

# Studying Protein Folding on the Grid: Experiences using CHARMM on NPACI Resources under Legion

Anand Natrajan, Michael Crowley, Nancy Wilkins-Diehr, Marty A. Humphrey, Anthony D. Fox, Andrew S. Grimshaw, Charles L. Brooks III

**Abstract** — One of the benefits of a computational grid is the ability to run high-performance applications over distributed resources simply and securely. We demonstrated this benefit with an experiment in which we studied the protein-folding process with the CHARMM molecular simulation package over a grid managed by a grid operating system, Legion. High-performance applications can take advantage of grid resources if the grid operating system provides both low-level functionality as well as high-level services. We describe in detail the nature of services provided by Legion for high-performance applications. Our experiences indicate that human factors continue to play a crucial role in the configuration of grid resources, underlying resources can be problematic, grid services must tolerate underlying problems or inform the user, and high-level services must continue to evolve to meet user requirements. Our experiment not only helped a scientist perform an important study, but also showed the viability of an integrated approach such as Legion's for managing a grid.

## I. INTRODUCTION

As available computing power increases because of faster processors and faster networking, computational scientists are attempting to solve problems that were considered infeasible until recently. Computational grids are becoming more pervasive platforms for running distributed jobs to solve such problems. A *computational grid* or a *grid* is a collection of distributed resources connected by a network. In such an environment, users, such as scientists, can access

Manuscript received April 10, 2001. This work supported in part by DARPA (Navy) contract #N66001-96-C-8527, DOE contract DE-FD02-96ER25290, DOE contract Sandia LD-9391, DOE contract D459000-16-3C, DARPA contract SC H607305A, Logicon (for the DoD HPCMOD/PET program through the NAVO MSRC) contract DAHC 94-96-C-0008, National Science Foundation Next Generation Software grant EIA-9974968, National Science Foundation NPACI grant ASC-96-10920, and a grant from NASA-IPG.

Anand Natrajan is with the Dept. of Computer Science at the University of Virginia, Charlottesville, VA 22904-4740, USA (e-mail: anand@virginia.edu).

Michael Crowley is with the Dept. of Molecular Biology at The Scripps Research Institute, La Jolla, CA 92037, USA (e-mail: crowley@scripps.edu).

Nancy Wilkins-Diehr is with the San Diego Supercomputing Center at the University of California at San Diego, La Jolla, CA 92037, USA (e-mail: wilkinsn@sdsu.edu).

Marty A. Humphrey is with the Dept. of Computer Science at the University of Virginia, Charlottesville, VA 22904-4740, USA (e-mail: humphrey@cs.virginia.edu).

Anthony D. Fox is with the Dept. of Computer Science at the University of Virginia, Charlottesville, VA 22904-4740, USA (e-mail: adf5j@virginia.edu).

Andrew S. Grimshaw is with the Dept. of Computer Science at the University of Virginia, Charlottesville, VA 22904-4740, USA (e-mail: grimshaw@virginia.edu).

Charles L. Brooks III is with the Dept. of Molecular Biology at The Scripps Research Institute, La Jolla, CA 92037, USA (e-mail: brooks@scripps.edu).

resources in a transparent and secure manner. When a user submits a job in a grid, the system runs the job on distributed resources and enables the user to access the results of the jobs as they execute and when they complete. Developers of grid infrastructures such as Legion [4] must conduct and present detailed case studies showing how users access grids. We present one such case study.

Legion is a grid operating system. It provides the services of a traditional operating system — process creation and control, interprocess communication, file system, security and resource management — on a grid. In other words, Legion abstracts the distributed, heterogeneous and potentially faulty resources of a grid by presenting users with the illusion of a single virtual machine [5]. In order to achieve this goal, Legion manages complexity in a number of dimensions. For example, it masks the complexity involved in running on machines with different architectures and different operating systems, managed by different software systems, owned by different organisations and located at multiple sites. In addition, Legion provides a user with high-level services in the form of tools for specifying what an application requires and accessing available resources.

In our experiment, a computational scientist accessed resources from NSF's National Partnership for Advanced Computational Infrastructure (NPACI) using the grid infrastructure provided by Legion. The application used was CHARMM (Chemistry at HARvard Molecular Mechanics) [3] [9], a popular general simulation package used by molecular biologists to study protein and nucleic acid structure and function. One large problem for which CHARMM is used is the study of the nature of the protein folding process. The scientist desired to study the energy and entropy of many folded and unfolded states of a certain protein, Protein L, to gather information about the behaviour of the protein during its folding process and to generate a protein-folding landscape. This study required multiple CHARMM jobs to be run with different initial parameters.

There were two clear goals for this experiment:

1. *Enhance the productivity of the user by solving a large and computationally challenging problem.* By accessing distributed grid resources, the user condensed the time required for performing his computations from a month (if he used the resources available at his organisation) to less than two days.

2. *Demonstrate a match between the mechanisms expected by the user and those provided by the grid infrastructure.* The user had to learn five commands or fewer in order to perform his computations on a variety of resources.

In the process of meeting these goals, we made a number of observations that affect grid infrastructure developers as well as grid users. In this paper, we present those observations in the context of the experiment.

Our primary observation was that the grid infrastructure must provide high-level services in addition to low-level functionality. Providing low-level functionality alone is not enough; without high-level services built on top of the underlying infrastructure, a user's productivity can fall tremendously. The novelty of this experiment is not that a large problem is being solved (the experiment continues to run at the time of writing), but the ease with which the user accessed grid resources and the low cognitive burden imposed on him by the grid infrastructure, Legion. In this paper, we describe how the user interacted with a grid using Legion services, what problems arose with the NPACI resources that were part of the grid and how Legion addressed those problems, and what lessons we learned regarding new functionality that can be provided to users.

In Section II, we present CHARMM in some detail. The purpose of the discussion is not so much to describe the scientific problem being solved (fascinating though that is), but to describe the characteristics of the application that make it desirable for running on a grid. In Section III, we describe Legion, especially in terms of the features that make it attractive for running high-performance applications. In Section IV, we show how the user interacted with Legion in order to run his jobs on the grid. In Section V, we present the results of our experiment. Also, we list the successes of our work, explain problems encountered and identify future directions. We conclude in Section VI.

## II. CHARMM

The process of protein folding is not well understood and the state-of-the-art methods of studying it are so computationally intensive that they are often not undertaken. One method is to calculate the free energy surface of the folding process. The calculation is designed to reveal the process by which a small protein (Protein L) folds up into its normal, three-dimensional configuration. The folding process occurs in nature every time a protein molecule is manufactured within a cell. The biophysics of folding must be understood in detail before the information can be used in developing ways of interacting with proteins to cure such diseases as Alzheimer's or cystic fibrosis.

The CHARMM molecular simulation package uses the CHARMM force field to model the energetics, forces and dynamics of biological molecules using the classical method of integrating Newton's equations of motion. Typical systems studied involve protein or nucleic acid molecules of several hundred to several thousand atoms and a bath of solvent, usually water, consisting of many thousands of molecules, for a total of 20000 to 150000 atoms. All chemical bonds and all interactions that do not involve bonds (for example, electrostatics) are used to model the system. These interactions

number in the millions to billions. For a typical simulation, hundreds of thousands to millions of timesteps of integration are required and, at each timestep, all interactions are determined. In a parallel run, all forces and all coordinates must be shared among all processors.

CHARMM is computation- as well as communication-intensive. In a single CHARMM job, hundreds of processes may communicate with one other while performing their computations. The processes communicate using Message Passing Interface (MPI), a standard for writing parallel programs [6] [10]. The parallel efficiency of the computation depends on the number and speed of the processors, and the speed and latency of the interconnect. On current-generation high-performance computers, processor speed has increased but interconnect speed has lagged. Therefore, CHARMM's performance degrades rapidly after 32 processors on almost all architectures except the T3E, on which it scales well to 128 processors. Therefore, we chose to run with 16 processors on most architectures, getting better than 95% parallel efficiency throughout the experiment on all high-performance architectures. For a 16-processor run, all processors communicate about 1 Mbyte of data at every timestep in a couple of all-to-all communications, and another 4 Mbyte in each-to-each communications. For a typical 16-processor run on a 375MHz Power3, approximately 3 timesteps occur per second. Each job requires a number of input files, some of which are a few Mbytes large, and generates a number of output files some of which are hundreds of Mbytes large. Thus, a single job requires powerful computation resources, fast network capabilities and fairly large amounts of disk space. In our experiment, the user required multiple (up to 400) CHARMM jobs to be run.

We decided to run the CHARMM jobs on a computational grid because the total amount of computing resources required made it unattractive to run at a single site. Typically though not necessarily, supercomputing centres such as the San Diego Supercomputing Center (SDSC) use queuing systems to control powerful computation resources connected by fast networks. Since such resources are exactly what CHARMM jobs require, our experiment was conducted on queuing systems. Nothing in CHARMM requires a queuing system; our choice of resources was governed by the coincidence that the kinds of resources that CHARMM requires are usually controlled by queues.

## III. LEGION

The Legion project is an architecture for designing and building system services that present users the illusion of a single virtual machine [5]. This virtual machine provides secure shared objects and shared name spaces. Whereas a conventional operating system provides an abstraction of a single computer, Legion aggregates a large number of diverse computers running different operating systems into a single abstraction. As part of this abstraction, Legion provides

mechanisms to couple diverse applications and diverse resources, thus simplifying the task of writing applications in heterogeneous distributed systems. This abstraction supports the performance demands of scientific applications, such as CHARMM. CHARMM runs as a legacy MPI application on queuing systems accessed by Legion. In the following subsections, we discuss some features of queuing systems and Legion’s support for legacy applications.

The grid chosen for running CHARMM was *npacinet*, a nation-wide grid consisting of heterogeneous resources present at multiple sites and administered by different organisations. The majority of the organisations contributing resources to *npacinet* are part of NSF’s National Partnership for Advanced Computing Infrastructure (NPACI) thrust. Legion has been managing this grid continuously for several months during which we have demonstrated Legion features numerous times, conducted tutorials on multiple occasions and supported various academic users running a variety of applications.

### A. Queuing Systems

Traditionally, queue systems have been used to schedule jobs on a cluster of nodes [2] [7] [11] [12]. When a user submits a job, the queue provides her with a ticket or job ID or token, which can be used to monitor the job at any later time. The ticket becomes invalid shortly after the job completes. Most queuing systems comply with a POSIX interface that requires three standard tools for running jobs: a submit tool (`qsub` in PBS, `bsub` in LSF, `llsubmit` in LoadLeveler), a status tool (`qstat` in PBS, `bjobs` in LSF, `llstatus` in LoadLeveler) and a cancel tool (`qdel` in PBS, `bkill` in LSF, `llcancel` in LoadLeveler). In addition, some queues provide other tools to check on the aggregate status of the queuing system, e.g., `bqueues` in LSF and `llq` in LoadLeveler. A queuing system’s status tool may report that a particular job is queued, running or terminated. If the execution of a particular job is deemed undesirable, the cancel tool can be used to terminate the job. Most queuing systems do not provide tools to access intermediate files or supply additional inputs. A user desiring such functionality must employ shared file systems or other file transfer tools. Queuing systems do not provide any support for checking aggregate progress of large sets of jobs. Users must check on the progress of each job individually or construct interfaces to monitor the progress of the entire set of jobs.

Part of the abstraction Legion provides is to hide the differences among

Fig. 1. Simple Application

queuing systems as well as between queuing and non-queuing systems. A user running over Legion does not have to know the particulars of every system on which a job could run. To appreciate why this abstraction is important, consider running a simple application, such as “Hello, world”, on different systems. We could run on a Unix or Windows system by writing a shell script or batch file such as the one in Fig. 1. However, if we wanted to run on a cluster of nodes controlled by Portable Batch System (PBS), we would have to modify the

application to construct a submission script as in Fig. 2. If we

```
#!/bin/ksh
#PBS -A anand
#PBS -c n
#PBS -m n
#PBS -N LegionObject
#PBS -r n
#PBS -l nodes=1:ppn=1:walltime=00:10:00
#PBS -p 1
#PBS -o test.o
#PBS -e test.e

echo 'Hello, world'
```

Fig. 2. Simple Application modified for PBS

decided to run on nodes controlled by Maui/LoadLeveler, we would have to construct a submission script as in Fig. 3. Not

```
#!/bin/ksh
# @ environment = COPY_ALL;MP_EUILIB=us
# @ account_no = met200
# @ class = express
# @ node = 1,1
# @ tasks_per_node = 1
# @ wall_clock_limit = 00:10:00
# @ input = /dev/null
# @ output = test.o
# @ error = test.e
# @ initialdir = /home/uxlegion

echo 'Hello, world'
```

Fig. 3. Simple Application modified for Maui

only are different queuing systems dissimilar, but the same queuing system installed at different sites may be dissimilar in terms of configuration parameters. Moreover, the tools for running special applications, for example MPI programs, may be different (`mpirun` versus `pam` versus `poe`). The different submission scripts required to run on different systems restrict a user in two significant ways:

1. The user is forced to learn the particulars of each queuing system, thus increasing his cognitive burden and increasing the time before he can start becoming productive on these systems.
2. When running large numbers of jobs, the user must construct submission scripts for running on each queuing system. The very act of creating a submission script *a priori* forces the user to construct a static schedule for running his jobs. Consequently, he cannot take advantage of dynamic load changes on resources to schedule jobs.

In addition to hiding the differences among queuing systems regarding submission scripts, Legion also hides differences among their submit, status and cancel tools. Also, Legion hides differences regarding the manner in which MPI jobs are run. Moreover, Legion provides tools and mechanisms for accessing intermediate files and viewing the aggregate status of large numbers of jobs. Finally, Legion does not require the user to log on to the various queuing systems to initiate jobs. In terms of convenience, single sign-on is one of the most useful features a grid operating system can provide.

## B. Legacy Applications

Legion supports running legacy applications on a grid. Legacy applications are those that have not been targetted specifically to a grid or Legion. Legion supports such applications “as is”, i.e., the user neither has to change a single line of code nor re-link the object code to run such an application. All Legion requires are the executables for the application for various architectures. A user who chooses this form of support understands the trade-offs for the convenience of not changing the application at all. One trade-off is that Legion can control very few aspects of the execution of the job after it is initiated. For example, Legion cannot provide restart support for a legacy application if the application itself does not write checkpointing data. However, Legion can and does provide support for starting the job, checking its status as reported by the underlying operating system or queuing system and terminating the job if necessary. In addition, Legion provides the ability to send in or get out intermediate files while the job is running.

## IV. EXECUTING CHARMM ON NPACI RESOURCES

The steps the user had to undertake to run CHARMM over Legion are illustrated in Fig. 4. All of these steps were performed after the user logged on (in the Unix sense) to a machine on which Legion had been installed. The shaded boxes represent the steps the user performed without Legion’s help. Of these, two, “Creating Jobs” and “Analysing Results”, are specific to the application. The third, “Creating Executables”, could have been performed with Legion’s help. The unshaded boxes represent steps performed with Legion’s help. The user had to learn one new command for each of the unshaded boxes. Learning four commands is a small price to pay for the ability to run multiple parallel jobs on distributed heterogeneous resources in a secure and fault-tolerant manner.

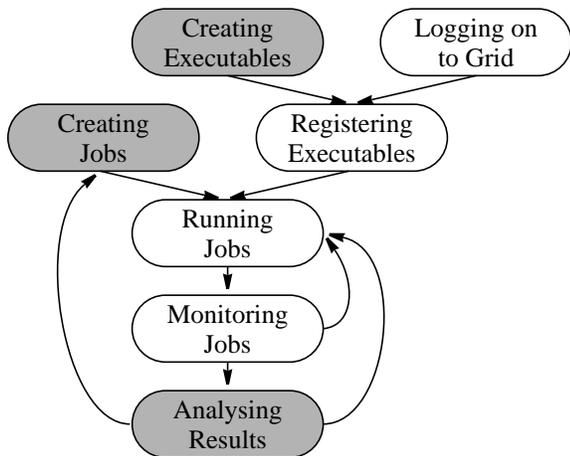


Fig. 4. Steps for running CHARMM over Legion

### A. Creating Executables for CHARMM

This step involved creating the executables for CHARMM. Recall that the user chose Legion’s legacy MPI support for CHARMM. If the user desired, he could have used

`legion_make`, a tool to compile the source code on machines or architectures of his choosing (in which case, he would have done so after “Logging on to Grid”). The resulting executables would still be “legacy code” because Legion would not require changing the source code or linking the object code against Legion libraries. Currently, `legion_make` works for applications that have relatively simple and standard make rules, i.e., it works for applications that use standard compilers and have straightforward local dependencies. Since CHARMM is not such an application, the user decided to compile CHARMM for different architectures without Legion’s help.

### B. Creating Jobs

This step involved creating a set of input files for each job. Clearly, this step is application-specific and requires no help from Legion.

### C. Logging on to Grid

In order to log on to the npacinet grid, the user ran the command `legion_login`, which required him to enter his Legion ID and password. Once the user logged in, Legion did not require him to log on to any other machine.

### D. Registering Executables

Registering executables is the process by which Legion can run a Unix or Windows executable. After an executable is registered with `legion_register_program`, Legion has the information necessary for selecting the appropriate executable to run on any particular machine. Multiple executables of different architectures may be registered with the same Legion object. The benefit is that a user can request Legion to “run” the object without having to manage which executable should be copied and run on which machine. For example, Legion will ensure that only a Solaris executable is copied and run on a Solaris machine.

### E. Running Jobs

After registering the executables for every architecture of interest, the user requested Legion to “run” the object with the command `legion_run`. This command has a number of parameters and options (details are in the Legion man pages accompanying the standard distribution [1]). Parameters for this command include the name of the object and parameters for the job, the names of input and output files for the job, and options such as number of nodes desired, tasks per node desired, duration, etc. Reasonable defaults are chosen for unspecified options. The user may specify a particular machine on which to run or let Legion choose the machine. Likewise, the user may choose to run on any machine of a particular architecture or let Legion make that decision.

The CHARMM user specified the input and output files for each job and the machines on which he desired to run. In addition, he specified the name of a “probe file” for monitoring the job. The user ran the `legion_run` command as many times as he wanted to initiate jobs. Although he chose different machines on which to run different jobs (effectively self-

scheduling his application dynamically), at no point did he have to write a single submit script, log on to any other machine<sup>1</sup>, copy executables and input/output files, or learn a new command for running jobs. The user could have initiated as many jobs as he desired concurrently; in practice, he initiated a few tens of jobs concurrently because i) the nature of the jobs imposed sequential dependencies, and ii) initiating multiple jobs is pointless when the next job is certain to be queued behind previous ones.

#### F. Monitoring Jobs

The user monitored each job in two ways. First, he started a console object for the Unix shell from which he initiated his jobs with one command, `legion_tty`. After the console object was started, output and error messages printed by the user's jobs or the queuing systems on remote machines became visible on the user's shell. Second, the user requested Legion to save a probe for every job. Using the probe and a tool called `legion_probe_run`, the user determined the status of each and every job as well as sent in and got out intermediate files at his leisure. If at any time the user determined that a job was not progressing satisfactorily, he terminated that job, corrected any problems and restarted that job.

#### G. Analysing Results

The final step involved analysing the results from each job. A basic analysis step involved determining whether each job actually ran to completion. The user made this determination by checking whether a certain output file contained specific lines in it. A large part of the subsequent analysis involved retrieving archived files and processing them by running CHARMM again. The subsequent steps were specific to the application and are outside the scope of this discussion.

## V. RESULTS

The experiment was conducted successfully over a period of two days. The user logged in to one machine at the University of Virginia on which Legion was installed<sup>2</sup>. From a single shell on that machine, he initiated as many jobs as he could, subject to the limitations discussed earlier. Some of the jobs failed, but a large number ran to completion successfully. Consequently, although the user did not manage to complete all of the jobs he desired initially, a significant fraction of the jobs were completed. The experiment showed the viability of running large, high-performance applications on a computational grid. In the following sections, we discuss how well Legion met the goals mentioned in Section I.

<sup>1</sup> In fact, in the current configuration of Legion on the NPACI machines, the user was not even required to own accounts on the machines. Legion ran his jobs as a generic user on those machines. In the future, the NPACI resources may insist that the user can run on their machines only if he has an account on them as well. Since respecting site autonomy is a critical part of the Legion philosophy, support the latter mode of operation is under progress.

<sup>2</sup> Legion was not installed on the user's machines at The Scripps Research Institute (TSRI) because of site-specific firewall restrictions.

#### A. Increasing User Productivity

A success of this experiment was that the grid was used to generate results for an actual scientific study. At the time of writing, around 80 of the desired 400 jobs had been completed. We demonstrated that Legion can be used to harness a vast amount of processing power harnessed for scientific users. In the final tally, 728 processors of different architectures and speeds were utilised for this experiment. The breakdown of these processors is:

- 512 IBM Blue Horizon Power3s running at 375MHz at San Diego Supercomputing Center (SDSC)
- 128 HP PA-8700 running at 440MHz at California Institute of Technology (CalTech)
- 24 IBM SP3 Power3s running at 375MHz at University of Michigan (UMich)
- 32 IBM Azure Power2s running at 160MHz at University of Texas (UTexas)
- 32 DEC Alpha EV56s running at 533MHz at University of Virginia (UVa)

In the future, we intend adding the following resources:

- 128 T3Es running at 300MHz at SDSC
- 88 T3Es running at 300MHz at UTexas
- 32 Sun HPC 10000s running at 400MHz at SDSC
- 32 dual-CPU Intel Pentium IIs running at 400MHz at UVA

Fig. 5. Breakup of CHARMM jobs completed

Rough estimates show that if the user had used the resources available at his organisation alone (128 SGI Origins), it would have taken one month to complete what was complete in less than two days on the grid. The number of jobs run on each resource is shown in Fig. 5. The vast majority of the jobs ran on the Blue Horizon at SDSC because that machine was by far the most powerful machine in the mix of available machines. Some of the machines did not contribute significantly to the results because of run-time problems (see below).

#### B. Simplifying Grid Access

Legion's ease of use could be measured in what the user had to do as well as what the user did *not* have to do to run his jobs. As described in Section IV, the user had to learn a mere four or five commands to run on the grid. The small number of commands is comparable to the number the user would have to learn for *each* queuing system had he chosen to run without Legion. During the experiment, the user did not have to log on to any of the queuing systems. He logged on to one machine at UVa on which Legion was installed. From a single shell on that machine, he initiated multiple jobs. Legion made the heterogeneous NPACI resources available to the user without

his having to know the details of how to run on each resource. The heterogeneity of the resources extended in a number of dimensions:

- 6 different organisations (UVa, TSRI, SDSC, UTexas, UMich, CalTech)
- 5 (6) different queue types (Maui, LoadLeveler, LSF, PBS, NQS)
- Up to 9 (11) different queuing systems
- Up to 5 (6) different architectures (IBM AIX, HP HPUX, Sun Solaris, DEC Linux, Intel Linux, **Cray Unicos**)

### C. Identifying and Eliminating Problems

A number of run-time problems caused fewer total jobs to complete. Minor organisational problems aside, the problems we encountered fell into two categories: network slowdowns and site failures. The Legion run-time system suffered no problems during the experiment, although a number of potential extensions were identified. Also, although the CHARMM user used the grid heavily, the remaining users on the same grid were unaware of the experiment. While the experiment progressed, other Legion users continued to run their usual jobs on the grid.

*Network Slowdown.* During the experiment, we experienced slowdowns in the network connections between UVa and SDSC. From around noon through about 3PM Eastern Time, medium-sized to large packets were transmitted from one site to the other with great difficulty. Preliminary investigation showed that packets of size equal to or greater than 8800 bytes were lost entirely. Packets in the size range 8000-8800 bytes suffered over 90% loss rates. The loss rates for packets of size less than 8000 bytes were lower but still significant. The implication for Legion was that some messages between objects had to be retransmitted a number of times to ensure that they were received correctly. Consequently, for the CHARMM user, monitoring jobs became a slow process. At one point, inquiring about the status of a job took nearly a minute to complete. Ordinarily, this process is almost instantaneous. Since the user could not monitor jobs quickly enough to start new ones, throughput was reduced.

*Site Failures.* Some of the NPACI sites experienced unforeseen failures. For example, at UMich, Legion encountered a number of errors that pointed to NFS failures. Since the ability to access permanent storage is important to Legion as well as CHARMM, the NFS failures reduced the throughput of CHARMM jobs. On the Blue Horizon machine at SDSC, the queuing system, Maui/LoadLeveler, had to be restarted a number of times because it became overloaded. During the time the queuing system was down, currently-running jobs continued to run. However, the queuing system could not inform anyone about the status about those jobs. Since “no information” is similar to what the queuing system reports when a job has been complete for a while, Legion assumed the jobs were complete and informed the CHARMM user accordingly. This erroneous reporting led the user to

believe that it was safe to access the output files from the job. However, on analysis of these jobs, the user discovered that the output files were only partially complete. At UMich, the purge policy in place removed CHARMM files as well as persistent state required by Legion objects. Without their persistent state, Legion objects can behave erroneously. Likewise, without the appropriate input files CHARMM cannot run as intended.

### D. Extending Legion

The experiences of the CHARMM user enabled us to identify potential extensions to Legion. These extensions would enhance Legion’s usability by building on low-level functionality already present.

*Graceful Error Handling.* Legion has been designed to mask many kinds of failures from end-users. While this strategy usually benefits the user, sometimes it is important for the grid infrastructure *not* to mask failures from the user. For example, the network failures discussed earlier were masked from the user who saw only gracefully-degraded performance. However, Legion also masked most site failures from the user, which often conveyed the mistaken impression that Legion itself had failed. Consequently, we are reviewing all aspects of error handling and propagation in Legion.

*Support for Archiving.* Although Legion permits users to specify input and output files at any time during the execution of a job, archival support is almost non-existent. In particular, there is no way for a user to specify that some files are meant to be stored on some kind of long-term storage after the job is complete. Instead, the Legion file solutions are that after a job is complete, the files are either copied out to the user’s local directories, or to Legion’s own distributed shared file system, or deleted. None of these solutions is satisfactory for jobs that generate large amounts of data. The user’s local directories or the individual components of the distributed file system may not have space to store large amounts of data. Moreover, the user may not want to copy files out the moment the job is done. Instead, scientific users generating large amounts of data, such as the CHARMM user, are likely to want to archive the data generated by their jobs on some long-term storage and access the data at their leisure. Since Legion developers did not anticipate such a need, currently, archiving has to be done by users themselves as part of their jobs.

*Support for Parameter-Space Studies.* We are making it possible for users to run parameter-space studies of parallel codes with a single command. The CHARMM user had to issue a fresh `legion_run` command for every job. Legion provides another tool, `legion_run_multi`, which enables multiple jobs to be started with one command. This tool works well if every job is a sequential program; we are looking to extend it to parallel programs.

*Web Interfaces.* We are developing a web portal that scientists can use to run CHARMM jobs. Currently, we have a Legion web interface for the entire grid. We intend adding a CHARMM-specific component to this interface.

### E. Observations

1. *High-Level Services:* The CHARMM experiment reassured us that a grid infrastructure must provide low-level functionality *and* high-level services. We consider it a significant advantage that using Legion, the user accessed heterogeneous resources controlled by multiple organisations with four or five commands and achieved order-of-magnitude speedup as compared to running at just one site.

2. *Human Factors:* The humans involved in the experiment were critical to the experiment's success. Three people were involved intimately with the continuous progress of the runs: the user, a Legion liaison and an NPACI liaison. The Legion liaison was present in case problems arose with Legion itself during the execution. Since Legion itself suffered no run-time problems, this person used Legion tools to identify site-specific problems as they arose. The NPACI liaison coordinated on-site efforts to keep the experiment running. Finally, administrators at individual sites ensured that problems were resolved as soon as possible by correcting misconfigurations, restarting services, increasing quotas, etc. Although this collaboration was rewarding, in the future the involvement of all parties except the user must be eliminated.

3. *Site Services:* The number of site failures that were identified was astonishingly high. Normally, users never expect services such as queues and operating systems to fail. Likewise, users rarely consider network failures when running their applications. However, running large numbers of high-performance jobs can stress-test every component of a grid. We discovered previously-ignored limits on the number of jobs queues can manage, credential expirations with file systems, queue-imposed limits on durations of jobs, purge policies, process table limits, quota exhaustions and numerous other problems, each of which could make a site unusable for continued running.

## VI. CONCLUSION

We demonstrated that Legion is a suitable environment for running large, high-performance jobs, such as CHARMM, on a grid. Legion provides a suite of tools for a grid that are similar to what traditional operating systems provide for a single system. Using these tools, users can start, monitor and terminate jobs on remote machines in a straightforward manner. The CHARMM runs on heterogeneous machines controlled by different organisations showed that Legion is

able to mask unwanted detail from the end-user, thus permitting him to focus on completing his work. Although we encountered a number of problems during the run, it is encouraging to note that none of the problems are unsolvable; solutions for each of them are forthcoming from the lessons we have learnt. These solutions can only improve and simplify a user's access to a computational grid.

### ACKNOWLEDGEMENT

We thank Dave Carver at UTexas; Sharon Brunett at CalTech; Victor Hazlewood, Larry Diegel and Kenneth Yoshimoto at SDSC; Tom Hacker and Rod Mach at UMich; and Katherine Holcomb and Norm Beekwilder at UVa for providing the CHARMM user with accounts, providing Legion accounts and allocations, and resolving problems if and when they arose at their respective sites.

### REFERENCES

- [1] —, "The Legion Manuals (v1.7)", University of Virginia, October 2000.
- [2] Bayucan, A., Henderson, R. L., Lesiak, C., Mann, N., Proett, T., Tweten, D., "Portable Batch System: External Reference Specification", Tech. Rep., MRJ Technology Solutions, November 1999.
- [3] Brooks, B. R., Bruccoleri, R. E., Olafson, B. D., States, D. J., Swaminathan, S., Karplus, M., "CHARMM: A Program for Macromolecular Energy, Minimization, and Dynamics Calculations", *J. Comp. Chem.*, vol. 4, 1983.
- [4] Grimshaw, A. S., Wulf, W. A., "The Legion Vision of a Worldwide Virtual Computer", *Comm. of the ACM*, vol. 40, no. 1, January 1997.
- [5] Grimshaw, A. S., Ferrari, A. J., Lindahl, G., Holcomb, K., "Metasystems", *Comm. of the ACM*, vol. 41, no. 11, November 1998.
- [6] Hempel, R., Walker, D. W., "The Emergence of the MPI Message Passing Standard for Parallel Computing", *Comp. Stds. and Interfaces*, vol. 7, 1999.
- [7] International Business Machines Corporation, "IBM LoadLeveler: User's Guide", September 1993.
- [8] Kingsbury, B. A., "The Network Queuing System (NQS)", Tech. Rep., Sterling Software, 1992.
- [9] MacKerell, A. D. Jr., Brooks, B. R., Brooks, C. L. III, Nilsson, L., Roux, B., Won, Y., Karplus, M., "CHARMM: The Energy Function and Its Parameterization with an Overview of the Program", *The Encycl. of Comp.Chem.*, vol. 1, 1998.
- [10] Snir, M., Otto, S., Huss-Lederman, S., Walker, D. W., Dongarra, J., *MPI: The Complete Reference*, MIT Press, 1998.
- [11] Zhou, S., "LSF: Load Sharing in Large-scale Heterogeneous Distributed Systems", *Proc. of Workshop on Cluster Computing*, December 1992.
- [12] Zhou, S., Wang, J., Zheng, X., Delisle, P., "Utopia: A Load Sharing Facility for Large, Heterogeneous Distributed Computer Systems", *Software Practice and Experience*, Vol. 23, No. 2, 1993.