

# Grid Computing and Biomolecular Simulation

Christopher J. Woods<sup>1</sup>, Muan Hong Ng<sup>2</sup>, Steven Johnston<sup>2</sup>, Stuart E. Murdock<sup>1,2</sup>, Bing Wu<sup>3,4</sup>, Kaihsu Tai<sup>4</sup>, Hans Fangohr<sup>2</sup>, Paul Jeffreys<sup>3</sup>, Simon Cox<sup>2</sup>, Jeremy G. Frey<sup>1</sup>, Mark S.P. Sansom<sup>4</sup>, Jonathan W. Essex<sup>\*1</sup>

<sup>1</sup>School of Chemistry and <sup>2</sup>Southampton E-Science Centre, University of Southampton, UK

<sup>3</sup>Oxford e-Science Centre and <sup>4</sup>Department of Biochemistry, University of Oxford, UK

## Abstract

Biomolecular computer simulations are now widely used not only in an academic setting to understand the fundamental role of molecular dynamics on biological function, but also in the industrial context to assist in drug design. In this paper, two applications of Grid computing to this area will be outlined. The first, involving the coupling of distributed computing resources to dedicated Beowulf clusters, is targeted at simulating protein conformational change using the Replica Exchange methodology. In the second, the rationale and design of a database of biomolecular simulation trajectories is described. Both applications illustrate the increasingly important role modern computational methods are playing in the life sciences.

## 1. Background

Grid computing is becoming increasingly important in the area of the life sciences. Two particular aspects dominate. First, distributed computing is a potentially powerful approach for accessing large amounts of computational power. Cycle stealers, which allow a PC user to donate the spare power of their computer, are now used in a wide range of scientific projects, e.g. the SETI@home study,<sup>1</sup> the CAN-DDO cancer screening project<sup>2</sup> and folding@home.<sup>3</sup> Cycle stealers are also becoming more widely used in the pharmaceutical industry, particularly for virtual screening projects. Second, large databases are used to hold the substantial amount of data now involved in the study of biological systems, and have found particular prominence in the field of bioinformatics. In this paper, two recent developments in each of these areas, as applied to biologically relevant problems, will be described.

## 2. Distributed Computing

While the use of cycle stealers can provide supercomputer-like resources, their use is limited to calculations that may be split into many independently parallel parts (i.e. coarsely parallel simulations). The distributed and unreliable nature of this resource makes it unsuitable for closely coupled parallel calculations. For these calculations, the speed

and latency of inter-processor communication are a bottleneck that cannot be overcome simply through the addition of more nodes. Unfortunately, a large number of chemical simulations require closely coupled parallel




Figure 1. The native (blue) and phosphorylated (red) conformations of NTRC. The site of phosphorylation (Asp54) is shown as spheres.

calculations, and are thus not suitable for deployment over a distributed computing cluster. An example of such a simulation is the investigation of protein conformational change. These simulations are typically performed using molecular dynamics (MD),<sup>4</sup> where the motions of the atoms are integrated over time using Newton's laws. These simulations cannot be

\* Author to whom correspondence should be addressed. Email: jwe1@soton.ac.uk, phone: +44 (0)23 8059 2794, fax: +44 (0)23 8059 3781




Figure 2. Progress of the simulation at each temperature as a function of simulation time. Iterations are run at each temperature; odd iterations are shown in blue and even iterations are shown in green. Our scheduler has to cope with extreme events, e.g. the complete failure of the distributed cluster after about 200 and 270 hours of simulation. The distributed cluster contains both fast and slow nodes. Some iterations can thus take a lot longer than others (visible here as longer bars). In addition, the owners of the PCs will also wish to use them (shown here as red dots). This will interrupt the calculation that temperature, again slowing it down relative to the other temperatures. Because of this, neighbouring temperatures will be ready to test at different times. This can lead to a loss of efficiency as completed temperatures wait for their neighbours. In the worst case this waiting can propagate, as occurs for temperatures around 310 K after 360 hours of simulation. To help prevent this, a catch-up cluster is used that identifies and reschedules slow temperatures (use of the catch-up cluster is shown in yellow).

broken up into multiple independent parts, as each nanosecond of MD must be run in series and in sequence.

The investigation of protein conformational change is important as it lies at the heart of many biological processes, e.g. cell signaling. Some bacteria regulate nitrogen metabolism using one such signalling pathway. Nitrogen Regulatory Protein C (NTRC)<sup>5</sup> plays a key role in this pathway. Changes in nitrogen concentration activate the kinase NTRB. This phosphorylates an aspartate residue in NTRC, causing it to change conformation (figure 1).<sup>5</sup> This change in conformation allows the NTRC to join together to form oligomers, which then help promote the transcription of genes. These genes are used to produce proteins that are used in nitrogen metabolism.<sup>5</sup> A key stage of this pathway is the change in conformation that occurs in NTRC when it is phosphorylated. It is difficult to study this conformational change experimentally as the phosphorylated form of NTRC has a very short lifetime.<sup>6</sup> It is thus desirable to simulate the NTRC protein and

encourage the conformational change on computer.

## 2.1 The Replica Exchange Method

We can use a distributed computing cluster to investigate protein conformational change via Replica Exchange simulations.<sup>7,8</sup> Multiple simulations of the protein are run in parallel, each running under a different condition, e.g. temperature. Periodically the simulations running at neighbouring temperatures are tested and swapped. This enables simulations at high temperatures, where there is rapid conformational change, to rain down to biologically relevant temperatures where conformational change occurs more slowly. The testing of neighbouring temperatures introduces a light coupling to the simulation, meaning that it no longer fits the archetypal coarsely parallel distributed computing model. This light coupling introduces inefficiencies to the scheduling of the simulation, as any delay in the calculation of one temperature can propagate

out to delay the calculation of all temperatures. To help overcome this, a catch-up cluster has been developed that monitors the simulation for temperatures that are taking too long to complete, and that are likely to negatively impact the overall efficiency of the simulation. Once identified, the calculation of these temperatures is rescheduled onto a small, yet fast and dedicated, computational resource so that they can ‘catch-up’ with the other temperatures (figure 2).

### 3. Experimental Details

Nuclear Magnetic Resonance (NMR) structures of the phosphorylated (1DC8) and unphosphorylated (1DC7) conformations of the NTRC protein were obtained from the protein databank.<sup>9</sup> Polar hydrogen atoms were added via WhatIf<sup>10</sup>. The proteins were solvated in 60<sup>3</sup> Å<sup>3</sup> boxes of TIP3P water and sodium ions were added via the XLEAP module of AMBER 7.0<sup>11</sup> to neutralise the system. The CHARMM27 forcefield<sup>12</sup> was used, and the systems

attempted between neighbouring temperatures every 2 ps, after the initial 20 ps of sampling that was used to equilibrate each replica to its initial temperature. A Langevin thermostat<sup>14</sup> and a Nose-Hoover Langevin piston barostat<sup>15</sup> were used to sample at constant temperature and pressure, while SHAKE<sup>16</sup> was used to constrain hydrogen bond lengths to equilibrium values. A 1 fs time step was used for the molecular dynamics integrator, and the non-bonded interactions were evaluated using a 12 Å cutoff and the particle mesh ewald sum.<sup>17</sup>

#### 3.1 Details of the Distributed Cluster

The simulations were run over the Condor<sup>18</sup> cluster provided by the University of Southampton. This cluster uses Condor<sup>18</sup> to make available the spare cycles of approximately 450 desktop computers running Microsoft Windows NT 5.1. The two replica exchange simulations were run simultaneously on this cluster. There was little competition between the two simulations for nodes as they




Figure 3. The number of nodes in use for the replica exchange simulation over just the distributed cluster (left), and over the distributed and catchup clusters (right).

minimised, then annealed from 100 K to 300 K. The systems were finally equilibrated for 100 ps at constant temperature (300 K) and pressure (1 atm). The final structures from equilibration were used as the starting structures for all of the replicas.

The temperatures for each replica were chosen using a custom program that optimised the temperature distribution such that a replica exchange move was accepted with a probability of 20%. The temperatures chosen ranged from several replicas below the target temperature of 300 K (290.1 K) to a maximum of 400 K. In total 64 replicas were used for each of the two proteins.

The simulations were conducted using NAMD 2.5.<sup>13</sup> A replica exchange move was

each required a maximum of 64 nodes out of the available 450.

#### 3.2 Implementation of the Catchup Cluster

The catchup cluster was implemented via dedicated dual Xeon 2.8 GHz nodes running Linux. Each Xeon processor was able to provide 2 virtual processors, allowing NAMD to run in parallel over 4 virtual processors per node. The fast catchup cluster consisted of two dual Xeon nodes, thus allowing it to catch up two replicas simultaneously.

To test the utility of the catchup cluster, it was only made available to the replica exchange simulation on the phosphorylated conformation of the protein (1DC8). As both replica exchange simulations were running simultaneously, any

differences in efficiency should thus be wholly attributable to use of the catchup cluster.

## 4. Results

Figure 3 shows the number of nodes in use during the replica exchange simulations on the phosphorylated and unphosphorylated conformations of NTRC. The initial phase of the simulation involved the 20 ps of equilibration of each replica to its initial temperature. This was broken down into 10 iterations of 2 ps. As there were no replica exchange moves during these first 10 iterations the replicas were all independent and thus the maximum number of 64 nodes were in use. However, there were efficiency problems during this phase of the simulation, as the unreliable nature of the distributed cluster caused several short periods of downtime that stopped both simulations. Frequent periods of downtime were common throughout the rest of replica exchange simulations.

The second stage of the simulations occurred when the replicas began to complete their 10<sup>th</sup> iteration. At this point each replica had to wait for its partner to complete 10 iterations such that the pair of replicas could be tested and potentially swapped. Owing to the range of processors available in the distributed cluster and the different impacts of downtime on each of the replicas, there was a large spread of times over which each replica completed 10 iterations. This meant that a large number of replicas were left waiting for a significant time for their partner to complete, and thus the number of nodes in use for each simulation dropped from the maximum of 64 down to approximately 30. If the efficiency is defined as the ratio of the number of nodes in use compared to the theoretical maximum, then the efficiency dropped from 100% down to about 47%.

After this dip in efficiency, the simulations then moved towards the final stage, which was a steady state, where the number of replicas running and the number of replicas waiting for their partner to complete reached a consistent range of values. This steady state was periodically disrupted by failure of the condor cluster, but was always quickly recovered once the disruption was over. The steady state for the simulation that used the catchup cluster had significantly more replicas running, and significantly fewer replicas waiting compared to the simulation that did not use the catchup cluster. The catchup cluster clearly improved the steady state number of nodes in use to approximately 50, compared to approximately




Figure 4. Histogram of the times to complete each iteration of the replica exchange simulation on the phosphorylated form of NTRC.

40 for the simulation that did not use the catchup cluster. This is an improvement in efficiency from 63% to 78%.

### 4.1 The Heterogenous Distributed Cluster

The distributed condor cluster consisted of a range of desktop computers with varying processor speeds. To investigate the effect of running the simulation on this heterogeneous cluster, the total simulation time for each iteration was histogrammed. The histogram of replica completion times for the phosphorylated form of the protein is shown in figure 4. This figure shows that while the majority of iterations completed in under ten thousand seconds (2.8 hours), there was a significant spread of replica completion times up to twenty thousand seconds (5.6 hours). This spread of completion times means that fewer iterations have been completed than would have been expected based on the speed of the processors used.

### 4.2 Comparison to Normal MD

The use of dual Xeons in the catchup cluster aiding the efficiency of a replica exchange simulation was compared to their use as dedicated nodes running a normal Molecular Dynamics simulation. The phosphorylated conformation of NTRC was simulated at 300 K using a standard MD simulation with an identical starting structure as the replica exchange simulation and identical simulation conditions. Over the same period of time as the replica exchange simulations were running, this MD simulation on a single dual Xeon node completed 1.9 ns of dynamics. This compares to a total of 10.5 ns of dynamics generated by the corresponding replica exchange simulation. However, the 1.9 ns of dynamics generated via

the MD simulation forms a single, self-consistent trajectory. In comparison, the 10.5 ns of sampling from the replica exchange simulations was formed over 64 individual trajectories of only 0.16 ns in length. The dedicated node has produced a single trajectory over ten times the length of those produced via the distributed condor cluster with catchup cluster. This is despite the dedicated node only running the MD approximately 3.5 times the speed of a typical node in the distributed cluster. There are two reasons for this discrepancy; first, as demonstrated in figure 4, the heterogeneous nature of the distributed cluster meant that there was a large spread in the amount of time needed to complete each iteration of the replica exchange simulation. This could be mitigated against by running each iteration twice at the same time on the distributed cluster and using the results from the first node that completed the calculation. The second reason for the discrepancy is that the distributed cluster was very unreliable, leading to large periods of time when the simulation was not running. The replica exchange simulations presented here were run at a time when the condor cluster was experiencing a higher than normal amount of downtime. It is anticipated that during normal operation the condor cluster would be more reliable, and that the steady state efficiency of the replica exchange simulations would be maintained throughout the majority of the simulation. However, the experience of running these simulations demonstrates that applications that use distributed clusters need to include estimates of downtime and the range of available resources when predicting how long a particular simulation will take to run. These results also demonstrate that a distributed computing resource is, unsurprisingly, not efficient compared to a dedicated computing resource. However, distributed computing typically provides resources that would otherwise not be available.

#### 4.3 Effectiveness of Replica Exchange

The primary aim of running these simulations over the distributed computing resource was to sample the conformational change induced by phosphorylating NTRC. The aim was to use replica exchange to swap simulations running at high temperature, where the conformational change occurs more rapidly, down to room temperature, where the simulation statistics are collected. Figure 5 shows the temperature for each replica of the unphosphorylated simulation as a function of iteration. Four of the replicas are highlighted. This figure shows that while the




Figure 5. Temperature of each replica as a function of iteration from the replica exchange simulation on the unphosphorylated conformation of NTRC. Four randomly chosen replicas are highlighted.

replica exchange moves were accepted with the desired frequency, the replicas themselves did not travel far in temperature. Instead each replica tended to oscillate around its initial temperature. No replicas from high temperature swapped down to room temperature. This shows that the replica exchange simulations need to be continued for many more iterations before the improved sampling of high temperature is able to be of use in enhancing the rate of sampling at room temperature.

### 5. Summary – Distributed Computing

Distributed computing provides a resource that is not ideally suited to a wide range of chemistry problems. The investigation of protein conformational change is one such problem. The replica exchange algorithm was used in an attempt to fit this chemistry problem to the distributed computing resource. The coupled nature of replica exchange simulations caused problems for the scheduling of the computation that were partially solved through the development of a dedicated catchup cluster. This cluster improved the efficiency of the replica exchange simulation from 63% to 78%.

### 6. BioSimGrid: A database for biomolecular simulation

As evidenced by the preceding example, computer simulations play a vital role in biochemical research. These simulations are computationally demanding and they produce huge amounts of data (up to ~10 GB each) that is analysed by a variety of methods in order to obtain biochemical properties. Generally, these data are stored at the laboratory where they have been computed in a proprietary format that is

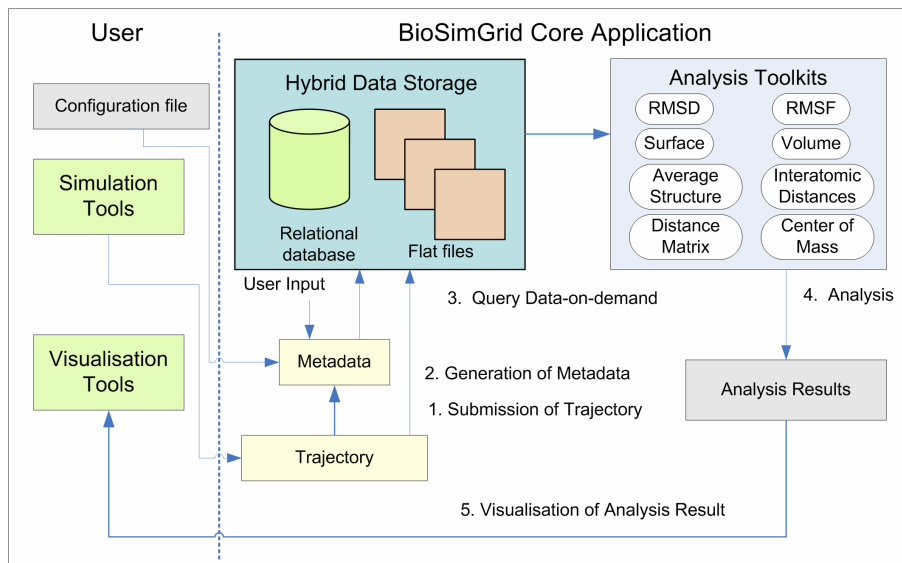


Figure 6. A schematic sketch of the work-flow in the BioSimGrid project.

unique to the simulation code that has been used. This constrains the sharing of data and results within the biochemistry community: (i) the different simulation results are usually not available to other groups and (ii) even if they are exchanged, for example via FTP, then the data can generally not be compared easily with post-processing tools due to the varying data formats. BioSimGrid will facilitate the comparative analysis of these simulations, allowing more general structure/dynamics/function relationships to be discovered.

BioSimGrid<sup>19,20,21</sup> seeks to tackle this problem by enabling biochemists to deposit their simulation data of varying formats to a shared repository. This will allow biochemists to retrieve a slice or part of a protein in a uniform way for post-simulation analysis. BioSimGrid also provides an integrated analysis environment. By providing a uniform data storage and data retrieval mechanism, different proteins can be compared easily.

Figure 6 demonstrates typical scenarios of using the BioSimGrid project. The completion of a biomolecular simulation delivers simulation data, which is called a "trajectory". A trajectory consists of many frames (corresponding to time steps in the simulation process) of simulation data recording the positions and velocities of all atoms. The first step is to submit the new trajectory and all the relevant meta-data (which describes the simulation and will allow sophisticated querying of all submitted trajectories) to the database. The extraction of the meta-data (from the simulation configuration files) and the trajectories (from the simulation data files) is fully automated, but the user has the option to provide additional information, such as publication references that cannot be extracted from the simulation configuration and data files.

Once the data is stored in the database, users can query different slices of one or more

trajectories and perform a number of standard analysis computations (a selection is shown in the figure) on these data. The work flow is then completed by the graphical display of the analysis results (either as vector graphics, bitmaps, movies or using interactive 3-dimensional environments such as Visual Molecular Dynamics<sup>22</sup> and PyMol<sup>23</sup>). The results are, of course, also available in text or data files.

The following section of the paper discusses two related projects on grid-enabled data storage. Section 8 describes the architecture of BioSimGrid where the data storage layer, the middleware layer and the user interface layer are discussed in detail. Section 9 gives a brief outline on the current issues and the future work on the next prototype, and we finally conclude in section 10.

## 7. Related Work

### 7.1 GridPP and the European DataGrid Project

GridPP<sup>24</sup> is a collaboration project between Particle Physicists and Computer Scientist from the UK and Cern aiming to build a Grid for Particle Physics. One of the key components of GridPP is the European DataGrid Project (EDG)<sup>25</sup> which deals with managing a large quantity of sharable data reliably, efficiently and scalably. EDG aims at enabling access to geographically distributed computing and storage facilities. It provides resources to process huge amounts of data from three disciplines: High Energy physics, Biology and Earth Observation. EDG has a file replication service to optimise data access by storing multiple copies of local data at several locations. This replication framework has an optimisation component to minimise file access by pointing access requests to appropriate



replicas and proactively replicating frequently used files based on access statistics.

As compared to DataGrid, BioSimGrid aims to provide a mechanism of data access at a finer granularity level, by delivering a slice of a trajectory rather than a whole file. Hence the concept of file replication of DataGrid can potentially be adopted and modified to suit a finer granularity level of data access.

## 7.2 OGSA-DAI

OGSA-DAI (Open Grid Services Architecture - Data Access and Integration)<sup>26</sup> is a grid middleware that has been developed to allow access and integration of heterogeneous data sources as though they were a single, logical data resource. It is still, however, an ongoing project that is yet to prove its maturity, and for this reason has not been adopted in the BioSimGrid project.

## 8. The Architecture of BioSimGrid

BioSimGrid seeks to fulfil the following criteria in its implementation:

- To minimise data storage, in order to store as many trajectories as possible in a fixed amount of storage space;
- to maximise data transfer rate, in terms of the speed of delivering data to the computational elements, in this case the post processing tools;
- to provide an abstraction of the data layer, where biochemists are freed from the complication of using and understanding data querying languages and the data storage structure in their scientific research;
- to provide a transparency of data location to the users, where actual physical location of the data is hidden.

As shown in figure 7, the architecture of BioSimGrid encompasses three layers: the data storage layer, the middleware and the user interface layer. Each of these will be described in the following sections.

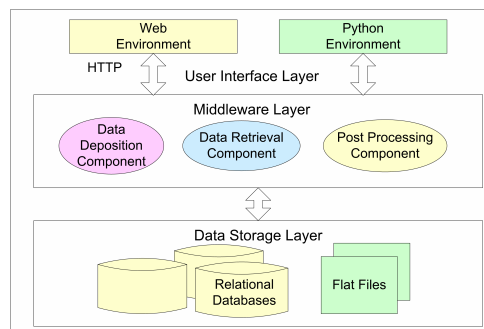


Figure 7. The architecture of BioSimGrid depicting the data storage layer, middleware layer and user interface layer

### 8.1 BioSimGrid Data Storage Layer

The data storage layer is responsible for managing the data on a single machine and exposes methods that are used by the Data Retrieval component to provide the user with data. This layer is required on each machine that is storing trajectory data, initially there will be six remote sites each running this layer. It provides an API that abstracts from the method used to store the data and provides simple access methods for both querying and retrieving data. The trajectory data is divided into two key sections, the metadata and the coordinate data.

#### Trajectory Metadata

The metadata is additional information about the trajectory that can either be supplied by the user, the input files or calculated at a later stage. It also includes the topology that describes the structure of the protein (chains and residues). This metadata is comparatively small and can be replicated across all sites using standard database replication tools. The advantage of replicating the metadata across all sites is so that a user can query all the trajectories stored in the system by querying and single machine and expect a timely response. This design also helps with scalability and load balancing; since the volume of metadata is small, additional nodes can be added to the system and easily incorporated by simply replicating the database. Since each node stores the topology of all the trajectories, users can use any node to query and process data helping to balance the load across the system.

#### Trajectory Coordinate Data

The coordinates for every atom for every timestep are stored resulting in a large volume of data which has to be managed. We have devised a fast, efficient way to store the

coordinates using flatfiles that reduces the storage requirements as well as improving performance results. This flatfile method was implemented using Python pickle<sup>27</sup> and it was compared with a commercial database (DB2) as well as an existing flatfile method (NetCDF<sup>28</sup>). The performance results are shown in Table 1. These results show that a flatfile method is well suited to our application for both random and serial data access.

We selected our own method for flexibility as a whole trajectory is broken into a set of files that are then replicated to at least one other node. This helps to load balance the coordinate data requests as well as provides offsite backups of the data. This abstraction layer also permits the use of different storage methods that can include compression and custom formats, which are completely transparent to the user.

Currently only the coordinates are stored using this method but the next version will store both coordinates as well as velocities.

|                                    | DB2   | NetCDF | Python pickle |
|------------------------------------|-------|--------|---------------|
| Size / GB                          | 7.5   | 3.0    | 3.0           |
| Random Access <sup>a</sup> / s     | 560.8 | 16.4   | 18.6          |
| Sequential Access <sup>b</sup> / s | 389.0 | 4.9    | 5.5           |

<sup>a</sup>A random frame is chosen and then read from a trajectory of 1000 frames. This is completed 1000 times with a different frame chosen each time.

<sup>b</sup>The same trajectory of 1000 frames is read frame by frame from start to finish.

Table 1. Summary of performance results comparing different flatfile methods with a commercial database (DB2)

## 8.2 BioSimGrid Middleware

The middleware of BioSimGrid is implemented on a modular architecture to enable easy extension and future plug in. It is written in Python<sup>27</sup>, a free, open-source and platform-independent high-level object-oriented programming language. Python is chosen for several reasons: (i) the biomolecular simulation community are moving towards Python as the preferred environment for post-processing analysis and several mature post-processing tools written in Python exist already (for example, MMTK<sup>29</sup> and PyMOL<sup>23</sup>). (ii) Python can easily integrate and interface to compiled codes so that other existing tools (typically written in Fortran or C) can be re-used immediately. (iii) Python comes with a substantial set of standard libraries that can be

used in this project and avoid re-coding common tasks.

## Data Deposition Component

The process of depositing a trajectory into the BioSimGrid database is completely automated and the complication of the underlying storage structure is abstracted from the users. One of the challenges is to cater for different simulation packages that produce simulation data in various file formats. To deal with this, the deposition component consists of different parsers for different simulation packages to parse the simulation data files into a generic input object. This object is then parsed through a validator to check for correct data type and their validity against various dictionaries (e.g. the existence of a residue in the dictionary). The process is completed when the validated generic input object is deposited into the flat files (coordinates and velocities) and database (metadata) through an importer. With the modular approach as shown in figure 8, new parsers can easily be added for any new simulation package if required. The underlying complexity of parsing, validating and importing of trajectory into the database is hidden from the users. A biochemist needs only to run five lines of code to deposit their trajectories by specifying the path for their simulation data files, as shown in figure 9.

For the next prototype, the data deposition component will be extended to cater for the distributed nature of the application. We envisage an implementation of multiple deposition points to avoid single point failure and performance bottlenecks. In this case, a global identifier will be assigned to uniquely identify a trajectory and facilitate the synchronisation of multiple metadata databases. To deposit a trajectory from a remote location the generic input object will be serialised at the deposition client and deserialised at the deposition server.

## Data Retrieval Component

The data retrieval component provides a single point of entry for all the trajectories stored on any of the sites. Each site will be running a data retrieval component and a user can use any site to query the data in the entire system. This component queries the local database to retrieve any metadata that is requested, so the user can query information about a trajectory on a different site without having the overhead of contacting the hosting site. This component abstracts the location of






Figure 8. The modular implementation of a data deposition component that includes a set of parsers, a validator and an importer. New parsers can be easily added to this modularised component

```
from bioSim.Settings import UserSettings
from bioSim.Deposit.NAMDDeposit import NAMDDeposit
filenames = {'parameters': '/path/paraFile',
             'topology': '/path/topoFile',
             'coordinates': '/path/coordFile'}
}
uSettings = UserSettings.UserSettings("guest")
NAMDDeposit.NAMDDeposit(uSettings, filenames)
```

Figure 9. An example of a user script to deposit a NAMD trajectory. The underlying complexity of parsing, validating and importing of trajectory into the database is hidden from the users

the trajectory data from the user and is responsible for getting coordinates from external sites if they are not stored locally.

Figure 10 shows how the data is transparently retrieved from a remote site so that it can be used by a user's script. In step 1 and 2 the user submits a script that requests for a set of coordinates from the Data Retrieval Component. This component first looks at the metadata database to retrieve the locations of the requested coordinate flatfiles (step 3). If the data is stored locally then it is returned otherwise a list of remote data source locations are returned to the Data Retrieval Component (step 6). A data source is then selected from the list and a request is made to the Data Retrieval Component on the remote site for the required data (step 7). As this source is listed as a valid data source it is guaranteed to store the data locally, hence it will not attempt to retrieve the data from another remote site. The data is then passed back to the requesting site (step 10) and the Data Retrieval Component returns the data to the user script (step 11) in the same way as a locally stored data set.

There are two key opportunities to save retrieval times when retrieving large amounts of data. The first is to look at the list of sites that store the trajectory and ask multiple sites to provide different parts of the trajectory. This will reduce the load on sites by distributing it across multiple sites as well as improving the speed that data is received.

The second is a cache (not implemented in the current prototype). Each frame that is retrieved from an external site will be stored using the same flatfile storage method. If a

whole trajectory is then cached it can be moved to the main database and marked as a valid location to retrieve data for that trajectory. So when a data query requires data that is not stored on the database then the cache is consulted first to see if it has been retrieved previously (step 4 and 5 in 10) if not then the hosting site is queried. There is a limit to the number of frames that are held in the cache and this is defined by a site-specific limit, which also includes the whole trajectories that are added to the local data store. The aim of storing whole trajectories on additional sites is to attempt to move the data closer to the processing. If a site continually requests a trajectory it makes sense to store the trajectory on that site.

Currently each site has an excess of storage space and we can utilise this space to gain a performance boost. However, more trajectories can still be added as temporary trajectories can simply be deleted and removed from the metadata database to make more room as required.

The Data Retrieval Component is not only responsible for getting the data from the distributed sources but it is also responsible for making the data transparently available to the users in an environment of their choice, in this case Python. This results in Python numeric arrays being made available to users who have no idea where or how the data is stored. This has currently been implemented and a series of analysis tools for the Post Processing Component have been built on this design. This design also permits extensions for other

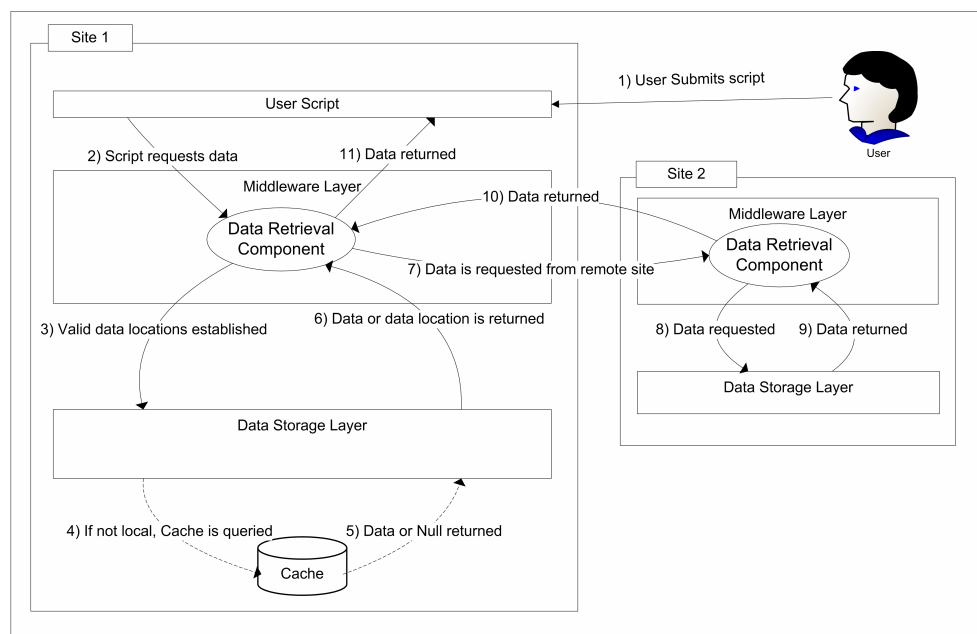


Figure 10. Schematic showing how a remote data request is returned. Caching mechanism is used to improve the speed of data transfer.

languages like Perl to assist the users to migrate and utilise the BioSimGrid project.

### Post Processing Component

For the post processing component, a set of analysis tools are written for standard and generic analysis on the simulation data, e.g. the calculation of Root Mean Square Deviation (RMSD) and the computation of the average structure and interatomic distances. Each analysis is exposed as a module and the modularity approach enables the tool set to be extended easily. An example of an analysis script is shown in figure 11 to demonstrate how to use the post processing tools. The fourth line of the script initialises the user settings. The fifth line specifies the setting for a frame collection – the part of the protein to be used to perform the analysis, in this case frames 100 to 200 from trajectory ‘BioSimGrid\_GB-STH\_1’. The seventh line requests a RMSD analysis by taking the frame collection as its input parameter. Finally, the last line specifies the format of the result to be generated, which in this case is the output of an image file in PNG format. The ease of selecting different data set and different post processing tools allows biochemists with little computational experience to perform an analysis on the simulation data and obtain meaningful results.

### 8.3 User Interface Layer

The BioSimGrid user application level offers two modes of interaction: via a graphical web based interface or via the Python scripting environment. The graphical interface is just another layer on top of the underlying Python codes. The scripting environment caters for advanced users who would like to connect to BioSimGrid in a scripting environment and utilise its data submission, retrieval and post-processing API in a fully programmable way. In this environment, biochemists can choose to run existing analysis toolkits provided by BioSimGrid. Alternatively, for more specific analysis, they can use the available data retrieval packages to write their own script. The graphical interface provides a more user-friendly environment to cater for novice users. It allows users to perform standard analysis runs and provides an overview of the available data and processing options. In this mode, a user first selects an analysis from a drop down menu, then proceeds to select a trajectory and the relevant frames on which to perform the analysis. All these operations are done by clicks of buttons on a web browser.

## 9. Current Issues and Next Prototype

BioSimGrid is in its early stage of development. Current prototypes that have been developed are based on architecture where both the application

```

from bioSim.DataRetrieval import FrameCollection, FCSettings
from bioSim.Analysis import RMSD
from bioSim.Settings import UserSettings
u = UserSettings.UserSettings('guest')
f = FCSettings.FCSettings(u, [['BioSimGrid_GB-STH_1', range(100, 201)]])
fc = FrameCollection.FrameCollection(f)
myRMSD = RMSD(fc)
myRMSD.createPNG()

```

Figure 11. An example of a user script to run a RMSD analysis using frames 100 to 200 from trajectory 'BioSimGrid\_GB-STH\_1']

and database server are implemented as client server architecture, running at a single location. We have modularised our components and have developed a basic set of functionalities of BioSimGrid for data deposition, data retrieval and analysis of post simulation data. The modularity approach of the components enables easy plug-in and future extension of various functionalities, such as adding more analysis tools or extending the data deposition tools to cater for new simulation result formats.

The next prototype of BioSimGrid will concentrate on tackling the geographically distributed databases and applications. Establishing secure asynchronous network communication, handling data latency and data recovery is non trivial in this case. We are investigating Python twisted framework<sup>30</sup> and Pyro<sup>31</sup> for programming network services and applications. For a more reliable data transmission, the next prototype will incorporate MD5<sup>32</sup> hashes to help manage corruptions in file transfer. We also envisage the use of standard protocols such as secure socket layer (openSSL) to provide secure point to point communication.

The issue of security is also a major concern in BioSimGrid. We envisage the use of digital certificate-based authentication to authorise users into the system and provide mechanism to set various permission levels for different user groups to authorise them to different resources and transactions.

In the future work, we plan to implement web service based interfaces in order to provide a platform and language independent way of accessing the existing middleware components.

## 10. Summary - BioSimGrid

In summary, BioSimGrid provides a trajectory storage system that allows users to submit simulation data from a wide range of simulation packages and to run cross simulation comparisons independent of the source of the data. We have developed the current version of the system together with biochemists who provide constant feedback on the usability of

the project, and we are currently expanding the user base and the number of available trajectories in the system.

## 11. Conclusion

Advanced computational methods and Grid computing are finding increasing use in the area of the life sciences. In the particular context of biomolecular computer simulations, we have extended the basic distributed computing model to the situation where the calculations are coupled, through the addition of a dedicated Beowulf cluster to catch-up on delayed simulations. This approach does yield an improvement in the overall simulation efficiency. We have also reported the development of a database for the storage and analysis of the large trajectories produced by these simulations. This database will not only allow for extensive and valuable comparisons to be made between related simulations, thereby yielding more a more reliable biochemical interpretation, but will also allow data to be readily shared between laboratories.

## Acknowledgements

For the work on distributed computing, we thank R. Gledhill, A. Wiley and L. Fenu for discussions and the EPSRC for funding comb-e-chem. For BioSimGrid, we would like to thank our collaborators D. Moss, C. Laughton, L. Caves, O. Smart and A. Mulholland. This project is funded by BBSRC.

## References

1. <http://setiathome.ssl.berkeley.edu>
2. <http://www.chem.ox.ac.uk/curecancer.html>
3. <http://www.foldingathome.org>
4. Leach, A.R., *Molecular Modelling, Principals and Applications*, Longman, Harlow, UK, 1996
5. Pelton, J.G., Kustu, S., Wemmer, D.E., *J. Mol. Biol.*, 292, 1095, 1999

6. Kern, D., Volkman, B.F., Luginbuhl, P., Nohaile, M.J., Kustu, S., Wemmer, D.E., *Nature*, 402, 894, 1999
7. Sugita, Y., Kitao, A., Okamoto, Y., *J. Chem. Phys.*, 113, 6042-6051, 2000
8. Hansmann, U.H.E., *Chem. Phys. Lett.*, 281, 140, 1997
9. <http://www.rcsb.org/pdb/>
10. Vriend, G., Hooft, R.W.W., Van Aalten, D., WhatIf, 1997
11. Pearlman, D.A., Case, D.A., Caldwell, J.W., Ross, W.S., Cheatham, T.E., Debolt, S., Ferguson, D., Seibel, G., Kollman, P., *Comput. Phys. Commun.*, 91, 1, 1995
12. Mackerrell, A.D., Bashford, D., Bellot, M., Dunbrack, L., Evanseck, M., Fischer, S., Gao, J., Guo, H., Ha, S., Joseph-Mccarthy, D., Kuchnir, L., Kuczera, K., Lau, F.T.K., Mattos, C., Michnick, S., Ngo, T., Nguyen, D.T., Prodhom, B., Reiher, W.E., Roux, B., Sclenkrich, M., Smith, J.C., Stote, R., Straub, J., Watanabe, M., Wiorkiewicz-Kuczera, J., Yin, D., Karplus, M., *J. Phys. Chem. B.*, 102, 3586, 1998
13. Kale, L., Skeel, R., Bhandarker, M., Brunner, R., Gursoy, A., Krawetz, N., Phillips, J., Shinozaki, A., Varadarajan, K., Schulten, K., *J. Comp. Phys.*, 151, 283, 1999. NAMD was developed by the Theoretical Biophysics Group in the Beckman Institute at Urbana-Champaign.
14. Paterlini, M.G., Ferguson, D.M., *Chem. Phys.*, 236, 243, 1998
15. Feller, S.E., Zhang, Y.H., Pastor, R.W., Brooks, B.R., *J. Chem. Phys.*, 103, 4613, 1995
16. Ryckaert, J.P., Ciccotti, G., Berendsen, J.C., 23, 327, 1977
17. Darden, T., York, D., Pedersen, L., *J. Chem. Phys.*, 98, 10089, 1993
18. Litzkow, M., Turning Idle Workstations into Cycle Servers, in Usenix Summer Conference, 1987, pages 381-384; Litzkow, M., Livny, M., Mutka, M., Condor – a hunter of idle workstations, in 8<sup>th</sup> International Conference of Distributed Computing Systems, 1988, pages 104-111.
19. Wu, B., Tai, K., Murdock, S., Ng, M.-H., Johnston, S., Fangohr, H., Jeffreys, P., Cox, S., Essex, J.W., Sansom, M.S.P., in Simon J. Cox, editor, Proceedings of UK e-science All Hands Meeting 2003, pages 412-419, Swindon, 2003. EPSRC.
20. Tai, K., Murdock, S., Wu, B., Ng, M.-H., Johnston, S., Fangohr, H., Cox, S.J., Jeffreys, P., Essex, J.W., Sansom, M.S.P., *Org. Biomol. Chem.* DOI: 10.1039/b411352g
21. <http://www.biosimggrid.org>
22. Dalke, A., Humphrey, W., Schulten, K., *J. Mol. Graph.* 14, 33, 1996
23. DeLano, W.L., The PyMOL molecular graphics system. DeLano Scientific, 2002. <http://www.pymol.org>
24. <http://www.gridpp.ac.uk>
25. <http://www.eu-datagrid.org>
26. <http://www.ogsa-dai.org.uk>
27. Drake Jr., F.L., van Rossum, G. Python library reference. Computer Science Department of Algorithmics and Architecture, CS-R9524, 1995. <http://www.python.org>
28. <http://my.unidata.ucar.edu/content/software/netcdf/index.html>
29. <http://starship.python.net/crew/hinsen/MMTK>
30. <http://www.twistedmatrix.com>
31. <http://pyro.sourceforge.net/index.html>
32. <http://www.faqs.org/rfcs/rfc1321.html>