Statement of Work: Fermi and IIT

Common Runtime Environment and Workflow Management

In spite of years study, a suitable and broadly applicable supercomputing solution has been elusive. This is especially true for Lattice calculation, where computing is large in scale often requiring a distributed computing over different computer clusters, temporary and input/output data are huge and cross a large number of file systems, and concurrent coordination and correlation is frequent for runtime analysis. The lattice community is distributed widely geographically, as well as its computing resources and data storages. Automating coordination between these resources is highly nontrivial. Conventional task scheduling systems for high performance computing, such as Maui, do not support cross resources coordination. Grid computing technologies can help. However, Grid computing is still in its infancy. A general Grid solution is not expected to be available in the near future. Also, lattice calculation has its own characteristics and requirements. Even a general Grid solution would be available, an in-depth understanding of how to tailor a general Grid solution for lattice calculation is necessary. In this study, we propose to adopt and extend current Grid technologies, especially the workflow technologies, to develop a common runtime environment for lattice calculation.

Grid computing is designed to coordinate computing resources across different virtual organizations to solve large scale, nontrivial applications. The Lattice calculation involves multiple data and computing resources combined in several mutually interdependent steps that form an analysis campaign. The specification of its steps and dependencies is known as a workflow. A workflow can be expressed as a directed acyclic graph. Grid technologies envisage that resources advertise their capabilities and that the Grid middleware can act as a broker between jobs with specified resource requirements and resources with specified capabilities. Grid technology is also available to monitor task execution and resource availability for efficiency and fault tolerance; schedule tasks execution in such a way to minimize communication and overlapping computing and data access; and execute a remote task over different local middleware and management systems.

Lattice QCD jobs have domain specific features, which make them morphable onto various sized parallel machines and fit well for Grid computing. On the other hand, however, given the complexity of current lattice QCD analysis campaigns, that can involve hundreds of input files and thousands of intermediate and final files, new tools need to be developed to manage the workflow closely, and to automate many aspects of executing analysis campaigns. We propose to develop these new tools based on Grid technologies and develop a common lattice calculation runtime environment that integrate these new tools with existing lattice computing environments. While this development is challenge and requires a combined expertise in both physics and computer science, it has a lasting impact and many `point' solutions are available. We will develop the common runtime environment based on our current success of the Grid Harvest Service (GHS) system for performance prediction and task scheduling of Grid computing (see www.cs.iit.edu/~scs/software.htm). GHS has a performance measurement

and prediction subsystem, a task allocation subsystem, a task scheduling subsystem, and a task execution subsystem. It can coordinate and execute tasks over different machines and middleware. GHS provides the technical foundation for the proposed lattice common runtime environment.

We plan to conduct the proposed system development step by step. The software developed will be made available to the lattice community at each development phase. Continued modification and optimization will be conducted based on user feedbacks. We plan to fully achieve our objectives:

- Gather requirements and describe common usage scenarios for the workflow management system based on experiences with the current LQCD software structure.

- Define a XML schema to support the description of workflow structures. Currently GHS uses XML files to describe the information of resources, files, and tasks.

- Apply existing XML parsing tools (e.g. JDOM and dom4j) into the validation of XML documents to aid the writing and validation of workflow.

- Combine our work and other workflow systems such as DAGMan, Pegasus, and Karajan to develop a unified workflow specification and visualization management system.

- Develop an execution system that transforms the workflow specifications into instructions and allocates and schedules tasks accordingly over various computing environments.

- Deploy a performance monitoring and scheduling system to monitor the progress of work being carried out within a workflow, monitor the system behaviors, and initiate a rescheduling if a fault or other abnormal behaviors occurs.

- Conduct lifetime management of temporary and intermediate data in order to maximize resource utilization (e.g. disk space, network bandwidth, memory, CPU), and adopt a workflow-based data prefetching to overlap computing and communication to improve the computing efficiency.

We plan to have an initial runtime environment available within two years, which can monitor and allocate lattice QCD jobs cross multiple computing domains, and support meta-task scheduling. We will deliver a fully functional, workflow supported runtime environment at the end of this research.