Project Description

1 Introduction

Funds are requested from the National Science Foundation's Major Research Instrumentation (MRI) program to develop a computing cluster to support the research programs at the University of Richmond in nuclear physics, astrophysics, and computer science. The instrument will also be available to senior users at other institutions for work on closely related projects. The research groups at Richmond support 7-10 undergraduates during the summer and the academic year (Richmond is a primarily undergraduate institution). These students routinely go on to careers in science and engineering. All 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, Szajda and Lawson in computer science).

The nuclear physics research is centered on unraveling the structure of the nucleon and the nature of quark confinement and is supported by DOE grant DE-FG02-96ER40980. Additional senior users in nuclear physics will work on the instrument. 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. It is supported by NSF grant AST-0507395. Users at both Brown and Wisconsin will be heavily involved in the project. In computer science, co-PIs Szajda and Lawson study means of ensuring the integrity of computations executing in distributed volunteer computing platforms. The cluster will facilitate the large-scale simulations necessary to evaluate novel integrity and data privacy mechanisms for these platforms, and to evaluate new *hierarchical* volunteer computing platforms. The computer science work is currently supported by NSF Cyber Trust grant IIS-0524239. The joint nature of this proposal will enable the computer scientists to leverage the work of Gilfoyle and Bunn for additional application data traces to use in their own volunteer computing simulations.

We have considerable experience with high-performance computing. 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 at the end of its useful life, but the infrastructure to support it is still sound and the proposed system will benefit from that investment. Co-PIs Szajda and Lawson now use a 10-node network of Linux machines, inadequate for the work proposed here because the machines are shared by faculty and students for research and course work. Szajda and Lawson are well versed in maintaining Linux networks, and Lawson has significant research experience specifically related to supercomputing. The University is committed to software and hardware support of the proposed cluster after the lifetime of the grant. See Dean's letter of support in Appendix C.

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 quark confinement. Quantum Chromodynamics (QCD) is a highly successful description of quarks at high 4-momentum transfers or Q^2 [1], but at energies where the nucleons exist (the non-perturbative region), it is a daunting challenge to solve [2]. At low $Q^2 < 0.5$ $(GeV/c)^2$ the "hadronic" picture of nuclei (*i.e.*, nuclei made of protons and neutrons) has been successful [3]. However, the transition region between these extremes is poorly understood and mapping the geography of this transition is an essential goal of nuclear physics [4].

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. We work 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. The particle detection system consists of drift chambers [5], Cerenkov detectors [6], scintillation counters [7] for time-of-flight measurements, and electromagnetic calorimeters [8]. There are about 33,000 detecting elements capable of acquiring about 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 (especially for our undergraduates). 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 1-3 events per second.

2.1 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 in terms of four elastic form factors (electric and magnetic ones for each nucleon) that characterize the distributions of charge and magnetization within the proton and neutron [9]. The form factors are stringent tests of non-perturbative QCD including calculations on the lattice and are connected to generalized parton distributions that enable us to perform nucleon tomography [10, 11, 12, 13].

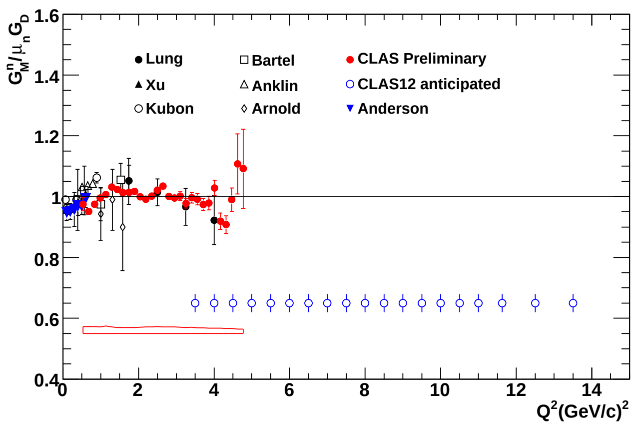

Figure 1: Preliminary results (red) for CLAS G^n_M measurement and expected data range and uncertainties (blue) for the CLAS12 G^n_M proposal.

We are part of a broad assault on the four elastic nucleon form factors at JLab [14]. Our focus is on G^n_M, the magnetic form factor of the neutron. To measure G^n_M we use the ratio of elastic $e-n$ to elastic $e-p$ scattering on deuterium which is less vulnerable to uncertainties than previous methods [15]. We have completed data collection and most of the analysis for a measurement of G^n_M in the range $\mathrm{Q}^2 = 0.2 - 5.0$ (GeV/c)^2 [15, 16, 17]. A report describing the analysis of two out of the three sets of running conditions is near the end of its Collaboration review. Preliminary results can be seen in Figure 1 along with some of the world data for G^n_M. Our group at Richmond has taken on the analysis of the third data set. We have also submitted a proposal to the JLab Program Advisory Committee (PAC) to make the same measurements at higher Q^2 as part of the JLab 12-GeV Upgrade.[1] Figure

[1] The DOE plans to upgrade the accelerator at JLab from a beam energy of 6 GeV to 12 GeV. The upgrade will

1 shows the anticipated data range and uncertainties. The proposal was approved by the PAC in August 2007 (spokesperson and contact person: Gilfoyle) [12, 19]. Completing the current analysis and preparing for the 12-GeV Upgrade will take advantage of the proposed cluster.

2.2 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 where there are few measurements [3, 20]. We need a baseline for the hadronic model so deviations at higher Q^2 can be attributed to quark-gluon effects with greater confidence [4, 21]. 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 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 beam helicity, and ϕ_{pq} is the azimuthal angle between the scattering plane (defined by the incoming and outgoing electron 3-momenta) and the reaction plane (defined by the 3-momentum transfer and the ejected proton momentum) [22, 23]. The unique, nearly-4π solid angle of CLAS creates an opportunity to extract the ϕ_{pq}-dependent structure functions f'_{LT}, f_{LT}, and f_{TT} in a model-independent way.

These structure functions are extracted using asymmetries that reduce our sensitivity to experimental effects. For example, the asymmetry $A'_{LT} = \rho'_{LT} f'_{LT}/(\rho_L f_L + \rho_T f_T)$ can be extracted using the $\sin\phi_{pq}$-weighted moments of the angular distributions measured with CLAS over the range $Q^2 = 0.2 - 5.0$ (GeV/c)^2 [22, 23]. We are studying the reaction in quasi-elastic kinematics first and later will investigate higher energy transfers. Our preliminary results for A'_{LT} show significant structure which is reproduced by a calculation from Jeschonnek and van Orden and disagrees with others [24, 25, 26]. The extraction and analysis of the other two structure functions (f_{LT} and f_{TT}) and investigations of different kinematic regimes are ongoing. This work is part of a CLAS Approved Analysis (spokesperson: Gilfoyle).[2] The analysis of these large data and the complex simulations of CLAS are computationally intensive and would use the proposed cluster. We note these data sets are the same ones used in the G_M^n analysis described in Section 2.1.

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 [4]. We have proposed a broad program of measurements to determine the mechanisms of confinement in forming systems. We use the nucleus as a "detector". Measuring the ratio of hadrons produced on nuclear targets relative to the production from deuterium will enable us to extract the lifetime of a deconfined quark after it is struck by a virtual photon and moving through the nucleus. The kinematic dependence of the transverse momentum broadening will enable us to measure the time interval required to form the hadronic color field around the struck quark. A proposal to do this experiment at high Q^2 as part of the 12-GeV Upgrade was approved by the JLab PAC in August, 2006 [27, 28]. 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.

---

require extensive changes to the accelerator and to CLAS to take advantage of the new physics opportunities. The JLab, 12-GeV-Upgrade is one of the highest priorities of the DOE office of Science in the next 20 years [18].

[2] A CLAS Collaboration member can write a proposal to analyze existing which is reviewed by a committee of Collaboration members, and defended before the Collaboration who then vote to approve the it.

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 with electrons [29] using a new method developed by Afanasev, *et al.* [30] and deuteron response functions calculated with the program DEEP [31]. The run time for these calculations 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 computing cluster.

We are 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 [32]. We will be responsible for design, prototyping, development, and testing of software for event simulation and reconstruction. The improved CLAS detector (CLAS12) will have prodigious software requirements and simulation is an essential aspect of the design and the eventual precision of CLAS12. The work will make significant use of the proposed cluster.

2.5 Additional Users 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 make better use of the proposed instrument. The number of students, undergraduate and graduate, that will learn high-performance computing will increase. Dr. M.F. Vineyard is a professor at Union College in Schenectady, NY (a primarily undergraduate institution like Richmond). and is a co-spokesperson on three approved experiments [16, 19, 33]. Dr. K.H. Hicks is a professor at Ohio University in Athens, OH. He typically has a postdoctoral fellow and 1-3 doctoral students in his group. He is co-spokesperson on three approved proposals at JLab [28, 34, 35] and two CLAS Approved Analyses [36, 37]. See the Appendix for letters of support from these users.

3 Astrophysics

The proposed astrophysics research 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 [38]. Polarimetry is the next frontier in CMB research, largely because polarization maps may provide a direct probe of an inflationary epoch in the extremely early Universe [39, 40]. A joint DoE/NASA/NSF Task Force recently advocated “a phased program to measure the large-scale CMB polarization signal expected from inflation” as its highest priority [41]. Similar opinions are expressed in the National Research Council’s decadal survey of astronomy and astrophysics [42] and their report *Connecting Quarks with the Cosmos* [43].

NASA’s *Beyond Einstein* road map for future astrophysics programs includes a dedicated CMB polarimeter known as the Einstein Inflation Probe (EIP) [44]. Co-PI Bunn is one of the leaders, along with Peter Timbie at Wisconsin and Gregory Tucker at Brown, of a NASA-funded Mission Concept Study for the EIP, the Einstein Polarization Interferometer for Cosmology (EPIC). With funding from NASA, 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 [45, 46]. 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. The research groups of Timbie and Tucker will be heavily involved in this research.

Any CMB polarization map can be divided into two components, a scalar E component and a pseudoscalar B component [39, 40], which probe different physical phenomena. Because it is insensitive to ordinary density perturbations (to linear order), the B component is predicted to be smaller than E by an order of magnitude or more over all angular scales. Detection of this weak B component is extremely important, because it provides a clean probe of other perturbations, especially the tensor perturbations produced by primordial gravitational waves. These tensor modes are predicted in inflationary models. If they are detected, we will direct evidence for inflation and will even be able to measure the energy scale and time of the inflationary epoch (perhaps as early as $\sim 10^{-33}$ s after the Big Bang). Aside from this tensor component, the dominant source of B polarization is expected to be gravitational lensing of E modes by large-scale structure [40] and will provide a valuable window on structure formation.

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. Although the B-type polarization signal is extremely weak, the next generation of experiments, including our proposed EPIC design, will have the raw sensitivity to detect it, by combining thousands of low-noise detectors with long integration times. The greater challenges are removal of non-cosmological foreground signals and control of systematic errors. The primary science goal of our proposed astrophysics research is to simulate interferometric CMB polarization observations to assess the effects of systematic errors.

3.1 Interferometric CMB polarimetry

In the past, both direct imaging telescopes and interferometers have been successful in CMB observations, and both technologies are 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. First, interferometers have reduced sensitivity to a variety of systematic errors [45, 46, 47]. MBI/EPIC will have simple, reflection-free optics, easily calculable, symmetric beam patterns with extremely low sidelobes, and no mechanical chopping of the telescope. Interferometers can measure the linear polarization Stokes parameters Q and U directly, without differencing signals from different detectors, mitigating leakage from the temperature anisotropy into the polarization channels which are 2-3 orders of magnitude weaker.

In addition, an interferometer can achieve higher resolution than is practical with a single dish. This is important for the EIP, as imperfections in the beam shape and pointing couple the CMB temperature anisotropy into the much weaker polarization signals [48]. These effects are mitigated with small beams, as the temperature is smooth on small scales. Separating the tensor and lensing contributions requires high-resolution maps [49].

Interferometry provides improved separation of the E and B polarization components [50]. In any incomplete sky map, there is "leakage" between the components [51, 52, 53], but inteferometry mitigates this leakage. Note 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 interferometer data consist of visibilities which have narrow window functions in Fourier space.

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 [47], 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 interferometers 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.*, [54]); 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 5.3.

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 [54], first 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 require real images. Traditional radio astronomy techniques such as the CLEAN algorithm [56] are well-suited for data sets with sharp features, not the diffuse nearly-Gaussian structure found in CMB maps. We plan instead to use maximum-entropy reconstruction, which has been well-developed in the CMB context [57]. Because maximum-entropy is a nonlinear method, the noise properties of the resulting maps can be computed only via simulation.

Component separation. The simulations described above 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. 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

[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 [55].

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.

3.3 Additional Users in Astrophysics

Co-PI Bunn will lead the computational astrophysics research. The key ingredients for the simulations have been developed in Timbie and Tucker's groups and elsewhere. Bunn and Richmond undergraduates will synthesize these elements into a working simulation pipeline. Timbie and Tucker's research groups will provide support, but the bulk of the work will be done at Richmond. We hope eventually to hire a postdoc to work full-time on the component separation problem. If we are able to do this, the postdoc would probably be based at Richmond and work closely with undergraduates but would also spend significant time at Brown and Wisconsin.

4 Computer Science

In a distributed volunteer computing platform, personal computers connected to the Internet volunteer otherwise idle processor cycles to assist in the execution of a large-scale distributed computation. In a typical setting, the supervisor of a distributed volunteer computing application recruits participants (volunteers). Each participant agrees to allow the supervisor to execute code on the participant's computer. Each participant then downloads code that serves as the local execution environment for assigned computational tasks which are typically small enough to be executed within a few hours by the participant. For a given computation, participants are chosen, tasks are assigned and transmitted, and, as tasks are completed, significant results are collected by the supervisor. Though task results may be related, the tasks themselves are independent, so no communication is necessary between participants.

Important application domains benefiting from these platforms include protein structure prediction (*e.g.*, Folding@Home [58], Rosetta@home [59]), quantum chemistry [60], physics [61], biomedical applications (*e.g.*, fightAIDS@home [62]), climate prediction [63, 64], and cryptanalysis [65]. Endeavors of a more academic nature include various mathematics applications (*e.g.*, prime searches [66, 67], graph theory [68]) and SETI@Home [69]. In addition, Apple's Xgrid [70] and the IBM-powered World Community Grid [71] are examples of commercial entities recognizing the importance of and supporting such volunteer platforms.

Distributed volunteer computing platforms facilitate computations once feasible only via expensive supercomputers. In recent years, state-of-the-art supercomputers have become increasingly based on tightly coupled cluster architectures, and these architectures comprise a majority of the most recent Top 500 list of supercomputer sites [72]. Nevertheless, distributed volunteer computing platforms remain a viable and relevant alternative to cluster-based supercomputers, as evidenced by the success of examples listed above. Distributed volunteer computing platforms are especially attractive for a variety of reasons: otherwise untapped resources are efficiently used; the cost of purchasing an expensive supercomputer, or the necessity of negotiating available time on an existing supercomputer, is eliminated; supporting costs (*e.g.*, power, cooling, and maintenance) of a supercomputer are essentially donated by participants.

Although these platforms provide an attractive source of computing power, providing assurance of results is difficult. The *integrity* of computations executed on these distributed computing platforms is a serious issue because the execution environment of each participant is, in general, outside the control of the computation supervisor. The intentional or unintentional corruption of results by participants is possible, even likely. Currently implemented platforms often attempt to address integrity by using *simple redundancy* (*i.e.*, assigning two copies of the same tasks) and/or by using code obfuscation. Neither solution is satisfactory. Code obfuscation is difficult in practice, and in some cases, it is impossible to implement. Simple redundancy, though easily implemented, is far from ideal as currently applied. Simple redundancy is expensive (at the very least it doubles the processing cost of the computation) and is easily subjected to collusion and denial-of-service attacks. Denial-of-service attacks are easily executed because if a result is not returned within a specified amount of time, then the supervisor will consider the task lost and will reassign. Furthermore, current implementations invite the threat of collusion. Few systems limit the number of user names an individual can collect or the number of outstanding tasks possessed per user. In principle, using this lack of limits, a single individual can easily obtain several thousand outstanding tasks. For example, the SETI@home project has experienced days during which more than 5000 new user names were assigned, and it boasts participants who have *averaged* more than 1000 tasks completed per day in the nine years since the project began.

Co-PIs Lawson and Szajda currently head a computer security research group at the University of Richmond (funded via NSF Cyber Trust grant number IIS-0524239) that is investigating protocols and algorithms for protecting the computation integrity of applications executed on distributed volunteer computing platforms. At its core, this work addresses the organization and coordination of untrusted components such that a computation succeeds whenever possible, and fails gracefully (*i.e.*, compromised results are detected) when not. The goal of this group is two-fold. First, to provide better solutions for computation integrity in existing platforms, which typically employ a strict client-server model. Second, to develop novel algorithms and protocols for ensuring computation integrity in future hierarchical platforms, which will be necessary due to the increasing number of participants of varying power, bandwidth, and storage characteristics (i.e., a *much* more heterogeneous participant pool). The long-term goal is to develop a "self-correcting" distributed volunteer computing platform, in which the platform (given sufficient information about the application and participant pool) will be able to detect and correct errors that result from a non-trivial amount of Byzantine activity.

The group's current efforts would benefit significantly from the use of a reliable high-performance cluster. Among these are experiments designed to investigate the application of hierarchical platforms to distributed volunteer computing, experiments to provide statistical characterization of different platforms and applications, and tests of techniques for providing enhanced privacy for certain applications involving sensitive data. One technique currently being used by the group for examining novel hierarchical architectures involves the use of genetic algorithms (GAs) for exploring the space of potential solutions. In an earlier group effort [73], mathematical analysis was used to find a task distribution strategy (based on assigning redundant work units) that was *practically optimal* in the sense that any other practical static redundancy-based distribution strategy that guaranteed a given adversary detection threshold would require more total processor cycles. Subsequent study of the same problem using genetic algorithm techniques revealed other distribution strategies that provided comparable protection under assumptions on the maximum proportion of malicious participants. Current efforts using similar techniques to examine hierarchical architectures have proved promising, but are computationally expensive. See Section 5.4 for a discussion of computing needs.

The security group is also interested in the proposed cluster because of its ability to facilitate new types of simulation work. As a specific example, the group is interested in using and modifying SimBA [74] (Simulator of BOINC Applications), a recent NSF-funded simulation framework for evaluating the performance of various task distribution strategies (e.g., [75, 76]) in volunteer computing platforms. An important aspect of the security work relies on intelligent distribution of tasks to meet predefined security and computation integrity goals. Using SimBA on the proposed cluster will allow the co-PIs to efficiently and accurately evaluate their own task distribution strategies, including strategies already developed (e.g., [73, 77]) and novel strategies (such as those determined by the recent GA work mentioned above), in a more realistic but carefully controlled setting. Another important part of the security research is the investigation of hierarchical platforms, as mentioned earlier. The proposed cluster will allow the co-PIs to easily modify SimBA to incorporate novel hierarchical platforms and execute large-scale tests of those platforms in a simulation context.

Furthermore, note that SimBA does not actually execute the associated distributed-computing applications, but instead uses execution traces from existing distributed-computing applications (*e.g.*, Predictor@Home). Because the applications of co-PIs Gilfoyle and Bunn can execute in a volunteer computing platform, co-PIs Lawson and Szajda will use the cluster experiments of Gilfoyle and Bunn to provide additional data traces for SimBA. We will use traces of the physics calculations and analyses to extrapolate and obtain additional trace inputs (*i.e.*, to simulate additional applications) for our own volunteer computing platform simulations on the cluster.

This project is fundamentally about developing efficient techniques for organizing and coordinating untrusted components such that the systems which they comprise produce measurably good outcomes. The potential interdisciplinary benefits are substantial — the list of applications that stand to benefit from our proposed advancement of these platforms include several important scientific applications. Sensitive genetic sequence comparisons and protein folding, for example, can potentially alter the way disease is treated by contributing to the development of new genetic therapies and powerful person-specific drugs. The potential impact extends to the realm of applications that have previously been dismissed as infeasible. Distributed volunteer computing platforms are self-upgrading systems — as the performance of personal computers and baseline devices grows, the power and speed of these platforms increases proportionally. As processor speeds continue to increase, the range of applications that become will practical is substantial. To realize this potential, however, the broad goal of providing provable security must be achieved. The work proposed here is an important step in that direction.

5 Research Instrumentation Needs

We request in this proposal funds for the purchase of a cluster of 28, dual-quad-core CPU computers supported by 6 terabytes 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 high-performance computing.

5.1 Current Computing Facilities

The current facilities in physics include a computing cluster developed in 2001 with NSF and University funds plus an array of workstations. The original system consisted of 49 remote machines (1.4 GHz) running the Linux operating system, a master node, 3 TByte of RAID storage, and 18 GByte of disk space on each node. 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 is now at the end of its useful life and only about 24 nodes still work and 2 out of 3 RAIDs. This is not surprising since the system is almost seven years old. At JLab, for example, replacement of failing nodes begins after two years and all machines are replaced after four years of use. The University employs a linux support expert who is responsible for software updates and administration.

The Department of Mathematics and Computer Science maintains a network of ten Pentium-based workstations running Linux (the primary platform for the simulation experiments performed by co-PIs Lawson and Szajda) with external NFS disk sharing. These workstations are shared by department faculty and students for both research and course work. Co-PIs Lawson and Szajda also maintain four additional workstations for undergraduate research assistants.

5.2 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 2.2. 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 CLAS simulation software is about 1-2 Hz on each remote node. With our remaining 24 nodes the current cluster will take about 20 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 (see Section 5.5). The JLab facilities are heavily subscribed, and our existing cluster is aging and falling short. 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 also note the astrophysics and computer science groups do not have alternative clusters to use.

We have learned several lessons from our previous experience. A major bottleneck in our data analysis is the speed of the switch (about 100 Mbps). Second-pass analysis requires us to copy the data from the RAID out to the remote nodes. We have to slow the analysis rate so the switch can keep up making it difficult for multiple users to take full advantage of the system. The switch in the proposed system will be at least 10 times faster. See Section 5.5.

The disk needs are large. We currently use 1.5 TBytes out of the remaining 2.0 TBytes available for nuclear physics. Adding the astrophysics and computer science users will only increase the demand for storage. On the remote nodes 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. We need more and faster remote nodes, a faster switch, and adequate storage.

5.3 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 $N_v \simeq N_{pix} \simeq 1 \times 10^4$.

The memory requirements for the astrophysics simulations will easily 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 4 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.

The time required for the simulations described in Section 3.2 are dominated by two steps: power spectrum estimation and maximum-entropy image reconstruction.

Power spectrum estimation. Naive implementations of this step would involve maximizing

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

the likelihood over a multidimensional parameter space, with each likelihood calculation scaling as $O(N_v^3)$. Fortunately, we can take advantage of the sparseness of the visibility-space covariance matrix [79] and use Monte Carlo Markov chains (MCMC) to replace a search of the entire likelihood parameter space [78]. Scaling the results of [79] to our parameters, we estimate that we can evaluate a single likelihood in about 1 second. A typical MCMC, requiring $\sim 10^3$ likelihoods, will therefore take roughly 20 min.

Maximum Entropy Image Reconstruction. The scaling properties of CMB maximum-entropy algorithms to the large data sets considered here are not well-known. As it happens, though, the limiting step in the algorithm (calculation of χ^2) involves manipulating a large sparse matrix quite similar to that required for a single likelihood evaluation of the power spectrum. 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 of order 1 hour.

The above estimates are for a full-scale production run of the simulations. In earlier testing stages, we can speed up the process in various ways (forgoing maximum-entropy reconstruction, working in a restricted likelihood parameter space, and reducing N_{tod}). With this streamlining, we can perform $\sim 10^3$ simulations (enough to draw reasonable statistical conclusions) in a few days. This is a reasonable turnaround time for testing purposes. For full-scale production runs, we estimate that we can perform $\sim 10^2$ simulations per week on the proposed cluster.

5.4 Computer Science Computing Needs

A substantial amount of the security group's research efforts rely on detailed simulation experiments. In previous related work by the group, simulation experiments were distributed across two, public, Windows-based labs with about 15 workstations each (the current 10-node Linux lab was not in existence at that time). None of the machines are dedicated to research and are volatile - subject to resource sharing with students and other faculty and to potential reboot. On user logout the Windows-based machines automatically reboot and revert to a standard image for security reasons. The only option is to start the simulation experiments each night at the close of lab and stop them when the labs reopen. A two-week simulation could be done in less than two hours on the faster, dedicated instrument.

Techniques for enhancing the data privacy of distributed volunteer computations require a great deal of simulation work. Our 2006 paper on privacy enhancements for the Smith-Waterman class of genetic sequence comparison algorithms [80] provided sequence modification techniques that allowed useful sequence comparisons without having to divulge the actual sequence data. We generated data from over 900 simulation runs, each of which took at least six hours and often as much as 2-3 days on a single 1GHz Pentium 4 PC. During peak processing periods, we were able to generate a maximum of around 60 GFLOPS using 10 similar machines (2 GFLOPS for a 1 GHz Pentium 4 [81]). The proposed cluster will provide a factor of 50 improvement in turnaround time. The instrument is much faster (see Section 5.5) and it can be used all day. This rapid turnaround time is essential for undergraduate involvement since they are limited in time and need quick feedback (see Section 6).

5.5 Proposed System

We now describe the proposed system. The components are listed in Table 1. A detailed quote for items 1-5 from the vendor LinuxLabs is in Appendix A. Below we discuss our reasoning behind the choice of the different components. The University is committed to providing $10,000 per year after the lifetime of the proposed grant to support the instrument. See Appendix C.

The dual, quad-core Xeon processors (item 1) were chosen because of their excellent cost-to-benefit ratio. Their clock speed is about 50% faster than the speed of the current remote nodes,

Item	Number	Description	Price($)
1	1	Dual, Quad-Core Xeon master node, 4 GByte RAM, 6 TByte RAID	11,400
2	27	Remote nodes, Dual, Quad-Core Xeon, 4 GByte RAM, 150 GByte storage	210,600
3	1	HP Procurve switch	4,460
4	5	cabinets	3,300
5	1	Nimbus OS license, installation, warranty, and shipping	28,900
6	-	Hardware items that cost less than $500	1,100
		Total Cost	259,760

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

and architectural improvements make them overall about 10 times faster than the current ones (6.8 gigaflops versus 0.7 gigaflops). Test results from the vendor (LinuxLabs) using LinPack (a benchmarking standard for solving large, square matrices) give an equivalent speed of 1.5 teraflops for the entire instrument. This is about 20 times the speed of the original Richmond cluster and 40 times the speed of the existing, depleted cluster (see Section 5.1). The Linux operating system will be used. It is a research-quality operating system commonplace in physics and computer science. 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, 24-node, cluster (about 20 days) for such a calculation. Calculations of this length are routine for the analysis of the CLAS data. The astrophysics projects require hundreds of power spectrum estimations and maximum entry reconstructions which each take about a CPU-hour on the current nodes. The memory (4 GByte for the 8 cores on 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 150-GByte hard drive (item 2) will be attached to each machine to provide storage. This space is needed to store data files for analysis, 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). Cabinets will hold the nodes (item 4). Hardware and software installation is required (item 5). 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 6. 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 4-6 years is the optimum lifetime (see Section 5.1).

6 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 users. 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 6-7 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 DOE (two grants), NSF (two grants), NASA (one grant), and Research Corporation (one grant). The Computer Science faculty consist of seven tenured or tenure-track professors and graduate about 7-9 majors each year. There are active research programs in security of distributed volunteer computing, computer organization, performance evaluation of high-performance systems, artificial intelligence, discrete-event simulation, data mining, and programming languages. The first two programs above are supported by two NSF grants. We all emphasize undergraduate involvement in research from early in the students' careers. Students involved in undergraduate research are more likely to attend graduate school and be successful after college [82].

We have been successful at starting our undergraduates on their research careers. Over the last two summers about 42 of our physics and computer science majors participated in research at Richmond, JLab, Yale University, Berkeley, and the University of Notre Dame. Our students have accomplished much in their research careers. Many in physics have presented their work at national meetings and some have even obtained travel support to conferences from the American Physical Society [83, 84, 85, 86, 87, 88, 89, 90, 91]. In computer science previous research related to this proposal has lead to nine independent study courses, two seminars and four honors theses. We have also been successful in attracting under-represented groups to work in our nuclear physics group. Of the six 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, astrophysics, and computer science groups involve 7-10 undergraduates in research every summer. In physics they are integral parts of our research program. 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 learn modern high-performance computing methods by using our existing cluster. We have also recruited computer science students to help with administration of the cluster; from setting up firewalls to updating the CLAS software. They have presented their work at national meetings [83, 84, 85, 86, 87, 88, 89, 90, 91]. It is worth noting that for our undergraduates an on-campus cluster is essential for rapid turnaround times. The computing farm at JLab is heavily used, and it can take a day or more to get a batch job started. Such a long turnaround time is a barrier to learning and productivity especially for these young people that an on-campus system will overcome.

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.

The computer science undergraduates are active participants in the research process, keeping abreast of recent developments through supervised reading of recent conference proceedings, by coding applications and simulations, by assisting in the managing and deployment of a local volunteer computing platform, by designing and running experiments, and by writing up results for publication. With the availability of a high-performance cluster, our research students will gain invaluable experience: helping to administer the cluster; writing simulation code to leverage the tightly-coupled nature of the cluster; and designing and executing large-scale experiments that would otherwise be infeasible using our existing, small Linux network.

In the classroom, previous work related to this proposal has led to enhancements to lower-level CS courses (Introduction to Computing, Data Structures, and Discrete Structures), upper-level CS courses (Algorithms, Security, Simulation, and Probability), and an interdisciplinary course in Bioinformatics taught by co-PI Lawson and Dr. Joe Gindhart from the Department of Biology. Future computer science and interdisciplinary course development and enhancement related to this proposal will ensue. This work will attract not only computer science majors, but also students from biology, chemistry, mathematics.

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 undergraduates working on data analysis and simulation issues.

7 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, astrophysics, and computer science groups to operate and maintain the proposed computer cluster. One of us (Gilfoyle) is responsible for maintaining the existing systems, two of us (Szajda and Lawson) assist in maintaining a 10-node Linux cluster of workstations, and we all have significant experience with Linux operating system. All faculty in our groups have considerable software experience in general and with the applications described above. The University administration has adequately supported our research efforts in the past and is committed to continuing to support the University's technology infrastructure (see Appendix C). 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 5 cluster-hours/week to over 100 cluster-hours/week if many simulations of the CLAS detector are required. The average is around 10-15 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. 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 vendor provides easy-to-use tools for allocating resources on the cluster. Software tools that are now routinely available from vendors will make management and software maintenance easier even for a 'large' cluster. Furthermore, co-PI Lawson has significant research experience in scheduling for high-performance computing systems that may be leveraged.

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.

8 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. Computer science students will be co-authors on articles submitted for publication in competitive, peer-reviewed conference proceedings. In addition, co-PIs Lawson and Szajda regularly participate in computer science teaching-related conferences, such as the ACM Special Interest Group on Computer Science Education (SIGCSE) Technical Symposium on Computer Science Education, and will disseminate teaching materials associated with this work.

9 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. Gilfoyle was co-PI on that grant entitled "MRI: 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. All of the Richmond nuclear physics work described here has made heavy use of the cluster (see Section 2) with the other JLab users [17, 29, 92, 93, 94]. All the Richmond students used the cluster [83, 84, 85, 86, 87, 88, 89, 90, 91].

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 work on two projects relevant to this proposal: analyzing systematic errors in CMB interferometry [47] and analyzing mosaicked interferometric observations [55]. It has also supported Bunn's work on dust contamination in CMB maps, on Sunyav-Zel'dovich surveys [95], an analysis of the astrophysical constraints on $f(R)$ theories of gravity [96], and a search for directionality in the WMAP data [97]. This grant has led to five refereed publications (Astrophys. J., Phys. Rev. D, and Monthly Notices of the RAS) [47, 55, 95, 96, 97] and three contributed presentations at national meetings of the American Astronomical Society, including one presentation by undergraduate Larson [98, 99, 100].

Co-PIs Lawson and Szajda are currently supported by NSF Cyber Trust grant IIS-0524239. To date, this grant has been used to support 11 undergraduates in related research projects, and has led to two refereed conference proceedings publications [80, 101] (one with undergraduate Michael Pohl as second author), one refereed workshop proceedings publication [102] (with undergraduate Ed Kenney as a co-author), and one non-refereed conference proceedings publication [103].