most of the remote nodes in the existing cluster. The Opteron processor does substantially more computations per clock cycle. The Linux operating system is a research-quality operating system that is commonplace in nuclear physics and astrophysics. The number of machines was chosen to reduce the time to for simulating the CLAS response to a reasonable value. To generate and analyze 40 million events, we estimate about 1 day compared with the time required for the existing cluster (about 6 days) for such a calculation. The astrophysics projects require hundreds of power spectrum estimations and maximum entry reconstructions which each take about a CPU-hour. The memory (2 GByte for each node) is needed because the reconstruction and simulation packages for the nuclear physics work use large amounts of memory and the astrophysics simulations work with large matrices. A 160-GByte hard drive (item 2) will be attached to each machine to provide storage. This space is needed to store data files for analysis, the output of the GSIM simulations to be analyzed by the reconstruction code, and the results of the astrophysical analysis. The fast ethernet switch (item 3) is needed to speed data transfer over the network (see Section 3.2). Backup power supplies (item 4) will prevent damage to the system in the event of a sudden power loss. The supercomputing laboratory has backup power, but there is a lag between power loss and the switch to backup power. Cabinets will hold the nodes (item 5). Hardware and software installation is required (item 6). The software for managing the cluster and submitting batch jobs is Nimbus Beowulf from Linux Labs in Atlanta, GA. This is the vendor who built the current cluster, and we have had a long relationship with them. A variety of other components each costing less than $500 (cables and tools) are included in item 7. We expect the system to have a 4-6 year lifetime. Our experience at Richmond and at JLab suggests that remote nodes will gradually fail over time and that four years is the optimum lifetime. See Section 4.1 for more details.

5 Impact of Project on Teaching and Research

This project will have a significant impact on the development of our students at Richmond and the institutions of the other senior personnel. We describe here the environment at Richmond, how the instrument will be used to train our students, and the impact at those other institutions.

The University of Richmond is a private, highly-selective, primarily-undergraduate, liberal arts institution in Richmond, Virginia with about 3000 undergraduates. A $36M expansion and renovation of the Gottwald Science Center was completed in spring 2006. All of the teaching and research spaces in Physics were renovated, and two new faculty positions were added in Physics (one instructor position and one tenure line). The Department of Physics consists of seven teaching faculty and graduates about 4-8 physics majors each year. The faculty are active in experimental nuclear physics, astrophysics, experimental and theoretical nuclear structure physics, surface and nano-physics, biological physics, and homeland security. There is considerable external support from the Department of Energy (two grants), the National Science Foundation (two grants), and NASA (one grant). We emphasize undergraduate involvement in research from early in the students' careers. Students who are involved in undergraduate research are more likely to attend graduate school and to be successful after college [67].

We have been successful at starting our undergraduates on their research careers. In the summer of 2006 fifteen of our physics majors participated in research at Richmond, JLab, Yale University, and the University of Notre Dame. Our students have accomplished much in their research careers. Eighteen have given presentations on their work at local, national, and international meetings in the last year. The students in the nuclear physics and astrophysics groups have given four presentations in the last year and thirteen over the last 2-3 years [68, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80]. Two of our graduates from 2006 are now attending graduate school in physics (Johns Hopkins University and the University of Kentucky) and three have gone into industry including one in the defense industry. Three more students from our nuclear physics group presented posters at the fall 2006 meeting of the Division of Nuclear Physics (DNP) of the American Physical Society as well as another six students from the experimental and theoretical nuclear structure groups at Richmond. The students in our group obtained travel support in fall 2005 and again in fall 2006 from the Conference Experience for Undergraduates program of the DNP. We have also been successful in attracting under-represented groups to work in our nuclear physics group. Of the four students who worked in our nuclear physics laboratory over the last two summers, two were women and two were African-American men.

The nuclear physics and astrophysics groups each year involve 5-6 students in research. They are integral parts of our research efforts and are deeply involved in many aspects of the physics programs. They receive training in sophisticated analysis methods for extracting signals from complex backgrounds, a range of programming languages (C, C++, FORTRAN, Perl, and IDL), and the Linux operating system. They also learn modern supercomputing methods and are submitting batch jobs to our existing cluster. They have developed their own Perl codes to control the submission of jobs, perform housekeeping in their own directories, and collect the results of the calculations at the end of the submission. It is worth noting that for the students working on JLab physics, this is the best chance they have of getting involved in this type of sophisticated data analysis. The computing farm at JLab is heavily used, and it routinely takes a day or more to get a batch job started. Such a long turnaround time is a barrier to learning and productivity, especially for undergraduates who may have only ten weeks in the summer to work on a project. On the Richmond cluster the turnaround time is typically less. Our Richmond students have performed simulations of the complex CLAS detector [75], analyzed hundreds of gigabytes of data from our experiments discussed in Section 2[68, 70, 72, 74], and made tens of thousands of fits to extract fiducial cuts used to select good events [69, 71, 73, 76]. They are involved in all parts of the analysis.

The astrophysics students have developed a large code base for simulating CMB sky maps and performing a wide variety of statistical analyses on them, including a variety of tests for non-Gaussianity as well as techniques based on wavelet and radon transforms. In the past, we have been able to perform this research on individual workstations, but we have reached the point where a more powerful computing cluster is necessary for further progress. Access to a cluster will be an invaluable resource for these students in their scientific training.

We have also recruited computer science students to help with administration of the cluster. Their duties can range from developing Perl scripts to managing batch jobs, to setting up firewalls to make the system more secure, to maintaining software. Like our other students they have presented their work at national meetings [81, 82].

The project will benefit a significant number of students beyond the University of Richmond. The groups at Union and Ohio together include a postdoctoral fellow, 1-3 graduates students, and many undergraduates. The astrophysics research groups headed by Timbie at Wisconsin and Tucker at Brown typically have a postdoctoral fellow, several graduate students, and sometimes undergraduates working on data analysis and simulation issues.

6 Project Management Plans

The system will be managed at least initially in the same manner the existing cluster is used now. Users will log into the master node to edit, compile, link, test, and execute their codes. They will submit jobs to the cluster from the master. All of students involved in the project, undergraduate and graduate, will have accounts on the master and be able to submit jobs.

The expertise exists in the University of Richmond nuclear and astro-physics groups to operate and maintain the proposed computer cluster. One of us (Gilfoyle) is responsible for maintaining the existing systems, and we all have experience with the Linux operating system. All members of our group have considerable software experience in general and with the codes used by CLAS and in astrophysics. The University administration has adequately supported our research efforts in the past and is committed to continuing to support the University's technology infrastructure. One member of the University's Information Services is a Linux expert who devotes half of his time to academic projects. He is responsible now for keeping the CLAS software up-to-date, updating the Linux software on the cluster and in our laboratory, and general troubleshooting. Finally, we have modeled many features of the proposed computer cluster after existing ones at Jefferson Lab or within the CLAS collaboration. There is a significant amount of expertise within the collaboration that we can call on. The anticipated operating costs are for power and Linux support staff. The University has covered those costs for the existing cluster since 2001 and will continue to do so.

The laboratory that will hold the cluster is complete and in regular use now. It has adequate electrical power and cooling for the proposed instrument. It is described in more detail in Section 4.1. The usage of the current clusters runs anywhere from 10 cluster-hours/week to over 100 cluster-hours/week if many simulations of the CLAS detector are required. The average is around

20-30 cluster-hours/week averaged over a full year with higher demand during the summer. We expect this average demand to increase with the proposed instrument. It will be faster and more reliable, and we have added the astrophysics component to the program. Over the last six years the downtime averages out to 3-4 days per month due to failed components, power outages, *etc.* The rate of failed components has, not surprisingly, increased recently as the system ages.

Currently we informally allocate time on the existing cluster. Users submit jobs when the cluster is open or work out a schedule with the other users. We also partition off subsets of the remote nodes for particular calculations. If demand is high, the Nimbus operating system from LinuxLabs has tools for queuing jobs.

We have attracted numerous other users from JLab. There are currently accounts for fifteen users from the CLAS Collaboration and other groups at JLab including the senior personnel in nuclear physics described in Section 2.5. We expect that we will have little trouble attracting new users to the proposed instrument.

See the letter from the University of Richmond dean in the supplementary materials committing the University to support instrument maintenance, operations, and housing.

7 Dissemination Plan

The work described above will be the subject of internal technical reports at JLab and ultimately publication in refereed journals. Our students will use their results as a springboard into their technical careers by presenting posters and talks at national and international meetings.

8 Results from Prior NSF Support

An NSF Major Research Instrumentation grant in 2001 provided funds (along with $24,000 in matching funds from the University) to purchase the existing cluster in our nuclear physics laboratory at Richmond. The title of the proposal was "RUI: Development of a Computing Cluster to Support the University of Richmond Nuclear Physics Research Program at Jefferson Lab" (#6030194) for $151,758 and for the period 6/01/2001 - 5/31/2003. Dr. M.F. Vineyard (one of the senior personnel in nuclear physics) was a co-PI on that project. All of the Richmond work described in this proposal has made heavy use of the cluster (see Section 2) with the other JLab users. A CLAS technical note on radiative corrections has been published [24] and the measurement of the deuteron structure functions (Section 2.1) is nearing completion and was the subject of a recent contributed talk [85]. A CLAS technical note on the G_M^n measurement is under Collaboration review. All the Richmond students used the cluster [68, 69, 70, 71, 72, 73, 74, 75, 76, 81, 82]. Calculations run on the cluster also led to two more CLAS technical reports by Jenkins [86, 87].

Co-PI Bunn is currently supported by NSF grant AST-0507395, "RUI: Cosmic microwave background analysis in the Post-WMAP era." This grant is for $110,000 over the period from 2005-2008. It has supported Bunn's recent work on two projects that are particularly relevant to this proposal: the development of a detailed framework for analyzing systematic errors in CMB interferometry [59] and a formalism for analyzing mosaicked interferometric observations [61]. In addition, it has supported work on dust contamination in CMB maps by Bunn and undergraduate Gary Larson, work by Bunn on probing the largest-scale perturbation modes in the Universe with Sunyaev-Zel'dovich measurements in distant galaxy clusters [83], and an analysis of the astrophysical constraints on alternative "$f(R)$" theories of gravity [84]. This grant has led to two refereed publications (Astrophys. J. and Phys. Rev. D), two papers currently under consideration in Phys. Rev. D, and three contributed presentations at national meetings of the American Astronomical Society, including one presentation by undergraduate Larson.