Supercomputing for Physics Research at the University of Richmond

G. P. Gilfoyle and E. Bunn
Physics Department
University of Richmond
Richmond, VA, 23173 USA

September 10, 2006

Abstract

We propose to develop a supercomputing cluster to support research programs in physics at the University of Richmond. The proposed research will address fundamental questions in nuclear physics and astrophysics. In addition, it will continue the department's history of training undergraduate students in high-performance computing and data analysis.

The proposed nuclear physics research consists of analysis of data from the CEBAF Large Acceptance Spectrometer (CLAS) at the Thomas Jefferson National Laboratory Facility. Co-PI Gilfoyle has been involved in research with CLAS for many years. The primary goal of this research is to probe the theory of quantum chromodynamics (QCD), which describes the interactions of quarks. QCD is quite successful at high energies, but it has proved difficult to apply to relatively low-energy systems. The experiments to be analyzed probe the transition between these two extremes in order to shed light on the nature of QCD.

The proposed research in astrophysics consists of performing simulations in support of the millimeter-wave bolometric interferometer (MBI), a polarimeter being developed for study of the cosmic microwave background (CMB) radiation. Co-PI Bunn is a leader of the effort to develop this instrument as well as a NASA-funded mission concept study for a potential future satellite-borne version. The proposed simulations will provide a quantitative assessment of the systematic errors in the proposed instrument design, which will be a crucial factor in deciding which technologies to pursue in future CMB polarimetry experiments.

1 Introduction

Funds are requested from the Defense University Research Instrumentation Program (DURIP) to develop a supercomputing cluster to support the research programs at the University of Richmond in nuclear physics and astrophysics. We recognize these sub-disciplines are not typically supported by the Army Research Office (ARO), but our conversations with the management of the ARO suggest that topics within our research areas are relevant to Army and Department of Defense (DoD) needs. Technical aspects of nuclear physics are already supported by the Defense Threat Reduction Agency (DTRA) within DoD. DTRA supports new research in radiological and nuclear detection, the agency is a center of excellence on nuclear weapon burst phenomenology, and it has stockpile stewardship responsibilities [1, 2, 3]. These areas require an understanding of nuclear radiation and how it interacts with matter; essential aspects of the research in the nuclear physics group at Richmond. Similarly, ARO supports research in image analysis and processing that is a component of the analysis of satellite data being performed in the astrophysics group at Richmond

[4]. The proposed funding will, as stated in the criteria for DURIP, 'enhance the institution's ability to educate, through the research to be conducted with the proposed equipment, future scientists and engineers in disciplines important to the DoD mission and will contribute to DoD research-related educational objectives' [4]. These 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 including some that are now in the defense industry using skills they honed in our laboratories. Both of the programs have external support from the US Department of Energy (DOE) (Gilfoyle in nuclear physics) and the National Science Foundation (NSF) (Bunn in astrophysics) that provide student and faculty stipends and travel and equipment money. The DURIP funds will take advantage of these existing programs funded by DOE and NSF. We also have considerable experience with supercomputing. One of us (Gilfoyle) was a co-principal investigator on a project that developed a 100-CPU 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. In the remainder of this proposal we will describe the research activities of the groups at Richmond, our research instrumentation needs, the impact on student training, and project management plans. We also include our curricula vitaes and a price quote for the proposed system.

2 Research Activities

2.1 Nuclear Physics

The research effort in nuclear physics is part of the program at the Thomas Jefferson National Accelerator Facility (TJNAF or JLab) in Newport News, VA. The primary goal of TJNAF is to reveal the quark and gluon structure of protons, neutrons, and atomic nuclei and to deepen our understanding of matter and, in particular, the confinement of quarks. In this section we describe the experimental environment and the specific physics programs.

TJNAF is a unique tool for basic research in nuclear physics. The central instrument is a superconducting electron accelerator with a maximum energy of 4-6 GeV, a 100% duty cycle, and a maximum current of 200 μA. These excellent beam characteristics allow for novel experiments that are being used to develop a quark-based understanding of nuclei. The electron beam is used simultaneously for scattering experiments in three halls that contain complimentary experimental equipment. Our research is done 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 [5] to determine the trajectories of charged particles, Cerenkov detectors [6] for the identification of electrons, scintillation counters [7] for time-of-flight measurements, and electromagnetic calorimeters [8] to identify electrons and to detect photons and neutrons. The six segments are instrumented individually to form six independent spectrometers. Together there are about 33,000 detecting elements. The CLAS detector is capable of acquiring approximately 1 terabyte of data per day. The CLAS detector was constructed and is operated by an international collaboration consisting of about thirty-five institutions. The Richmond group has been part of the CLAS Collaboration since its inception, and has been actively involved in the construction of the spectrometer and the development of the physics program.

The analysis of these large data sets requires several stages. The initial phase of the data analysis is done on the JLab computer farm where components of CLAS are calibrated, events are reconstructed, and particles are identified with comparatively loose criteria with a program called RECSIS and stored on tape. The next phase is to refine the first pass results and perform more sophisticated analysis to investigate the underlying physics. There is considerable demand for the computing resources at JLab so we have developed our own local computing cluster so we can analyze our data in a timely fashion. The large, first-pass, data sets require considerable storage space to hold the data and a large number of nodes running in parallel to perform the analysis productively. We also simulate the response of the detector so we can separate real physics effects from artifacts of the detector. We do this step using the GEANT simulation package from CERN adapted for use with CLAS and called GSIM. We generate Monte Carlo events and then pass that simulated data through our analysis codes to study how the CLAS acceptance alters the distributions of the original simulated events. This stage of the analysis also requires considerable disk space to store the Monte Carlo events and, perhaps more importantly, considerable computing power. We can detect real data with CLAS at a rate of 4,000 events per second, but our simulation can generate Monte Carlo events at only about 2 events per second.

The physics projects are all focused on the central missions of JLab. Protons and neutrons are comprised of three valence quarks held together by the exchange of gluons, along with a quark-gluon sea of particles that are constantly popping into and out of existence. We want to learn how quarks, gluons and the quark-gluon sea combine to form nucleons and what forces mediate these interactions. According to Quantum Chromodynamics (QCD), the fundamental theory of particle physics, quarks are bound together by the strong or color force. QCD is a highly successful description of quarks at very high 4-momentum transfers or Q^2 [9], but at the energies where proton and neutrons exist (the non-perturbative region), it has, so far, proved to be daunting challenge to solve [10]. At low 4-momentum transfers ($Q^2 < 0.5$ (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 [11]. However, the transition region between these extremes is poorly understood in the range $Q^2 \approx 0.5 - 10.0$ (GeV/c)^2 and studying this transition is an essential goal of nuclear physics as described by the Nuclear Science Advisory Committee (NSAC)[1] [12]. One of the primary missions of JLab is to map the geography of this unknown region and make the measurements that will challenge theorists to solve QCD at low energies and develop a quark-gluon based picture of nucleons and nuclei. Below we discuss each research project, summarize its current status, and plan the future. We also describe how the proposed instrument will be used in each project.

2.1.1 Out-of-Plane Structure Functions of the Deuteron

As mentioned above, the hadronic model of nuclear physics has been successful at low Q^2, but it is not well-developed in the GeV region even though we expect it to be valid. There has simply been little data to challenge theory. The relative importance of relativistic corrections, final-state interactions, meson-exchange currents, and isobar configurations is unknown. We need measurements to provide a baseline for the hadronic model so deviations at higher Q^2 can be attributed to quark-gluon effects with greater confidence. This is an important step in unraveling the quark-gluon substructure of nucleons and nuclei according to NSAC and the JLab Program Advisory Committee (PAC)[12, 13]. To this end we are investigating the out-of-plane structure functions of the deuteron using the reaction $d(\vec{e}, e'p)n$ with CLAS. The deuteron is an essential testing ground because it is the simplest nucleus. The cross section for the reaction can be written as

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

---

[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. NSAC is chartered under the Federal Advisory Committee Act (FACA).

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 $\vec{q}$. This angle ϕ_{pq} can be viewed as the angle between the plane defined by the incoming and outgoing electron and the plane defined by the ejected proton and neutron. The structure functions are an essential meeting ground between theory and experiment and the unique, nearly-4π solid angle of CLAS creates an inviting opportunity to study f'_{LT}, f_{LT}, and f_{TT}. These structure functions depend on ϕ_{pq} and have not been extensively investigated in the past. They represent a model-independent measurement of a little-studied part of the deuteron cross section and probe its wave function and the charged currents within the deuteron.

The three structure functions are being 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 $\mathrm{Q}^2 = 0.2 - 5.0$ $(\mathrm{GeV/c})^2$. We are studying the reaction in quasi-elastic kinematics first and later will investigate higher energy transfers. In quasi-elastic kinematics, the electron projectile scatters nearly elastically off the proton within the deuteron. The analysis of the structure function f'_{LT} for quasielastic kinematics is far along and will be completed in the near future. We are refining our simulations of CLAS to ensure that we understand the detector response. Our preliminary results show we can observe small asymmetries with good precision in quasi-elastic kinematics. We also observe significant deviations from the hadronic model at modest 4-momentum transfer ($\mathrm{Q}^2 \approx 0.5$ $(\mathrm{GeV/c})^2$. 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 will follow. This work is part of a CLAS Approved Analysis entitled 'Out-of-Plane Measurements of the Structure Functions of the Deuteron'. The CLAS collaboration has a procedure where collaboration members can analyze existing data sets with official collaboration approval. The collaboration member writes a proposal describing an analysis project, it is reviewed by a committee of collaboration members, and then defended before the full collaboration. Gilfoyle is the spokesperson on this analysis project that was approved by the collaboration in November, 2003 [14]. The analysis of the data and the simulations of CLAS are computationally intensive and would use the proposed cluster.

2.1.2 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 magnetic form factor, G_M is the magnetic form factor, and

$$\tau = \frac{Q^2}{4M^2} \qquad \text{and} \qquad \epsilon = \frac{1}{1 + 2(1+\tau)\tan^2(\frac{\theta}{2})} \tag{3}$$

where M is the nucleon mass and θ is the electron scattering angle. There are a total of 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. These observables also provide stringent tests of non-perturbative QCD and are connected to generalized parton distributions (GPDs) via sum rules. 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. This new generalization of the parton distributions permits us to simultaneously extract information about the longitudinal and transverse parton structure of hadrons. With exclusive measurements one can determine the longitudinal momenta and transverse position of the partons inside the proton, 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 if its important tests will be the accuracy that 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 Jefferson Lab [18]. All four elastic form factors are needed to untangle the different quark contributions. Our focus is on the magnetic form factor of the neutron. For 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 and we will have consistency checks between different detector components and an overlap with our previous CLAS measurements. 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 manuscript describing that work is under collaboration review. We have submitted a Letter-of-Intent to the JLab Program Advisory Committee (PAC) to make the same measurements at higher Q^2.[2] A Letter-of-Intent is a preliminary proposal for beam time at JLab. It provides an opportunity for researchers to get feedback from the PAC on the quality of their idea before they 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 now have one year to develop a full beam time proposal for a future PAC. We have been involved with the analysis of the G_M^n data for several years now because the deuteron structure function analysis project described in Section 2.1.1 uses the same data set. Gilfoyle is the spokesperson and contact person for the Letter-of-Intent. The simulations necessary to prepare the full proposal will take advantage of the computing cluster proposed here.

2.1.3 Quark Propagation and Hadron Formation

The confinement of quarks inside hadrons is perhaps the most remarkable feature of QCD. The quest to understand confinement quantitatively and in terms of intuitive physical pictures is an essential goal of modern nuclear physics. Much experimental attention has been focused on understanding confinement through hadron spectroscopy. Alternatively, the subject is often introduced through sketches of the hadronization processes by string-breaking. This picture is confirmed by lattice calculations using static quarks, showing that the gluon field is concentrated in a flux-tube (or string). In hadronization, the color string stretches until $q\bar{q}$ pairs tunnel up from the vacuum, thwarting the struck quark's attempt to escape to isolation. Unfortunately, the real picture with full QCD is more complicated than this simple picture, and lattice calculations of dynamical quarks are not yet possible. Experimental information is necessary to guide models of hadronization, to elucidate the mechanism of confinement.

We have proposed a broad program of measurements and analyzes to determine the mechanisms of confinement in forming systems. The essential experimental technique is to employ nuclei as analyzers of hadronization processes. In this approach, the hadron is formed from energetic quarks over distance scales ranging from 0–10 fm, i.e. the dimensions of atomic nuclei. The power of this technique arises from several factors: (1) the dimensions of the analyzing medium are per-

---

[2] The DOE plans to upgrade the electron 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 out of 28 for the DOE office of Science in the next 20 years [21].

fectly matched to the distance scales of the hadronization process; (2) sophisticated knowledge about nuclear currents and properties can be exploited; (3) detailed experimental data on deep inelastic scattering from the nucleon provide a quantitative baseline against which to compare the data on nuclei. In essence, 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 hardrons produced relative to the production from deuterium (hadronic multiplicity ratio R_M^h) and transverse momentum broadening Δp_T^2 are the two primary observables. These are known to be sensitive to different kinematic regimes (which gives us the opportunity to study their evolution) as well as the average distance through the nuclear medium $\langle L \rangle$. 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 τ_p is the lifetime of a deconfined quark, and it will be determined by analyzing the transverse momentum broadening Δp_T^2 as a function of $\langle L \rangle$ and ν, the energy transferred by the electron beam. The length of the linear region yields the production time. The formation times τ_f^h 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_M^h by using τ_p and models for interactions of pre-hadrons and hadrons in nuclei. Studying the systematic behavior of formation times for a variety of hadrons will determine the mechanisms by which hadrons are formed. A proposal to do this experiment at high Q^2 was approved by the JLab PAC August, 2006 [23]. Gilfoyle is a spokesperson on that proposal and will be responsible for analyzing the π^0, η, and η' production. As this topic is developed for future workshops and updates, we will use the proposed computing cluster to simulate the physics and the upgraded CLAS detector response to prepare for the 12-GeV Upgrade.[3]

2.1.4 Technical Projects

The measurements of the nuclear reactions described above are subject to radiative corrections. A new computer code called DEEP_EXCLURAD for calculating these corrections developed by one of us (Gilfoyle) can be used for exclusive reactions using electrons. Radiative corrections are usually calculated using the formalism originally developed by Schwinger and Mo and Tsai [24, 25]. In that approach, it was assumed that only the scattered electron was detected (inclusive scattering). That method suffers from several shortcomings. First, detecting the ejected hadron (exclusive reactions) alters the phase space that is allowed for the final radiated photon. Second, more structure functions can contribute in exclusive reactions. Only the longitudinal and transverse pieces contribute in the Schwinger/Mo and Tsai method while for the out-of-plane production analysis (see Section 2.1.1) there are three other components associated with the f_{LT}, f_{TT} and f'_{LT} structure functions. Third, the Schwinger/Mo and Tsai approach relies on an unphysical parameter to split the hard and soft regions of the radiated photon's phase space and cancel the infrared divergence. The method used here relies on a covariant procedure of infrared divergence cancellation which does not require the splitting [26].

These calculations are being done with a modified version of the computer program EXCLURAD written by Afanasev, *et al.* [27]. The code was originally written for the $p(e, e'\pi^+)X$ reaction and it has been modified to work for the $d(\vec{e}, e'p)n$ reaction as part of the deuteron-structure-function program described in Section 2.1.1. A separate program called DEEP is being used to calculate the deuteron response functions [28]. This routine uses the covariant spectator theory and

---

[3] The schedule for the JLab 12-GeV Upgrade is still tentative, but we now expect first beams around 2014. In the meantime, this proposal will have to be updated for future PAC reviews.

the transversity formalism to calculate the unpolarized, coincidence cross section for $d(\vec{e}, e'p)n$. The program calculates the ratio of the cross section in a particular kinematic bin to the plane-wave-impulse-approximation (PWIA) result. Details on running the code can be found in Reference [29]. Note that because of the adaptive stepsize used in performing some of the integrals in the code, 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. The Department of Energy has approved the initial phases of the process that will double the beam energy of the electron accelerator and modify the experimental equipment in Hall B to take full advantage of the new physics opportunities. 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. The online data rate is expected to be 20 kHz with 10 kBytes for each of those events and less than 15% deadtime [30]. We will collect about 1 petabyte of data each year.

Event simulation is an essential aspect of the design of CLAS12 and eventual precision of the detector. For many experiments, the quality of the results will be limited by systematic uncertainties instead of statistical ones. Accurate, precise calculations of the CLAS12 acceptance and response are important to keep those systematic uncertainties small. To do that we expect to generate about four times as much simulated, Monte Carlo data as CLAS12 collects. The CLAS12 simulation will produce data more slowly than the detector itself (see Section 3.2 below for the current status of the simulation speed in CLAS), so the contribution of university groups to this effort is essential. The same issues arise in the design and prototyping phase of the project we are just entering. The CLAS12 software group (Gilfoyle is a member) is now beginning the effort to simulate events. The work will definitely make significant use of the proposed cluster during much of its lifetime.

2.2 Astrophysics

In the past ten years, cosmology has become a data-rich, relatively high-precision branch of astrophysics, with a “standard model” that is well-constrained by a variety of observations. For example, we now have reliable estimates of parameters such as the age and expansion rate of the Universe and the densities of dark matter and dark energy. With continued advances in observing technology in recent years, we are poised to make further advances in the coming decade.

Observations of temperature anisotropy in the cosmic microwave background (CMB) radiation have been among the leading causes of this revolution, with major contributions coming from the NASA-funded satellite missions COBE [31] and WMAP [32] as well as a large number of suborbital experiments. 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 [33, 34, 35, 36, 37, 38]. 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., [39]), probe the ionization history of the Universe [40] and the details of recombination [41], and measure gravitational lensing due to large-scale structure [42]. Most exciting of all, polarization maps may provide a direct probe of an inflationary epoch in the extremely early Universe [43, 44].

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" [45]. Similar opinions are expressed in the National Research Council's decadal survey of astronomy and astrophysics [46] and their report *Connecting Quarks with the Cosmos* [47].

NASA has recognized CMB polarimetry as a top priority. The agency's *Beyond Einstein* road map for future astrophysics programs includes three satellite-borne "Einstein Probes," one of which is to be a dedicated CMB polarimeter known as the Einstein Inflation Probe (EIP) [48]. Co-PI Bunn is one of the leaders, along with researchers at Brown University and the University of Wisconsin, of a NASA-funded Mission 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 [49, 50]. The work proposed herein is to develop data analysis and simulation tools in support of MBI and EPIC.

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 [43, 44], 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 1). 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 1 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 2.2.1, we briefly outline some differences between interferometric CMB observations and traditional imaging systems, and in Section 2.2.2 we describe the particular simulations to be performed.

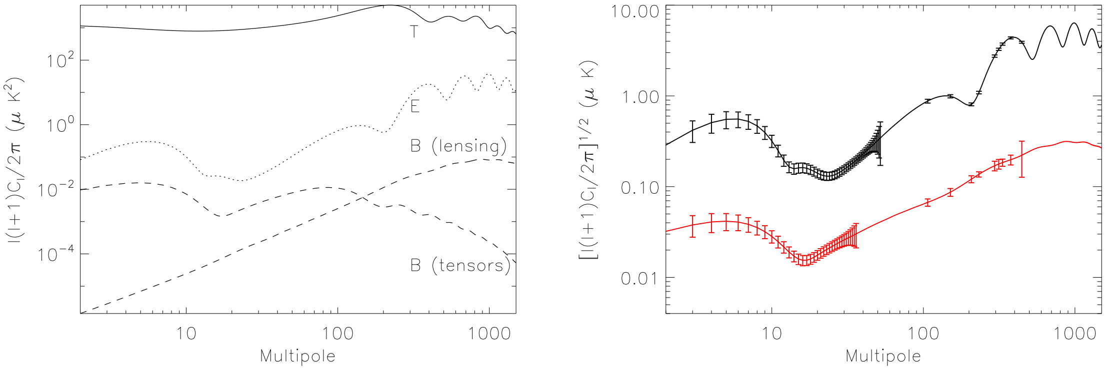

Figure 1: The left panel shows CMB angular power spectra for temperature anisotropy, E polarization, and B polarization. For the B spectrum, the tensor and lensing contributions are shown separately. These spectra are based on the best-fit model from WMAP [32], with the tensor-to-scalar ratio taken to be 0.1. The right panel shows forecasted errors on the E (upper) and B (lower) signals for EPIC [49]. 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.

2.2.1 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:

1. **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.

2. **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.

3. **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.

4. **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 $\sim 10'$ [51]. 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 [52].

5. **Clean separation of E and B modes**. In any incomplete sky map, there is some "leakage" between E and B components [53, 54, 55]. 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 [56].

6. **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 [51] but do not in an interferometer [57]. Since the unpolarized signal is orders of magnitude larger than the polarized signal (Fig. 1), this makes a big difference.

2.2.2 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 [57], 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. Although our work will focus on the MBI/EPIC design, the conclusions should be applicable to other CMB interferometers.

Along with the construction of the prototype interferometer MBI-4, the MBI/EPIC research group (headed by Co-PIs Bunn at Richmond, Tucker at Brown, and Timbie at Wisconsin) has been developing a data analysis pipeline. Many steps in the pipeline can be performed with adaptations of publicly-available parallelized code (e.g., [58]); 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. Although we are focusing on the MBI design in particular, we expect the results to be applicable to CMB interferometers in general.

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 3.3.

1. **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.

2. **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.[4] 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 [58], originally developed for traditional imaging systems, can be adapted for this purpose.

3. **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.

4. **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 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 [60]. 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 [61]. 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.

2.2.3 Component separation

The simulations described in the preceding section will be our primary focus over the period funded by this grant. 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

---

[4] 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 [59].

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 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, our computational needs and present a detailed rationale for the proposed system.

The proposed instrument will be used by the faculty and students in the University of Richmond nuclear and astrophysics groups. Undergraduates will be intimately involved in the use of the proposed system. The University of Richmond's primary mission is undergraduate education (there are no graduate students in the Physics Department) and the University is committed to research involving undergraduates as a means to broaden our students education beyond the scope of the classroom. In recent years we have routinely had 5-6 students each summer involved in nuclear and astrophysics. The proposed instrument will make their research (and the faculty's) more productive and will train them on the most modern physics instruments.

3.1 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, dual-processor machines running the Linux operating system and 3 TByte of RAID storage (29 of the 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 acts as a firewall. The remaining machines in the cluster are relieved of responding to outside network traffic and access (and security) can be controlled and monitored through a single computer. It resides in a laboratory designed specifically for its needs which require a 50-ton, 60,000-BTU air conditioner to cool the room, an upgraded electrical panel, and a connection to the building's backup power. The system we now have is near the end of its useful life. We started out with 49 remote nodes and one master and have 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 even swapped parts to keep them going, but their days are numbered. We have also replaced the master node, fileserver, and several power supplies in the RAID. It is worthwhile to consider 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 worth fixing and they replace it. All nodes are replaced after four years of use. The Richmond cluster is now five years old.

The remainder of the laboratory consists of nine linux machines for student and faculty use. The software used in the both the nuclear physics and astrophysics research is non-proprietary. The University has hired information services staff with linux expertise. They provide support for updating software and problem solving. See Program Management in Section 5.

3.2 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 2.1.1. The goal of this CLAS experiment is to measure the out-of-plane structure functions of the deuteron. The analysis of this data is far along and the simulation of the response of the CLAS relevant to this reaction channel is underway. 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 $R_{gsim} \approx 2.4\ Hz$ on each remote node. This code uses about 90% of the CPU time. The entire process requires event generation, simulation (with GSIM), reconstruction (with RECSIS), final pass analysis, and various formating and housekeeping tasks. The 34 nodes in the current cluster will take about 6 CPU-days to complete this simulation; about one week or more depending on the load on the cluster. The actual calendar time can be much longer because of several competing activities. Other members of the group will be simultaneously performing simulations relevant to their studies. Second-pass analysis of the data will be performed on the same machines. Finally, some CPU time will be used for development, testing, and data transfer. These other activities can increase the time needed to perform the simulation by factors of several. The proposed cluster will reduce that CPU time down to about one day. The six-fold improvement comes from a factor of 4-5 increase in the computational speed and a nearly 50% increase in the number of working nodes. This dramatic improvement means that students and faculty can perform several calculations in a week and more precisely test ideas for analysis and simulation. They will be more productive. This is especially important for projects like the deuteron structure function analysis project because the data set consists of really three subsets: two different beam energies and two different magnetic field settings for the lowest beam energy. The Jefferson Lab facilities are inadequate for all the computing demands of the CLAS collaboration and our existing cluster falls short of our needs. 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 also considered the effect of maintaining the existing cluster and adding it to the proposed one. This increases the complexity of the project here and only improves the performance by about 10%.

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 copy files from the RAID out to the cluster. If we don’t wait for 2 minutes between batch job submissions the system will hang (some data sets require 80-90 batch submissions). 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 makes better use of the system. To conclude, 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.

3.3 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 → visibilities. The scaling properties of standard map-making algorithms are nontrivial. However, we can estimate the time required for this step by comparing with the 2003 flight of the BOOMERanG telescope [62], 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 estimate that 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 [63]. Second, Monte Carlo Markov chains (MCMC) can be used to replace a search of the entire likelihood parameter space for a maximum [64]. Typically, only a few thousand likelihoods need to be evaluated for each MCMC.

Hobson & Maisinger [63] 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. For our benchmark parameters, N_v is larger by a factor of 4, so the algorithm should take 8 times longer. With 100 2.2-GHz processors, 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 chi-squared. 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.

3.4 Proposed System

We now describe the proposed system that will satisfy our computational needs. The components are listed in Table 3. A detailed quote for items 1-6 is in the Appendix from Linux Labs. 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 Appendix 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 (according to the vendor) 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. The proposed cluster will generate and analyze 40 million events in a much shorter time. We estimate about 1 day compared with the time required for the existing cluster (about 6 days). The proposed machines reduce the time for an acceptance calculation by a factor of six. 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 (RECSIS and GSIM) 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 RECSIS 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, fruitful 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 3.1 for more details.

4 Student Contributions and Training

The University of Richmond is a private, highly-selective, primarily-undergraduate, liberal arts institution in Richmond, Virginia. There are about 3000 undergraduates. The Department of Physics consists of seven teaching faculty and we graduate 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. 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 eleven over the last two years [65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75]. 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.

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. In the last year two of those students were women and two were African-American. The skills they learn are relevant to DoD needs and will 'enhance the institution's ability to educate ... future scientists and engineers in disciplines important to the DoD mission' [4].

5 Project Management Plans

The system will be properly maintained. The expertise exists in the University of Richmond Nuclear and Particle Physics Group to develop 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 has adequately supported our research efforts in the past and we expect them to continue supporting the University's technology infrastructure. One member of the University's Information Services is a linux expert and he 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.

References

[1] 'Basic Research Program in Chemical, Biological, Radiological and Nuclear Defense of the Joint Science and Technology Office, Physical Science and Technology Division, Defense Threat Reduction Agency', U.S. Army Research Office, Chemical Sciences Division, March, 2006.

[2] 'WMD Defeat Technology', Exhibit R-2a, RDT&E Project Justification, Feb, 2005. `http://www.dod.mil/comptroller/defbudget/fy2006/budget_justification/pdfs/rdtande/DTRA%200602716BR%20R-2A.doc`, last accessed Aug 30, 2006.

[3] 'DTRA Fact Sheet', `http://www.dtra.mil/newsservices/fact_sheets/display.cfm?fs=dtra`, last accessed Aug 30, 2006.

[4] 'Fiscal Year 2007 Defense University Research Instrumentation Program (DURIP)', Army Research Office, 2006.

[5] M. D. Mestayer *et al.*, Nucl. Inst. and Meth. **A449**, 81 (2000).

[6] G. Adams *et al.*, "The CLAS Cerenkov Detector", Nucl. Instr. and Meth. **A465**, 414 (2001).

[7] E.S. Smith *et al.*, Nucl. Inst. and Meth. **A432**, 265 (1999).

[8] M. Amarian *et al.*, "The CLAS Forward Electromagnetic Calorimeter", Nucl. Instr. and Meth. **A460**, 239 (2001).

[9] G.P. Lepage and S.J. Brodsky, Phys. Rev. D **22**, 2157 (1980).

[10] R. Gupta, 'Introduction to Lattice QCD', hep-lat/9807028, (1998).

[11] S.Gilad, W. Bertozzi, and Z.-L. Zhou, Nucl. Phys. **A631**, 276 (1998) and references therein.

[12] 'Opportunities in Nuclear Science: A Long-Range Plan for the Next Decade', DOE/NSF Nuclear Science Advisory Committee, April 2002.

[13] PAC 14 Few-Body Workshop, JLab Program Advisory Committee, Williamsburg, VA, July, 1998.

[14] G.P. Gilfoyle, W.K. Brooks, B.A. Mecking, S.E. Kuhn, L.B.Weinstein, and M.F. Vineyard, 'Out-of-Plane Measurements of the Structure Functions of the Deuteron', CLAS Approved Analysis (2003).

[15] M.N. Rosenbluth, Phys. Rev. **79**, 615 (1950).

[16] M.Diehl, 'Introduction to Generalized Parton Distributions', `wwwcompass.cern.ch/compass/publications/2004_yellow/Body/diehl_new.pdf`, last accessed Sep 4, 2006.

[17] J.D. Ashley, D.B. Leinweber, A.W. Thomas, and R.D.Young, Eur. Phys. J A **19** s01, 9 (2004).

[18] C.E. Hyde-Wright and K.deJager, Ann. Rev. Nucl. Part. Sci. **54**, 217 (2004).

[19] W. Brooks and M.F. Vineyard, *The Neutron Magnetic Form Factor from Precision Measurements of the Ratio of Quasielastic Electron-Neutron to Electron-Proton Scattering in Deuterium*, Jefferson Lab Experiment E94-017 1994.

[20] J.D. Lachniet, *A High Precision Measurement of the Neutron Magnetic Form Factor Using the CLAS Detector*, thesis, Carnegie Mellon University.

[21] 'Office of Science Strategic Plan, February, 2004', US Department of Energy Office of Science, February, 2004.

[22] 'Measurement of the Neutron Magnetic Form Factor at High Q^2 Using the Ratio Method in Deuterium', G.P. Gilfoyle, M.F. Vineyard, S.E. Kuhn, J.D. Lachniet, L.B.Weinstein, K. Hafidi, M. Holtrop, M. Garcon,Jefferson Lab Letter-of-Intent LOI-12-06-107 (2006).

[23] 'Quark Propagation and Hadron Formation', K. Hafidi, J.Arrington, L. El Fassi, D.F. Geesaman, R.J. Holt, B. Mustapha, D.H. Potterveld, P.E.Reimer, P. Solvignon, K. Joo, M. Ungaro, G. Niculescu, I. Niculescu, W.K. Brooks, M. Holtrop, K. Hicks, T. Mibe, L.B. Weinstein, M. Wood, G.P. Gilfoyle, H. Hakobyan, Jefferson Lab Experiment PR-12-06-117 (2006).

[24] J.Schwinger, Phys. Rev., **76**, 898 (1949).

[25] L.W.Mo and Y.S.Tsai, Rev. Mod. Phys., **41**, 205 (1969).

[26] D.Y.Bardin and N.M.Shumeiko, Nucl. Phys. B, **127**, 242 (1977).

[27] A.Afanasev, I.Akushevich, V.Burkert, and K.Joo, Phys.Rev., **D66**, 074004, 2002.

[28] J.Adam, Jr., F.Gross, S.Jeschonnek, P.Ulmer, and J.W. Van Orden, Phys. Rev. C, **66**, 044003 (2002).

[29] G.P. Gilfoyle and A. Afanasev, 'Radiative Corrections for Deuteron Electro disintegration', CLAS-Note 2005-022.

[30] *Pre-Conceptual Design Report (pCDR) for The Science and Experimental Equipment for The 12 GeV Upgrade of CEBAF*, Jefferson Lab report 2004.

[31] G.F. Smoot *et al.*, Astrophys. J. Lett., **396**, L1 (1992).

[32] C.L. Bennett *et al.*, Astrophs. J. Supp., **148**, 1 (2003).

[33] J.M. Kovac *et al.*, Nature, **420**, 772 (2002).

[34] A. Kogut *et al.*, Astrophys. J. Supp., **148**, 161 (2003).

[35] L. Page *et al.*, astro-ph/0603450, submitted to Astrophys. J. (2006).

[36] A.C.S. Readhead *et al.*, Science, **306**, 836 (2004).

[37] E.M. Leitch *et al.*, Astrophys. J., **624**, 10 (2005).

[38] D. Barkats *et al.*, Astrophys. J. Lett., **619**, L127 (2005).

[39] W.H. Kinney, Phys. Rev. D, **58**, 123506 (1998).

[40] M. Zaldarriaga, Phys. Rev. D., **55**, 1822 (1997).

[41] P.J.E. Peebles, S. Seager, and W. Hu, Astrophys. J. Lett., **539**, L1 (2000).

[42] M. Zaldarriaga and U. Seljak, Phys. Rev. D, **58**, 023003 (1998).

[43] M. Zaldarriaga and U. Seljak, Phys. Rev. D, **55**, 1830 (1997).

[44] M. Kamionkowski, A. Kosowsk, and A. Stebbins, Phys. Rev. D, **55**, 7368 (1997).

[45] J. Bock, S. Church, M. Devlin, G. Hinshaw, A. Lange, A. Lee, L. Page, B. Partridge, J. Ruhl, M. Tegmark, P. Timbie, R. Weiss, B. Winstein, & M. Zaldarriaga, Task Force on Cosmic Microwave Background Research, available at http://arxiv.org/abs/astro-ph/0604101 (2006).

[46] Astronomy & Astrophysics Survey Committee (National Research Council), *Astronomy & Astrophysics in the New Millennium*, National Academies Press (2001).

[47] Committee on the Physics of the Universe (National Research Council), *Connecting Quarks with the Cosmos: 11 Science Questions for the New Century*, National Academies Press (2003).

[48] http://universe.nasa.gov/program/probes.html

[49] P. Timbie *et al.*, "The Einstein Polarization Interferometer for Cosmology (EPIC) and the Millimeter-wave Bolometric Interferometer (MBI)," to appear in Proceedings of the Fundamental Physics with CMB Workshop, UC Irvine, New Astronomy Reviews, in press (2006).

[50] A.L. Korotkov *et al.*, Proc. SPIE, Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III; Jonas Zmuidzinas *et al.*, eds., **6275**, 285 (2006).

[51] W. Hu, M.M. Hedman, and M. Zaldarriaga, Phys. Rev. D, **67**, 043004 (2003).

[52] L. Knox and Y.-S. Song, Phys. Rev. Lett., **89**, 011303 (2002).

[53] A. Lewis, A. Challinor, and N. Turok, Phys. Rev. D, **65**, 023505 (2002).

[54] E.F. Bunn, Phys. Rev. D, **65**, 043003 (2002). See also erratum at Phys. Rev. D, **66**, 069902 (2002).

[55] E.F. Bunn, M. Tegmark, M. Zaldarriaga, and A. de Oliveira-Costa, Phys. Rev. D, **67**, 203501 (2003).

[56] C.-G. Park and K.-W. Ng, Astrophys. J., **609**, 15 (2004).

[57] E.F. Bunn, astro-ph/0607312, submitted to Phys. Rev. D (2006).

[58] http://crd.lbl.gov/~borrill/cmb/madcap/

[59] E.F. Bunn and M. White, astro-ph/0606454, submitted to Astrophys. J. (2006).

[60] J. Högbom, Astrophys. J. Supp., **15**, 417 (1974).

[61] K. Maisinger, M.P. Hobson, and A. Lasenby, Mon. Not. R. Astron. Soc., **290**, 313 (1997).

[62] S. Masi *et al.*, astro-ph/0507509 (2005).

[63] M.P. Hobson and K. Maisinger, Mon. Not. R. Astron. Soc., **334**, 569 (2002).

[64] N. Christensen, R. Meyer, L. Knox, and B. Luey, Class. Quant. Grav., **18**, 2677 (2001).

[65] R.Burrell and G.P.Gilfoyle, ‘Momentum Corrections for the CLAS E5 Data Set’, University of Richmond Symposium, April 21, 2006.

[66] K.Greenholt and G.P.Gilfoyle, ‘Fiducial Cuts on CLAS for the E5 Data Set’, Bull. Am. Phys. Soc., Fall DNP Meeting, FR.0006(2005).

[67] R.Burrell and G.P.Gilfoyle, ‘Momentum Corrections for the E5 Data Set’, Bull. Am. Phys. Soc., Fall DNP Meeting, FR.00068(2005).

[68] K.Greenholt and G.P.Gilfoyle, ‘Generating Fiducial Cuts for CLAS E5’, University of Richmond Symposium, April 14, 2005.

[69] R.Burrell and G.P.Gilfoyle, ‘Momentum Corrections for the CLAS E5 Data Set’, University of Richmond Symposium, April 14, 2005.

[70] K.Greenholt and G.P.Gilfoyle, ‘Generating Fiducial Cuts for CLAS E5’, Bull. Am. Phys. Soc., April Meeting, S.06(2005).

[71] R.Burrell and G.P.Gilfoyle, ‘Momentum Corrections for the CLAS E5 Data Set’, Bull. Am. Phys. Soc., April Meeting, S.07(2005).

[72] G. Larson, E.F. Bunn, V. Kasliwal, M. McCann, ‘Filtering Dust Contamination from CMB Data with Wavelets and Radon Transforms’, American Astronomical Society Meeting, Washington, DC, January 2006.

[73] M. McCann, E.F. Bunn, R. Chan, V. Kasliwal,‘The Three-Point Correlation Function of Galactic Dust: Implications for Microwave Background Non-Gaussianity’, American Astronomical Society Meeting, Denver, CO, June 2, 2004.

[74] V. Kasliwal, E.F. Bunn, R. Chan, M. McCann, ‘The Bispectrum of Galactic Dust: Implications for Microwave Background Non-Gaussianity’, American Astronomical Society Meeting, Denver, CO, June 2, 2004.

[75] M. McCann, ‘E.F. Bunn, The Three-Point Correlation Function as an Indicator of non-Gaussianity in the Cosmic Microwave Background Radiation’, University of Richmond Symposium, April 20, 2005.

Dr. Gerard P. Gilfoyle

Physics Department
University of Richmond, VA 23173
ggilfoyl@richmond.edu

3234 Kensington Avenue
Richmond, VA 23221
804-289-8255

Degrees
Ph.D., University of Pennsylvania, 1985 - 'Resonant Structure in ^{13}C(^{13}C,^{4}He)^{22}Ne', H.T. Fortune, adviser.
A.B., cum laude, Franklin and Marshall College, 1979.

Experience
2004-present - Professor of Physics, University of Richmond.
2002-2003 - Scientific Consultant, Jefferson Laboratory.
1999-2000 - AAAS Defense Policy Fellow.
1994-1995 - Scientific Consultant, Jefferson Laboratory.
1993-2004 - Associate Professor of Physics, University of Richmond.
Summer, 1988 - Visiting Research Professor, University of Pennsylvania.
1987-1993 - Assistant Professor, University of Richmond.
1985-1987 - Postdoctoral Research Fellow, SUNY at Stony Brook.
1979-1985 - Research Assistant, University of Pennsylvania.

Research and Teaching Grants
1990-present - US Department of Energy ($1,361,000).
2002-2003 - SURA Sabbatical Support ($10,000).
2002-2003 - Jefferson Laboratory Sabbatical Support ($28,335).
2001-2002 - National Science Foundation ($175,000).
1995-1997 - National Science Foundation($14,986).
1994-1995 - CEBAF Sabbatical Support ($24,200)
1992-1995 - National Science Foundation ($49,813).
1989-1991 - Research Corporation($26,000).
1987-2002 - University of Richmond Research Grants($13,082).

Selected Service
2006 - present - Chair, Nuclear Physics Working Group, CLAS Collaboration.
2005 - Reviewer, National Science Foundation (Nuclear Physics).
2003 - present - Southeastern Universities Research Association Trustee.
2002 - present - Reviewer, CLAS Collaboration.
2002 - Reviewer, Civilian Research and Development Foundation.
2002 - 2003 - American Physical Society Task Force on Countering Terrorism.
2000 - 2006 - Chair, Department of Physics.
2000 - Reviewer, US Department of Defense.
1999 - Reviewer, Department of Energy EPSCoR Program.
1997 - Chair, Jefferson Laboratory CLAS Collaboration nominating committee.
1996 - Chair, review panel, National Science Foundation, Instrumentation and Laboratory Improvement Program.

Honors
2003 University of Richmond Distinguished Educator Award.
Phi Beta Kappa, 1978.

Selected Listing of Refereed Publications

1. M. Battaglieri, R. De Vita, V. Kubarovsky, et al. (The CLAS Collaboration), 'Search for $\theta^+(1540)$ pentaquark in high statistics measurement of $\gamma p \rightarrow \overline{K}_0 K^+ n$ at CLAS', Physical Review Letters 96, 042001 (2006).
2. P. Rossi, et al. (The CLAS Collaboration), 'Onset of asymptotic scaling in deuteron photodisintegration', Physical Review Letters, 94 012301 (2005).
3. D. Protopopescu, et al. (The CLAS Collaboration), 'Survey of A'_{LT} asymmetries in semi-exclusive electron scattering on ^{4}He and ^{12}C', Nuclear Physics, A748, 357 (2005).
4. S. Stepanyan, *et al.* (The CLAS Collaboration), 'Observation of an Exotic $S = +1$ Baryon in Exclusive Photoproduction from the Deuteron', Physical Review Letters **91**, 252001 (2003).
5. K. Joo, *et al.* (The CLAS Collaboration), 'Measurement of Polarized Structure Function σ'_{LT}) for $p(\vec{e}, e'p)\pi^0$ from single π^0 electroproduction in the Delta resonance region', Physical Review C, Rapid Communications, **68**, 032201 (2003).
6. B. Mecking, *et al.*, (The CLAS Collaboration), 'The CEBAF Large Acceptance Spectrometer', Nucl. Instr. and Meth., **503**/3, 513 (2003).
7. G.P.Gilfoyle and J.A.Parmentola, 'Using Nuclear Materials to Prevent Nuclear Proliferation', Science and Global Security **9**, 81 (2001).
8. G.P.Gilfoyle, 'A New Teaching Approach to Quantum Mechanical Tunneling', Comp. Phys. Comm., **121-122**, 573 (1999).
9. G.P.Gilfoyle, M.S.Gordon, R.L.McGrath, G.Auger, J.M.Alexander, D.G.Kovar, M.F. Vineyard, C.Beck, D.J. Henderson, P.A.DeYoung, D.Kortering, 'Heavy Residue Production in the 215 MeV ^{16}O+^{27}Al Reaction', Phys. Rev., **C46**, 265(1992).

Selected Presentations

1. 'Measurements of the Fifth Structure Function of the Deuteron', CLAS Collaboration Meeting, March 3, 2006.
2. 'Nuclear Physics Working Group Report', plenary session CLAS Collaboration Meeting, June 18, 2005.
3. 'Picking Winners: Emerging Technologies for Countering Terrorism', colloquium presented at Virginia Commonwealth University, Richmond, Virginia, April 22, 2005.
4. 'Out-of-Plane Measurements of the Structure Functions of the Deuteron', plenary session of the CLAS Collaboration Meeting, November 13, 2003.
5. 'Maintenance and Upgrading of the Richmond Physics Supercomputing Cluster', V.Davda and G.P.Gilfoyle, Program and Abstracts for the Fall 2003 Meeting of the Division of Nuclear Physics of the American Physical Society, Tucson, AZ, Oct 30 - Nov 1, 2003.
6. 'Putting the Genie Back in the Bottle: Nuclear Non-proliferation in the New Millenium', Funsten Series Lecture presented at the Science Museum of Virginia, January 29, 2003.
7. 'A New Teaching Approach to Quantum Mechanical Tunneling', Presented at the Conference on Computational Physics 1998, September 2-5, 1998, Granada, Spain.

Emory F. Bunn

Physics Department
University of Richmond, VA 23173
ebunn@richmond.edu

816 St. Christopher's Rd.
Richmond, VA 23226
804-289-6486

Education:

Ph.D. Physics, University of California, Berkeley, 1995.
M.A. Physics, University of California, Berkeley, 1993.
A.B. Physics, Princeton University, 1989 (magna cum laude).

Dissertation:

"Statistical Analysis of Cosmic Microwave Background Anisotropy."
Dissertation Adviser: Professor Joseph Silk.

Professional Experience:

- **University of Richmond**, Physics Department. Assistant Professor. (2002-present)
- **St. Cloud State University**, Physics and Astronomy Department. Assistant Professor. (1999-2002)
- **Bates College**, Physics and Astronomy Department. Assistant Professor. (1996-1999)
- **U.C. Berkeley**. Postdoctoral Researcher. (1995–1996)

Selected Publications:

- A.L. Korotkov, J. Kim, G.S. Tucker, A. Gault, P. Hyland, S. Malu, P.T. Timbie, E.F. Bunn, E. Bierman, B. Keating, A. Murphy, C. O'Sullivan, P.A.R. Ade, C. Calderon, and L. Piccirillo, "The Millimeter-wave Bolometric Interferometer," Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III (J. Zmuidzinas *et al.*, eds.), Proc. SPIE, 6272, 62750X (2006).
- E.F. Bunn, "Systematic Errors in Cosmic Microwave Background Interferometry," submitted to Phys. Rev. D (2006).
- E.F. Bunn and M. White, "Mosaicking with Cosmic Microwave Background Interferometers," submitted to Astrophys. J. (2006).
- E.F. Bunn, "Probing the Universe on Gigaparsec Scales with Remote Cosmic Microwave Background Quadrupole Measurements," Phys. Rev. D, 73, 123517 (2006).
- J.C. Baez and E.F. Bunn, "The Meaning of Einstein's Equation," Am. J. Phys., 73, 653 (2005).

- E.F. Bunn, “Separating E from B,” New Astron. Rev., 47, 987 (2003).
- E.F. Bunn, M. Zaldarriaga, M. Tegmark, and A. de Oliveira-Costa, “E/B Decomposition of Finite Pixelized CMB Maps,” Phys. Rev. D, 67, 023501 (2003).
- E.F. Bunn, “Detectability of Microwave Background Polarization,” Phys. Rev. D, 65, 043003 (2002). See also erratum at *Physical Review D*, 66, 069902 (2002).

Funding:

- Co-I. on NASA Grant Proposal “Data Analysis for the Einstein Inflation Probe,” submitted to NASA under Research Announcement NNH06ZDA001N-BEFS ($250,000, pending).
- Sole P.I. on NSF Grant “RUI: Cosmic Microwave Background Analysis in the Post-WMAP Era,” 2005-2008 ($109,000 awarded).
- Co-I. on NASA Grant “Mission Concept Study for the Einstein Polarization Interferometer for Cosmology,” 2003-2005 ($200,000 awarded).
- Funding from NASA to support work on data analysis issues for the upcoming Planck satellite mission, 2001-2006 (approx. $10,000/year awarded).
- Sole P.I. on NSF Grant AST-0098048, “RUI: Cosmic Microwave Background Data Analysis,” 2001-2004 ($112,000 awarded).
- Sole P.I. on Research Corporation Cottrell Award CC5341, “Statistical Characterization of Foregrounds for Microwave Background Observations,” 2002-2005 ($34,000 awarded).

A Cluster Price Quote

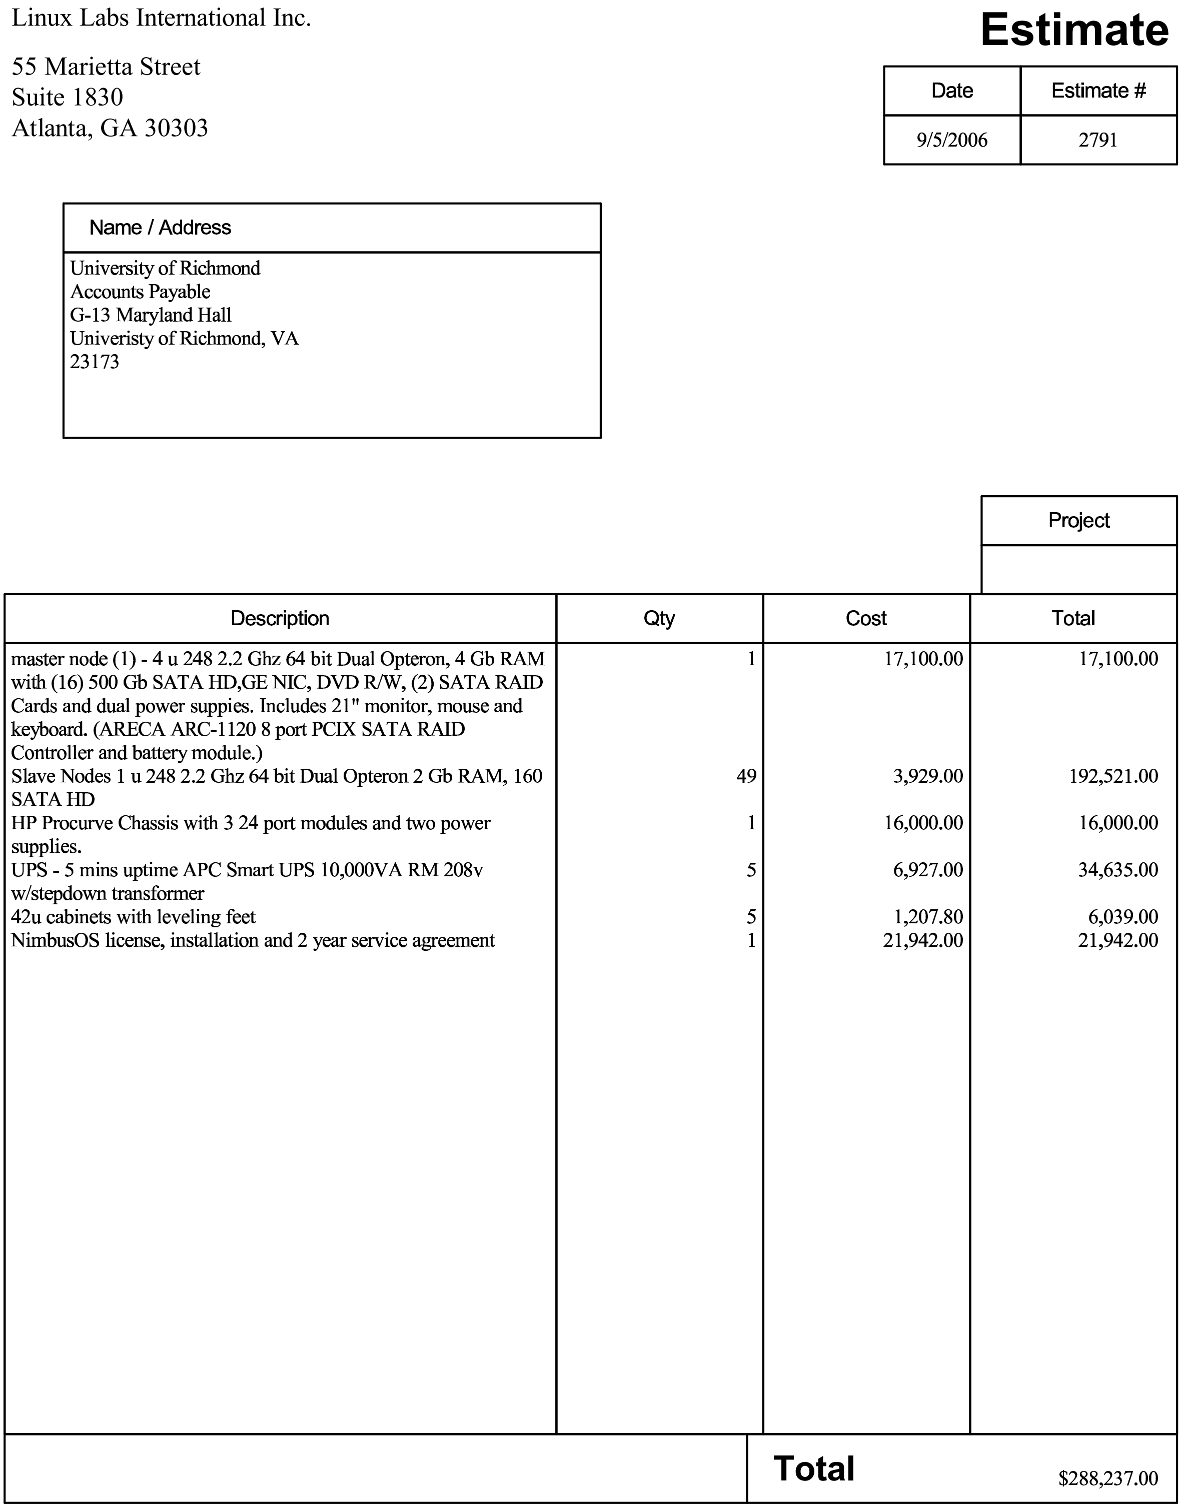

Linux Labs International Inc.

55 Marietta Street
Suite 1830
Atlanta, GA 30303

Estimate

Date	Estimate #
9/5/2006	2791

Name / Address
University of Richmond Accounts Payable G-13 Maryland Hall Univeristy of Richmond, VA 23173

Project

Description	Qty	Cost	Total
master node (1) - 4 u 248 2.2 Ghz 64 bit Dual Opteron, 4 Gb RAM with (16) 500 Gb SATA HD,GE NIC, DVD R/W, (2) SATA RAID Cards and dual power suppies. Includes 21" monitor, mouse and keyboard. (ARECA ARC-1120 8 port PCIX SATA RAID Controller and battery module.)	1	17,100.00	17,100.00
Slave Nodes 1 u 248 2.2 Ghz 64 bit Dual Opteron 2 Gb RAM, 160 SATA HD	49	3,929.00	192,521.00
HP Procurve Chassis with 3 24 port modules and two power supplies.	1	16,000.00	16,000.00
UPS - 5 mins uptime APC Smart UPS 10,000VA RM 208v w/stepdown transformer	5	6,927.00	34,635.00
42u cabinets with leveling feet	5	1,207.80	6,039.00
NimbusOS license, installation and 2 year service agreement	1	21,942.00	21,942.00
		Total	$288,237.00