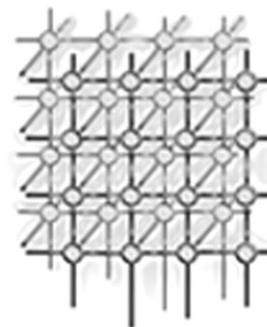


The Virtual Laboratory: a toolset to enable distributed molecular modelling for drug design on the World-Wide Grid



Rajkumar Buyya^{1,*,\dagger}, Kim Branson², Jon Giddy³ and
David Abramson⁴

¹*Grid Computing and Distributed Systems Laboratory, Department of Computer Science and Software Engineering, The University of Melbourne, Parkville, Melbourne, Australia*

²*Structural Biology, Walter and Eliza Hall Institute, Royal Parade, Parkville, Melbourne, Australia*

³*Welsh e-Science Centre, Department of Computer Science, Cardiff University, Cardiff, U.K.*

⁴*School of Computer Science and Software Engineering, Monash University, Caulfield Campus, Melbourne, Australia*

SUMMARY

Computational Grids are emerging as a new paradigm for sharing and aggregation of geographically distributed resources for solving large-scale compute and data intensive problems in science, engineering and commerce. However, application development, resource management and scheduling in these environments is a complex undertaking. In this paper, we illustrate the development of a Virtual Laboratory environment by leveraging existing Grid technologies to enable molecular modelling for drug design on geographically distributed resources. It involves screening millions of compounds in the chemical database (CDB) against a protein target to identify those with potential use for drug design. We have used the Nimrod-G parameter specification language to transform the existing molecular docking application into a parameter sweep application for executing on distributed systems. We have developed new tools for enabling access to ligand records/molecules in the CDB from remote resources. The Nimrod-G resource broker along with molecule CDB data broker is used for scheduling and on-demand processing of docking jobs on the World-Wide Grid (WWG) resources. The results demonstrate the ease of use and power of the Nimrod-G and virtual laboratory tools for grid computing. Copyright © 2003 John Wiley & Sons, Ltd.

KEY WORDS: Grid computing; drug design; application scheduling; grid economy

*Correspondence to: Rajkumar Buyya, Grid Computing and Distributed Systems Laboratory, Department of Computer Science and Software Engineering, The University of Melbourne, Parkville, Melbourne, Australia.

[†]E-mail: rajkumar@buyya.com



1. INTRODUCTION

Computational Grids [1] enable the sharing of a wide variety of geographically distributed resources including supercomputers, storage systems, databases, data sources and specialized devices owned by different organizations in order to create virtual enterprises and organizations. They allow selection and aggregation of distributed resources across multiple organizations for solving large-scale computational and data intensive problems in science, engineering and commerce. The parallel processing of applications on wide-area distributed systems provide a scalable computing power. This enables exploration of large problems with huge data sets, which is essential for creating new insights into the problem. Molecular modelling for drug design is one of the scientific applications that can benefit from the availability of a large computational capability.

Drug discovery is an extended process that can take as many as 15 years from the first compound synthesis in the laboratory until the therapeutic agent or drug, is brought to market [2]. Reducing the research timeline in the discovery stage is a key priority for pharmaceutical companies worldwide. Many such companies are trying to achieve this goal through the application and integration of advanced technologies such as computational biology, chemistry, computer graphics and high-performance computing (HPC). Molecular modelling has emerged as a popular methodology for drug design—it can combine computational chemistry and computer graphics. Molecular modelling can be implemented as a master–worker parallel application, which can take advantage of HPC technologies such as clusters [3] and Grids for large-scale data exploration.

Drug design using molecular modelling techniques involve screening a very large number (of the order of a million) of ligand[‡] records or molecules of compounds in a chemical database (CDB) to identify those that are potential drugs. This process is called molecular docking. It helps scientists to predict how small molecules, such as substrates or drug candidates, bind to an enzyme or a protein receptor of known three-dimensional (3D) structure (see Figure 1). Docking each molecule in the target chemical database is both a compute and data intensive task. It is our goal to use Grid technologies to provide cheap and efficient solutions for the execution of molecular docking tasks on large-scale, wide-area parallel and distributed systems.

While performing docking, information about the molecule must be extracted from one of the many large chemical databases. As each chemical database requires storage space in the order of hundreds of megabytes to terabytes, it is not feasible to transfer the chemical database to all resources in the Grid. Also, each docking job only needs a ligand or module record, not the whole database. Therefore, access to a chemical database must be provided as a network service (see Figure 2). The chemical databases need to be selectively replicated on a few nodes to avoid any bottleneck due to providing access to the database from a single source. Intelligent mechanisms (e.g. CDB broker) need to be supported for selecting optimal sources for CDB services depending on the location of resources selected for processing docking jobs.

Fundamentally, drug design is a computational and data challenge problem since it involves screening millions of compounds in chemical databases. Screening each compound, depending on structural complexity, can take from a few minutes to hours on a standard PC, which means

[‡]An ion, a molecule or a molecular group that binds to another chemical entity to form a larger complex.

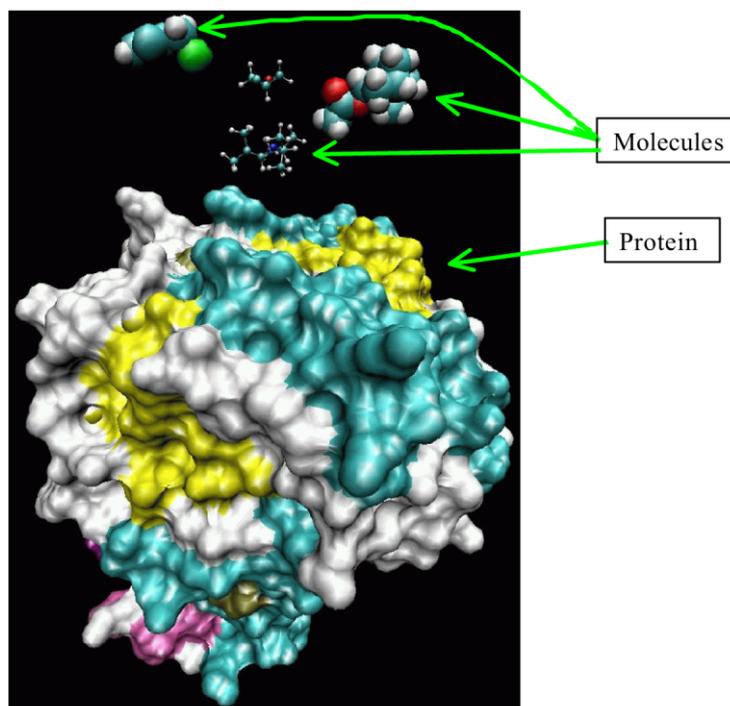


Figure 1. X-ray crystal structure of a target protein receptor and small molecules to be docked.

screening all compounds in a single database can take years! For example, we are looking into a drug design problem that involves screening 180 000 compounds. Each job screening a compound is expected to take up to three hours of execution time on a desktop computer (e.g. Pentium-based Linux/Windows PC). That means, if we aim to screen all these compounds on a single PC, it can take up to 540 000 hours, which is roughly equivalent to 61 years! If we use a typical cluster-based supercomputer with 64 nodes, we can solve this problem in one year. The problem can be solved with a large scale Grid of hundreds of supercomputers within a day. If we use a massive network of peer-to-peer style Grid computing infrastructure such as SETI@Home [4], the drug discovery problem could be solved within a few hours.

The rest of this paper is organized as follows. A high-level operational model for molecular modelling on the Grid is presented in Section 2. A layered architecture for building the Virtual Laboratory environment is presented in Section 3. It leverages the existing Grid technologies and supports new tools that are essential for Grid-enabling the chemical database and the docking application on distributed resources. Formulation of molecular docking as a parameter sweep application is presented in Section 4. The results of two experiments on scheduling molecular docking

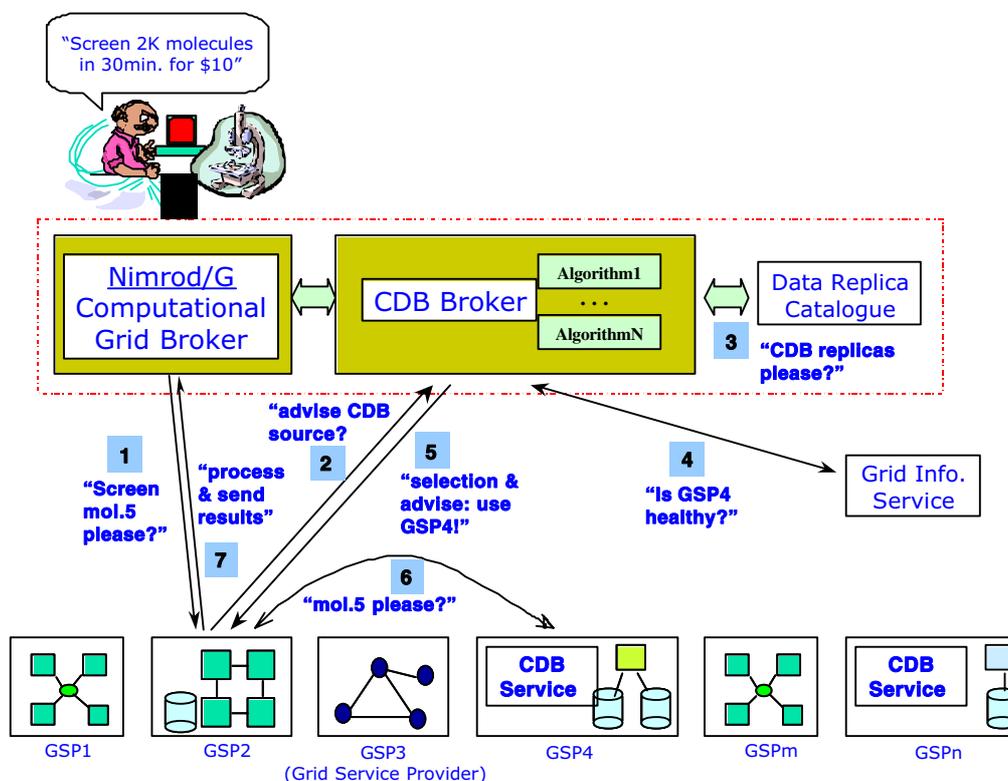


Figure 2. Resource brokering architecture for screening molecules on distributed resources.

jobs for processing on the World-Wide Grid (WWG) [2] testbed resources are presented in Section 5. The final section summarizes the paper along with suggestions for future works.

2. OPERATIONAL MODEL

The Virtual Laboratory tools transform the existing molecular modelling application (without the need for making any changes to it) into a parameter sweep application for executing jobs docking molecules in the CDBs in parallel on distributed resources. The parametrized application contains multiple independent jobs, each screening different compounds to identify their drug potential. These jobs are computationally intensive in nature and only a small proportion of the execution time is spent on data communication (e.g. fetching molecular information on demand from remote databases). Applications expressed with this task-farming computational model have high computation to communication ratio.



Hence, they can tolerate high network latency, which makes them suitable for executing in parallel on Internet-wide distributed resources.

A high-level operation model of docking molecules on the Grid is shown in Figure 2. The drug designer formulates the molecular docking problem, submits the application to the Grid resource broker (e.g. Nimrod-G [6,7]) along with performance and optimization requirements—‘screen 2000 molecules within 30 min and the available budget for processing is \$10’ [7]. The broker discovers resources, establishes their cost and capability, and then prepares a schedule to map docking jobs to resources. Let us say, it identified a Grid Service Provider (GSP), say GSP2 and assigned a job of screening a molecule 5 to it. A job has a task specification that specifies a list of operations to be performed. To process a job on GSP2, the broker dispatcher deploys its Agent on resource GSP2. The agent executes a list of commands specified in the job’s task specification. A typical task specification contains necessary commands to copy executables and input files from the user machine, substitution of parameters declared in the input file, execution of the program and finally copying results back to the user. It can also contain special commands for accessing the input data from the remote database. For example, a docking task can contain a special command (e.g. an instruction to fetch molecule record from the CDB) to make a request to the data broker (e.g. CDB broker) for a molecule record. The data broker looks at the replica catalogue for a list of sites providing CDB services, checks the status of those sites, and selects a suitable site (e.g. a node with fast network connectivity) and recommends the same. The molecule fetch command can then request the CDB service provider for a molecule record and write the molecule structure to a file that acts as an input to the docking program. After executing the docking program, the agent executes commands related to copying docking results to the user home node.

3. ARCHITECTURE—THE SOFTWARE STACK

The Virtual Laboratory builds on the existing Grid technologies and tools for performing data intensive computing on distributed resources. It provides new tools for managing and accessing remote chemical databases as a network service. There are many scientific and commercial applications (e.g. molecular modelling, high-energy physics events processing and financial investment risk-analysis) that explore a range of scenarios. Instead of explicitly developing them as parallel applications using interfaces such as MPI, they can be composed as parameter sweep applications using tools such as Nimrod [8]. Such application jobs can be executed in parallel on distributed resources using the Nimrod-G resource broker (see Figure 2). A layered architecture and the software stack essential for performing molecular modelling on distributed resources is depicted in Figure 3. The components of the Virtual Laboratory software stack are:

- the DOCK software for molecular modelling [9];
- the Nimrod Parameter Modelling Tools [8,10] for enabling DOCK as a parameter sweep application;
- the Nimrod-G Grid Resource Broker [6,7] for scheduling DOCK jobs on the Grid;
- chemical database (CDB) management and intelligent access tools:
 - CDB database lookup/index table generation;
 - CDB and associated index-table replication;

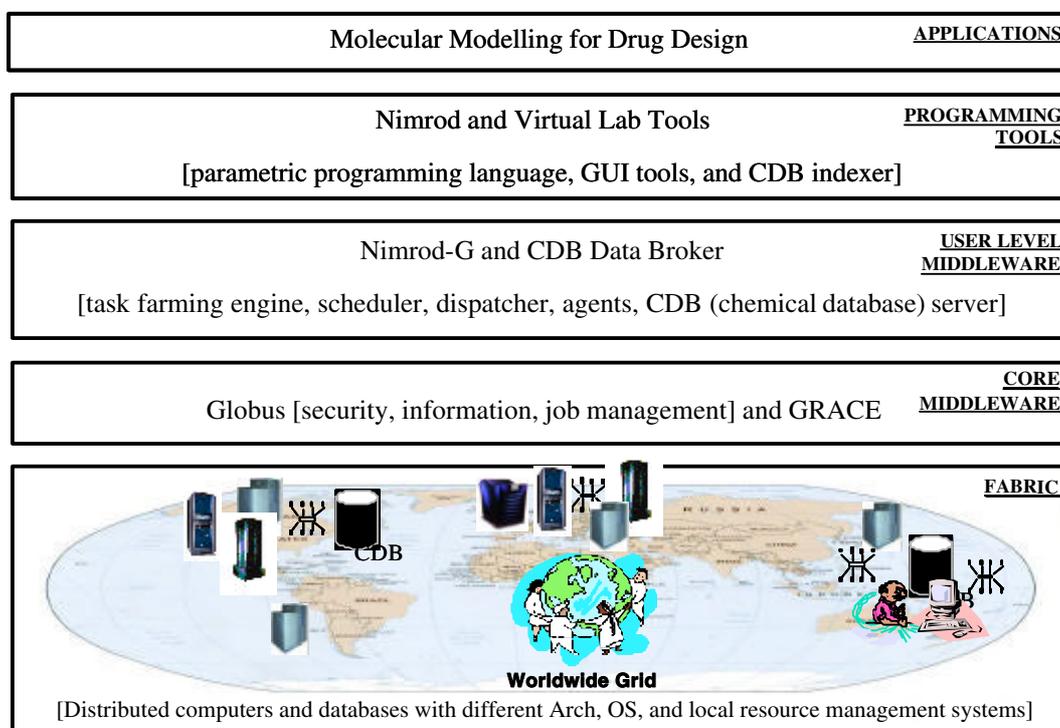


Figure 3. Layered architecture of Virtual Laboratory for drug design.

- CDB replica catalogue for CDB resource discovery;
- CDB servers for providing CDB services;
- CDB broker for selecting a suitable CDB service (replica selection);
- CDB clients for fetching molecular records (data movement);
- the GrACE software for resource trading toolkit [11];
- the Globus middleware for secure and uniform access to distributed resources [12].

The Grid resources (e.g. multiprocessors or clusters) at each location are generally presented as a single entity using resource management systems such as OS-fork, LSF, Condor and SGE.

3.1. Docking code

The original docking code developed by researchers at the University of California in San Francisco (UCSF) is one of the most popular molecular docking applications [13]. The program evaluates the

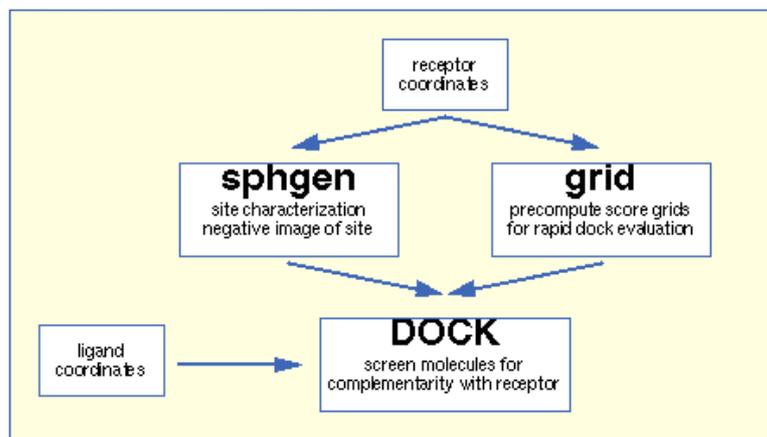


Figure 4. Relation between key programs in the *dock* suite.

chemical and geometric complementarities between a small molecule and a macromolecular binding site. It explores ways in which two molecules, such as a drug and an enzyme or protein receptor, might fit together. Compounds that might bind tightly to the target receptor must have complementary chemical and spatial natures. Thus docking can be seen as a 3D puzzle searching for pieces that will fit into the receptor site. It is important to be able to identify small molecules (compounds), which may bind to a target macromolecule. This is because a compound, which binds to a biological macromolecule, may modulate its function and with further development eventually become a drug candidate. An example of such a drug is the anti-influenza drug Relenza which functions by binding to influenza virus attachment proteins thus preventing viral infection.

The relationship between the key programs in the dock suite is depicted in Figure 4 [13]. The receptor coordinates at the top represent the 3D structure of a protein. The molecular modeller identifies the active site and other sites of interest, and uses the program *sphgen* to generate the sphere centres, which fill the site [14]. The program *grid* generates the scoring grids [9]. The program *dock* matches spheres (generated by *sphgen*) with ligand atoms and uses scoring grids (from *grid*) to evaluate ligand orientations [9,14]. It also minimizes energy-based scores [15]. The focus of our work is on docking molecules in CDB with receptors to identify potential compounds that act as a *drug*. Hence, discussion in this paper is centred on the execution of the program *dock* as a parameter sweep application on world-wide distributed resources.

The docking code is highly portable—we have been able to compile and produce executables for Sun-Solaris, PC Linux, SGI IRIX and Compaq Alpha/OSF1 architectures. For docking on heterogeneous resources, the Nimrod-G broker selects the correct executable automatically based on the resources it discovers at runtime.



3.2. Nimrod-G tools

The Nimrod-G toolkit provides a suite of tools and services for creating parameter sweep applications, managing resources and scheduling applications on the world-wide distributed resources. It supports a simple declarative programming language and GUI tools for creating scripts and parametrization of application input data files and a programmable grid resource broker for processing jobs on grid resources.

3.2.1. Tools for creating parameter sweep applications

The Nimrod-G toolkit provides a suite of tools and services for creating parameter sweep applications, managing resources and scheduling applications on the world-wide distributed resources. It provides a simple declarative programming language or GUI tools for parametrization of application input data files and creation of task-script to be performed by each job; and a programmable grid resource broker for processing jobs on Grid resources.

The steps involved in distributed parametric execution are:

- (i) parametrize input files;
- (ii) prepare a plan file containing the commands that define parameters and their values;
- (iii) generate a run file, which converts the generic plan file to a detailed list of jobs;
- (iv) schedule jobs for processing on distributed machines; and
- (v) control and monitor the execution of the jobs.

The application execution environment handles online creation of input files and command line arguments through parameter substitution. The GUI tools supported by enFuzion, a commercial version of Nimrod, can also be used for parametrizing applications. enFuzion uses the same syntax as Nimrod for parameter specification [10]. Both Nimrod and enFuzion have been successfully used for performing parameter studies in a single administrative domain such as clusters. Nimrod-G [6,7] extends the capabilities of Nimrod and EnFuzion with the addition of powerful resource discovery, trading and scheduling algorithms [16]. In Section 4, we discuss the capabilities of Nimrod-G tools by composing a molecular modelling program as a parameter sweep application for docking compounds in CDB databases and processing docking jobs on the Grid.

3.2.2. Nimrod-G Grid resource broker for scheduling DOCK jobs on Grid

The Nimrod-G resource broker identifies the user and application processing requirements and selects Grid resource combination in such a way that the user requirements are met [6,7,17]. It performs resource discovery, selection, scheduling, dispatching of docking jobs to remote resource, starting and managing the execution of jobs and gathering results back to the home node. The sub-modules of our resource broker are:

- a persistent task farming engine;
- a grid explorer for resource discovery;
- a resource trading manager for establishing access price;
- a scheduler that maps jobs to resources using scheduling algorithms;

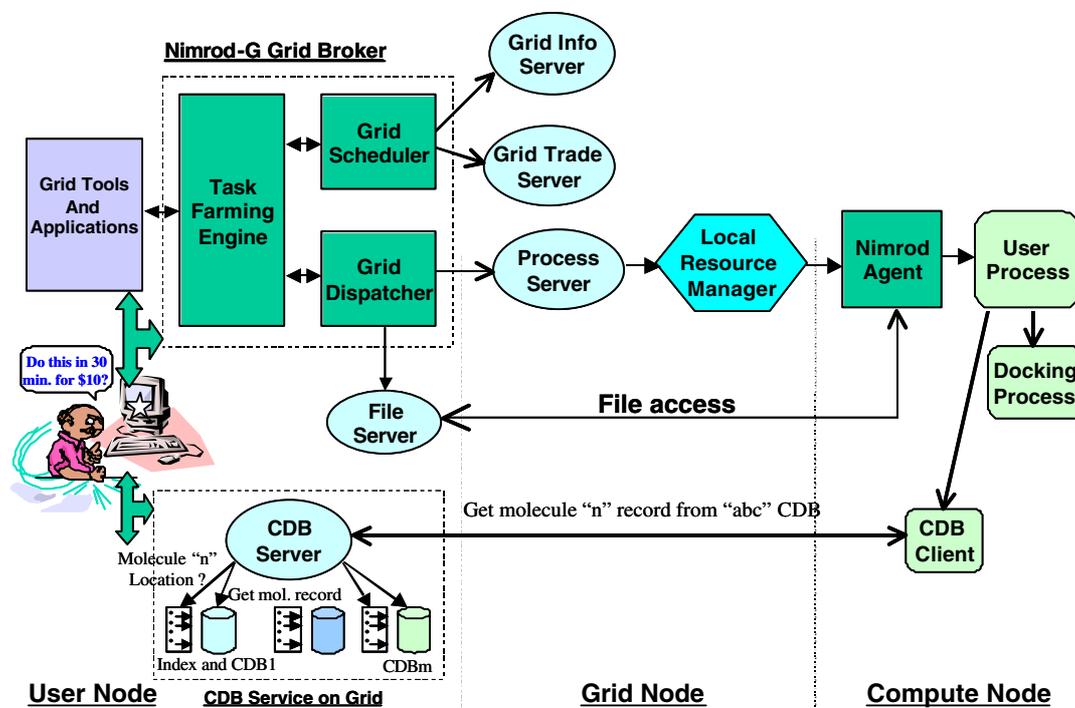


Figure 5. Deployment of Virtual Laboratory components at runtime and their interaction.

- a dispatcher and actuators for deploying agents on grid resources; and
- agents for managing execution of Nimrod-G jobs on grid resources.

When Nimrod-G deploys its agents on the Grid node, it is submitted to the local resource manager, which then allocates a compute node[§] to it for executing the job.

The Nimrod-G components, grid explorer for resource discovery and dispatcher for submitting jobs to remote resource, are implemented using the Globus GIS and GRAM services. The dispatcher actually initiates the execution of Nimrod-G agents on remote resources, which takes care of all the operations associated with the execution of an assigned job. The interaction between the components of the Nimrod-G runtime machinery and Grid services during runtime is shown in Figure 5. The Nimrod-G broker supports deadline and budget constrained (DBC) scheduling algorithms driven by a computational economy and user requirements [16]. In Section 5, we discuss the results of the

[§]It is a node/CPU that the local resource manager allocates to the Nimrod-G agent for job processing.



Nimrod-G broker scheduling a molecular modelling application on the Grid with DBC time and cost optimization scheduling algorithms.

3.3. Chemical database management and intelligent access tools

The chemical databases contain records of a large number of molecules from commercially available organic synthesis libraries and natural product databases. The molecules in the CDB are represented in MOL2 file (.mol2) format [18], which is a portable representation of a SYBYL [19] molecule. The MOL2 file is an ASCII file that contains all the information needed to reconstruct a SYBYL molecule. Each ligand record in a chemical database represents the 3D structural information of a compound. The numbers of compounds in each CDB can be in the order of tens of thousands and the database size be anywhere from tens of megabytes to gigabytes and even terabytes. We have developed tools for turning the CDB into a network service and accessing them from remote resources. They include tools for indexing ligand records in the CDB, a multithreaded CDB Server for serving requests for molecule records, and a tool for fetching molecule/ligand record from remote CDB via the network [20].

When a chemical database is available from more than one source (replica site), a suitable strategy such as a source with high network speed or lightly loaded, can be used for selecting one of them. It is likely that multiple users from different locations issue requests for accessing the CDB, the server should be able to process such simultaneous requests concurrently. Therefore, we have developed a multithreaded CDB server that can service requests from multiple users concurrently. An interaction between a Grid node and a node running the CDB server while performing docking is shown in Figure 5. We developed and implemented protocols shown in Figure 6 for interaction between the CDB clients and the server. Both figures illustrate the operational model and the flow of control between CDB clients and servers.

When the Nimrod-G schedules a docking job for processing on one of the Grid resources, it actually submits an agent to Globus GRAM, which acts as a process server. The process server either executes it by forking it as process on time shared system or submits to the site resource manager such as PBS, LSF and Condor, which allocates a compute node for the agent and starts its execution. The agent contacts the Nimrod-G dispatcher for job task information, which contains instructions for executing a job. It copies input files, performs parameter substitution, executes programs (e.g. CDB client to fetch a molecule record from the remote CDB server and docking program) and ships results back to the Nimrod-G user. When the CDB server receives a request for molecule record, it reads the molecule record from the chemical database and sends back to the client.

Instead of searching molecule records sequentially in a database, we have built tools for creating index-tables for each CDB represent using the MOL2 format along with the record size information. The CDB index file, organized in binary format, contains the starting address (byte location) of a molecule record and record size of all molecules in sequence. When a molecule record is requested, the CDB server first looks at the CDB index file to identify the record location and its size. It then fetches the molecule record from the CDB file with a single read operation and thus improving the access and response speed.

It is possible to screen virtual combinatorial databases in their entirety. This methodology allows only the potential compounds to be subjected to physical (manual) screening and/or synthesis in laboratories, which is extremely time-consuming and resource-intensive.

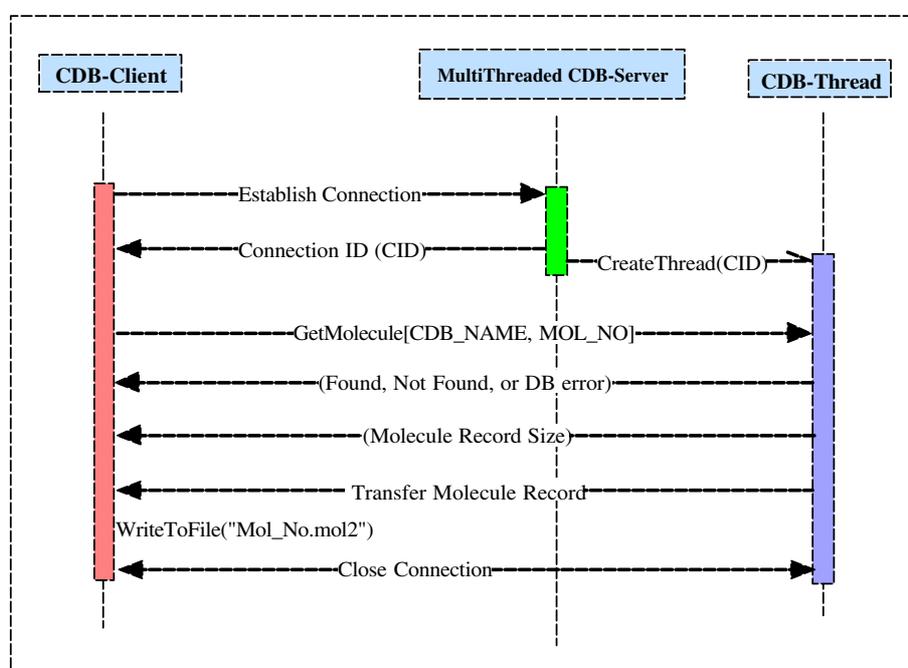


Figure 6. Protocols for interaction between the CDB clients and the server.

4. APPLICATION COMPOSITION

A docking application having an ability to screen a molecule for each execution can be composed as a task-farming, parameter sweep application for distributed execution. This can be achieved by using Nimrod-G parameter specification language to parametrize docking application input data and files. There is no need to make any changes to the existing (sequential) docking application neither does it need to be developed as parallel application explicitly for distributed execution. The users just need to parametrize the input data and files appropriately and define a Nimrod-G plan file once. Note the value of parameters can be changed while launching the application execution. The plan file specifies the parametric tasks and the types of the parameters and their values for these tasks. A parametric task consists of a script defined using a sequence of simple commands, providing the ability to copy files to and from the remote node, perform parameter substitutions in input files, execute certain programs and copy output files back to the user home node. The parametric plan can be submitted to the Nimrod-G runtime machinery, which creates independent docking jobs and schedules jobs for concurrent execution on distributed resources. It takes care of replacing the actual value of parameters in the parametrized input files before executing docking jobs.



```

score_ligand          yes
minimize_ligand       yes
multiple_ligands      no
random_seed           7
anchor_search         no
torsion_drive         yes
clash_overlap         0.5
conformation_cutoff_factor 3
torsion_minimize      yes
match_receptor_sites  no
random_search         yes
. . . . .
. . . . .
maximum_cycles        1
ligand_atom_file      $1.mol2
receptor_site_file    ece.sph
score_grid_prefix     ece
vdw_definition_file   parameter/vdw.defn
chemical_definition_file parameter/chem.defn
chemical_score_file   parameter/chem_score.tbl
flex_definition_file  parameter/flex.defn
flex_drive_file       parameter/flex_drive.tbl
ligand_contact_file   dock_cnt.mol2
ligand_chemical_file  dock_chm.mol2
ligand_energy_file    dock_nrg.mol2

```

Molecule to be screened

Figure 7. A configuration input file for docking application.

A sample configuration input file of the docking application is shown in Figure 7. It specifies docking configuration parameters and molecule to be docked by indicating a name of the file in which molecule record is stored using the parameter variable 'ligand_atom_file'. To perform a parameter sweep of different molecules, the value specified by the parameter variable 'ligand_atom_file' needs to be parametrized. This is accomplished by replacing the current value, which represents the name of a file containing molecule record, by a substitution place marker. The place marker T consists of a dollar-sign (\$) followed by the name of the parameter controlling the substitution, optionally surrounded by braces.

A parametrized input file with several attributes replaced by substitution place markers is shown in Figure 8. For example, a place marker called for the parameter 'ligand_number' has replaced the first part of the 'ligand_atom_file' attribute value. The actual value of these parameters is defined in the Nimrod-G plan file that contains parameter definition and task specification. The parameter definition section of the plan file is shown in Figure 9. Each parameter is defined by a keyword 'parameter', followed by the parameter name, an optional label and a parameter type. The remaining information on each line defines valid values for the parameter.

The parameter, for example, 'database_name' has a label and is of type *text*. Its valid values are listed and the user will be able to select one of the values for the duration of the entire experiment. Most of



```
score_ligand           $score_ligand
minimize_ligand       $minimize_ligand
multiple_ligands      $multiple_ligands
random_seed           $random_seed
anchor_search         $anchor_search
torsion_drive         $torsion_drive
clash_overlap         $clash_overlap
conformation_cutoff_factor $conformation_cutoff_factor
torsion_minimize     $torsion_minimize
match_receptor_sites $match_receptor_sites
random_search         $random_search
. . . . .
. . . . .
maximum_cycles       $maximum_cycles
ligand_atom_file     ${ligand_number}.mol2
receptor_site_file   $HOME/dock_inputs/${receptor_site_file}
score_grid_prefix    $HOME/dock_inputs/${score_grid_prefix}
vdw_definition_file  vdw.defn
chemical_definition_file chem.defn
chemical_score_file  chem_score.tbl
flex_definition_file flex.defn
flex_drive_file      flex_drive.tbl
ligand_contact_file  dock_cnt.mol2
ligand_chemical_file dock_chm.mol2
ligand_energy_file   dock_nrg.mol2
```

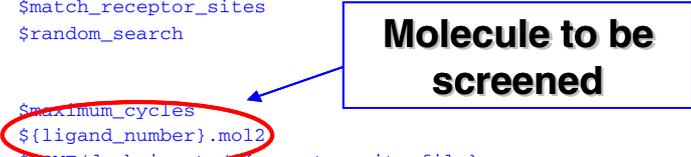


Figure 8. Parametrization of a configuration input file.

the remaining parameters are single values, either text strings or integers, selected by the user, but with default values provided if the user does not wish to choose a value.

The range parameter, 'ligand_number', used to select the molecule, is defined as an integer variable with bounds. For example, to process the first 2000 molecules in the CDB, this range parameter can vary from 1 to 2000 with the step size of 1.

The parameters 'receptor_site_file' and 'score_grid_prefix' indicate the data input files. Their values indicate that data input files are located in the user home directory on Grid nodes. Instead of pre-staging, these files can be copied at runtime by defining necessary 'copy' operations in the job's 'nodestart' or 'main' task (see Figure 10). However, it is advisable to copy or 'pre-stage' large input files in the beginning of application execution instead of copying them during execution of every job. This saves transmission time particularly when those files are going to be used for docking with many databases.

The plan file is submitted to a job generation tool, such as the EnFuzion Generator, in order to create a run file that contains specific instances of jobs to be run, which is then submitted to the Nimrod-G runtime machinery for processing on the Grid. The run file contains a job for each combination of parameters. Hence the number of jobs is the product of the number of values chosen for each parameter. Since most of the parameters except 'ligand_number' are single-valued, they have no effect on the number of jobs.



```

parameter database_name label "database_name" text select oneof "aldrich"
"maybridge" "maybridge_300" "asinex_egc" "asinex_epc" "asinex_pre"
"available_chemicals_directory" "inter_bioscreen_s" "inter_bioscreen_n"
"inter_bioscreen_n_300" "inter_bioscreen_n_500" "biomolecular_research_institute"
"molecular_science" "molecular_diversity_preservation"
"national_cancer_institute" "IGF_HITS" "aldrich_300" "molecular_science_500"
"APP" "ECE" default "aldrich_300";
parameter CDB_SERVER text default "bezek.dstc.monash.edu.au";
parameter CDB_PORT_NO text default "5001";
parameter score_ligand text default "yes";
parameter minimize_ligand text default "yes";
parameter multiple_ligands text default "no";
parameter random_seed integer default 7;
parameter anchor_search text default "no";
parameter torsion_drive text default "yes";
parameter clash_overlap float default 0.5;
parameter conformation_cutoff_factor integer default 5;
parameter torsion_minimize text default "yes";
parameter match_receptor_sites text default "no";
. . . . .
. . . . .
parameter maximum_cycles integer default 1;
parameter receptor_site_file text default "ece.sph";
parameter score_grid_prefix text default "ece";
parameter ligand_number integer range from 1 to 2000 step 1;

```

Molecules to be screened

Figure 9. A plan file defining parameters type and their values.

It is also possible to set concrete values for each of the parameters at runtime when job Generator is invoked. For the parameter 'ligand_number', the user may choose not to select all values from 1 to 2000, but may select a subset of these values. By default, this generated 2000 jobs, each docking a single molecule.

The second part of Nimrod-G plan file is task specification that defines a series of operations that each job needs to perform to dock a molecule (see Figure 10). The 'nodestart' task is performed once for each remote node. Following that, the files copied during that stage are available to each job when it is started. The 'main' task controls the actions performed for each job.

The first line of the 'main' task performs parameter substitution on the file 'dock_base', creating a file 'dock_run'. This is the action that replaces the substitution place markers in our input file with the actual values for the job.

As each docking operation is performed on a selected molecule in the CDB database, it is not necessary to copy such large databases on all Grid nodes. Hence, not only is the molecule file named in the configuration file, we also go to particular lengths to copy only the data for the molecule being tested. The executable script 'get_molecule_fetch' (see Figure 11) is also created using parameter substitution, and runs the 'vlab-cdb-get-molecule' executable, which fetches the molecule record from the CDB molecule server based on the parameter 'ligand_number'. The molecule record is saved in a file whose name is the same as integer value of the 'ligand_number' parameter and 'mol2' as its



```
task nodestart
  copy ./parameter/vdw.defn node:.
  copy ./parameter/chem.defn node:.
  copy ./parameter/chem_score.tbl node:.
  copy ./parameter/flex.defn node:.
  copy ./parameter/flex_drive.tbl node:.
  copy ./dock_inputs/get_molecule node:.
  copy ./dock_inputs/dock_base node:.
endtask
task main
  node:substitute dock_base dock_run
  node:substitute get_molecule get_molecule_fetch
  node:execute sh ./get_molecule_fetch
  node:execute $HOME/bin/dock.$OS -i dock_run -o dock_out
  copy node:dock_out ./results/dock_out.$jobname
  copy node:dock_cnt.mol2 ./results/dock_cnt.mol2.$jobname
  copy node:dock_chm.mol2 ./results/dock_chm.mol2.$jobname
  copy node:dock_nrg.mol2 ./results/dock_nrg.mol2.$jobname
endtask
```

Figure 10. Task definition of docking jobs.

```
#!/bin/sh
$HOME/bin/vlab-cdb-get-molecule.$OS $CDB_SERVER $CDB_PORT_NO ${database_name}.db $ligand_number
```

Figure 11. Parametrization of script for extracting molecule from CDB.

extension. For instance, if the parameter `ligand_number` value is 5, then molecule record will be saved in a file '5.mol2'.

The main code is the 'dock' executable. Note that in the 'execute' command, there are pseudo-parameters that do not appear in the plan file. These include environment variables, such as 'HOME', as well as other useful parameters, such as 'OS' indicating the operating system on the node. This allows us to select the correct executable for the node. If the 'dock' executable files do not exist on Grid nodes, they need to be copied at runtime as part of the job's 'nodestart' task similar to copying input files.

The `dock_run` file created in the substitution step previously is now provided as the input configuration file for the docking process. The output files are then copied back to the local host and renamed with another pseudo-parameter, the unique 'jobname' parameter.

5. SCHEDULING EXPERIMENTATIONS

We have performed scheduling experiments from a grid resource in Australia along with four resources available in Japan and one in U.S.A. Table I shows the list of resources and their properties, Grid



Table I. The WWG testbed resources used in scheduling experiments, job execution and costing.

Organization and location	Vendor, resource type, # CPU, OS, hostname	Grid services and fabric, role	Price (G\$/CPU s)	Number of jobs executed	
				TimeOpt	CostOpt
Monash University, Melbourne, Australia	Sun: Ultra-1, 1 node, bezek.dstc.monash.edu.au	Globus, Nimrod-G, CDB Server, Fork (Master node)	—	—	—
AIST, Tokyo, Japan	Sun: Ultra-4, 4 nodes, Solaris, hpc420.hpcc.jp	Globus, GTS, Fork (Worker node)	1	44	102
AIST, Tokyo, Japan	Sun: Ultra-4, 4 nodes, Solaris, hpc420.hpcc.jp	Globus, GTS, Fork (Worker node)	1	44	102
AIST, Tokyo, Japan	Sun: Ultra-4, 4 nodes, Solaris, hpc420-1.hpcc.jp	Globus, GTS, Fork (Worker node)	2	41	41
AIST, Tokyo, Japan	Sun: Ultra-4, 4 nodes, Solaris, hpc420-2.hpcc.jp	Globus, GTS, Fork (Worker node)	1	42	39
AIST, Tokyo, Japan	Sun: Ultra-2, 2 nodes, Solaris, hpc220-2.hpcc.jp	Globus, GTS, Fork (Worker node)	3	11	4
Argonne National Lab, Chicago, U.S.A.	Sun: Ultra -8, 8 nodes, Solaris, pitcairn.mcs.anl.gov	Globus, GTS, Fork (Worker node)	1	62	14
Total experiment cost (G\$)				17 702	14 277
Time to finish experiment (Min.)				34.00	59.30

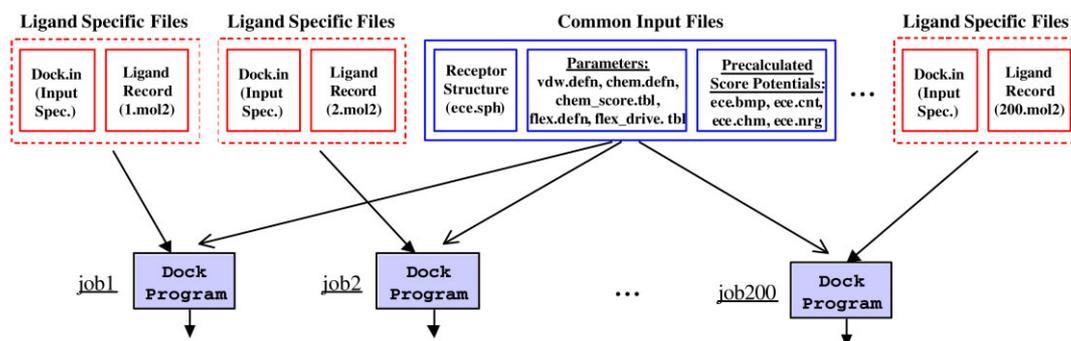


Figure 12. Static and dynamic input files of docking program.

services, access cost or price in terms of Grid dollar (G\$) per CPU-second and the number of jobs processed on resources with deadline-and-budget constrained (DBC) time optimization (TimeOpt) or cost optimization (CostOpt) strategies. The resource price in terms of G\$ is assigned arbitrarily at runtime in these experiments, however, they can be set to match the power of resources and job turn around time as valued in supercomputing centres such as the Manchester computing services [21]. The G\$ can be equated to real money or tokens charged to users for accessing resources. In the current scenario, the users get allocation of tokens via funding from the project sponsoring agents or partnerships. There are supercomputing centres that sell tokens to commercial users and the value of tokens correspond to the quantity of resource allocations. It is also possible to price resources based on the real world economic models [22] that are driven by the supply and demand for resources.

We have performed a trial screening 200 molecules (from the *aldrich_300* CDB) on a target receptor called endothelin converting enzyme (ECE), which is involved in hypotension. The 3D structure of the receptor is derived from homology modelling using related receptor structures whose 3D structures have been solved by x-ray crystallography experiments. In these experiments, for faster evaluation purposes, the range parameter 'ligand_number' is defined with the bounds 1 and 200 and the step size as 1, which produces 200 jobs for docking molecules. As shown in Figure 12, the *dock* program takes two different types of inputs files: (i) *common input files*, the same files are required for all docking jobs and (ii) *ligand specific input files*, which vary from one job to another. The large common input files (receptor structure and pre-calculated score potentials) are pre-staged on resources instead of copying them at runtime. The files are copied using the *globus-rcp* command and stored in the directory location '\$HOME/dock_inputs/' on resources as specified by the parameters 'receptor_site_file' and 'score_grid_prefix' (see Figure 8). The two application-specific executable files, 'dock' and 'vlab-cdb-get-molecule' invoked in the task scripts (see Figure 10 and Figure 11) are also pre-staged. The executable files are stored in the '\$HOME/bin/' directory on resources.

We conducted deadline and budget constrained scheduling experiments for two different optimization strategies [16].

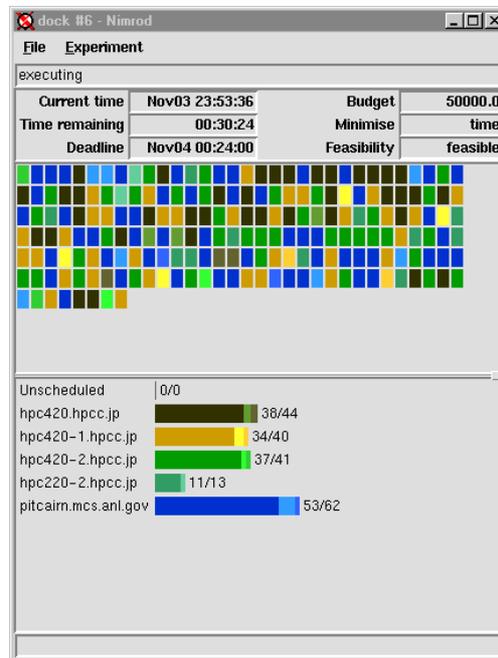


Figure 13. A snapshot of the Nimrod-G monitor during 'Optimize for Time' scheduling.

1. Optimize for time—this strategy aims to produce results at the earliest possible time before a deadline and within a budget limit. It process as many jobs as possible cheapest resources for the deadline period and uses expensive ones just to meet the deadline.
2. Optimize for cost—this strategy aims to minimize the cost of execution (spending from the given budget) and complete the experiment on or before the deadline. It uses all resources aggressively as long as it can afford them and tries to process all jobs at the earliest possible time.

In both experiments, we have set 60 min as the deadline limit and 50 000 G\$ as the budget limit at runtime using the Nimrod-G scheduler steering and control monitor. The value of these constraints can be changed at anytime during the execution, of course not less than the time and budget that is already spent!

The first experiment, *Optimize for time* scheduling, was performed on 3 November 2001 at 23:23:00, Australian Eastern Standard Time (AEST), with a 60 min deadline and finished on 3 November 2001 by 23:57:00. A snapshot of the Nimrod-G monitoring and steering client taken a few minutes (~5 min) before the completion of application processing is shown in Figure 13. This experiment took 34 min to finish the processing of all jobs using resources available at that time with an expense of 17 702 G\$. Figure 15 shows the number of jobs processed on different resources selected depending on their cost

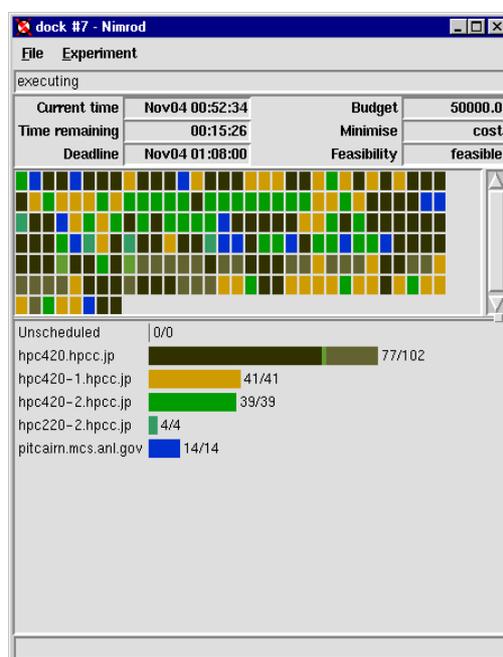


Figure 14. A snapshot of the Nimrod-G monitor during 'Optimize for Cost' scheduling.

and availability. Figure 16 shows the corresponding expenses of processing on resources. Figure 17 shows the number of jobs in execution on resources at different times. From the graphs it can be observed that the broker selected resources to ensure that the experiment was completed at the earliest possible time given the availability of resources and the budget limitations. After 30 min, it discovered that it could still complete early without using the most expensive resource, hpc220-2.hpcc.jp.

It should be noted that for each job scheduled for execution on the Grid, the Nimrod-G runtime machinery (actuator) deploys Nimrod-G agents on remote resources. The Nimrod agents setup runtime environments (generally in scratch file space, 'tmp') on remote resources and execute commands specified in the task definition script (see Figure 10). The docking parameter files and ligand specific files are transferred from the home node, bezek.dstc.monash.edu.au in this case. The agent uses *http* protocols to fetch files via the http-based file server running on the home node. All parameter variables in the parametrized input files (see Figure 8 and Figure 9) are substituted by their concrete values before processing. The ligand record is fetched from the CDB database server running on the home node. The agent then executes the *dock* program and stores output files in the scratch area. The required output files are then transferred to the home node and stored with the job number as their extension. All these steps involved in the execution of the *dock* program on Grid resources were completely hidden from the user.

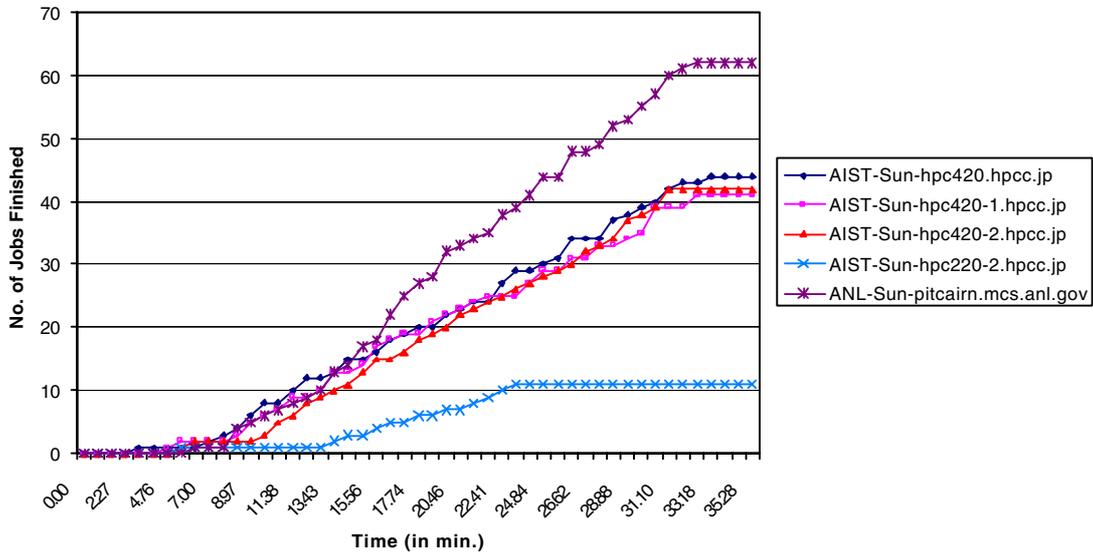


Figure 15. Number of jobs processed on Grid resources during DBC time optimization scheduling.

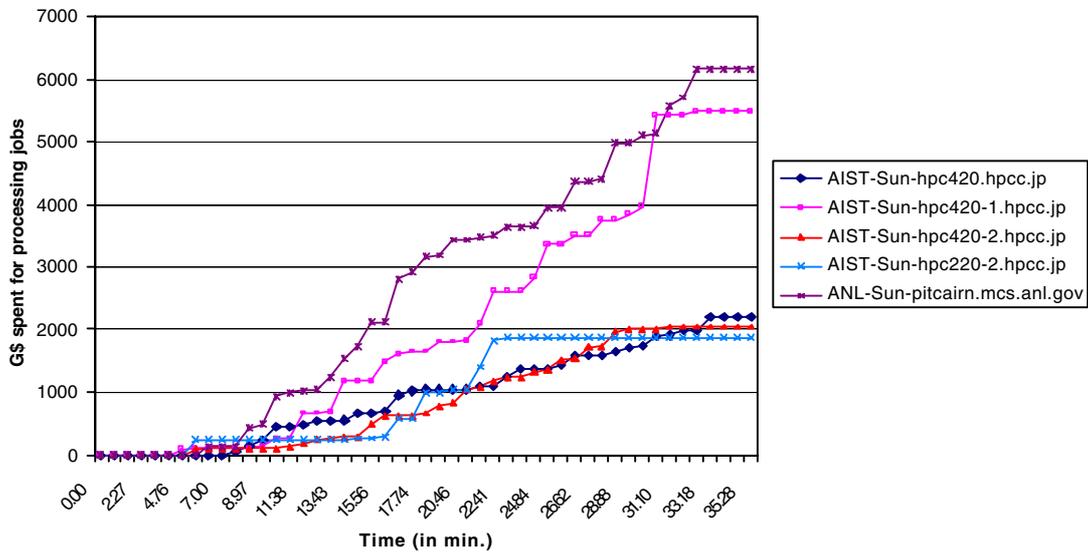


Figure 16. The amount spent on resources during DBC time optimization scheduling.

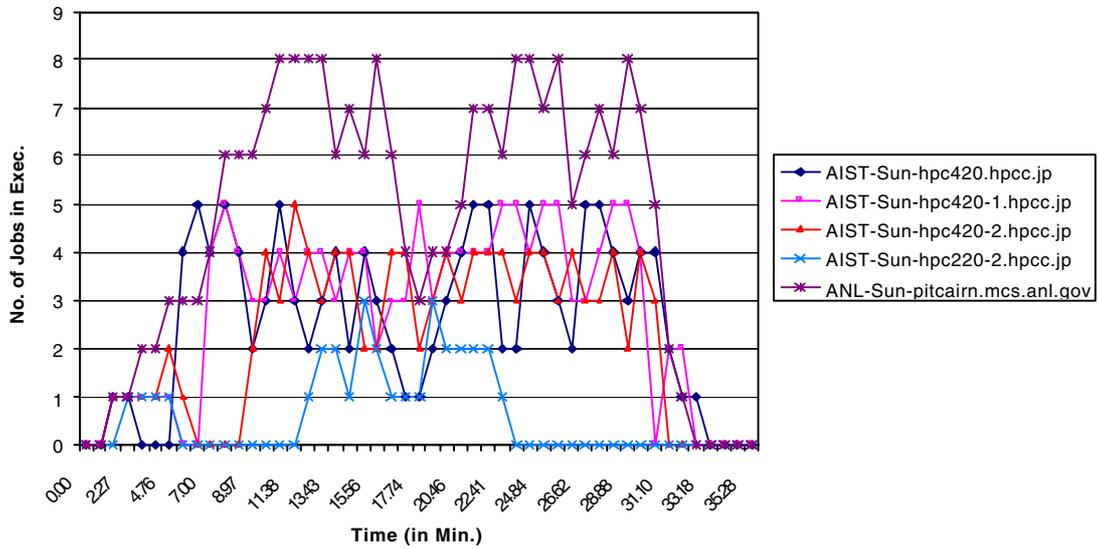


Figure 17. Number of jobs in execution on Grid resources during DBC time optimization scheduling.

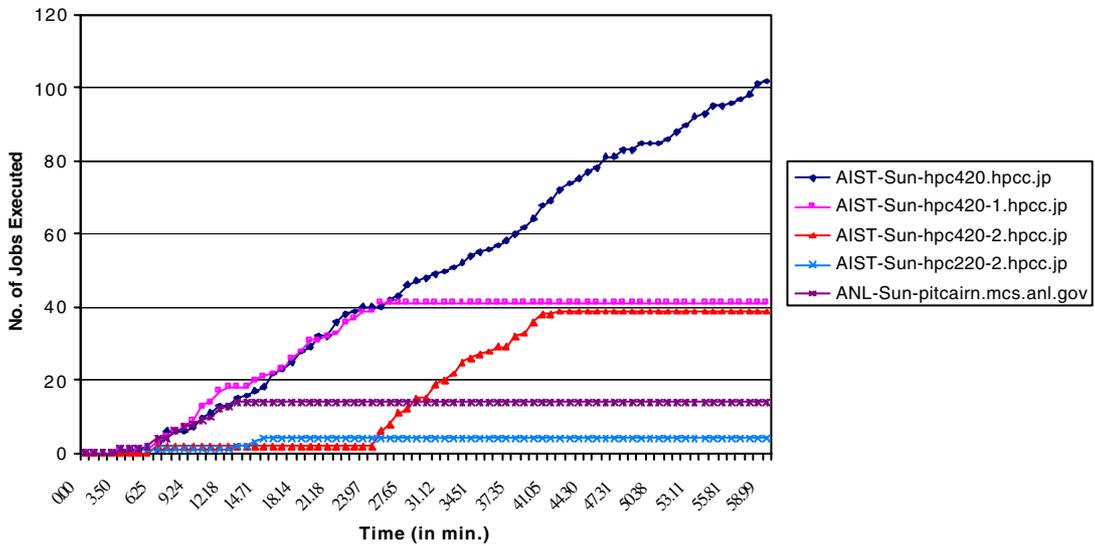


Figure 18. Number of jobs processed on Grid resources during DBC cost optimization scheduling.

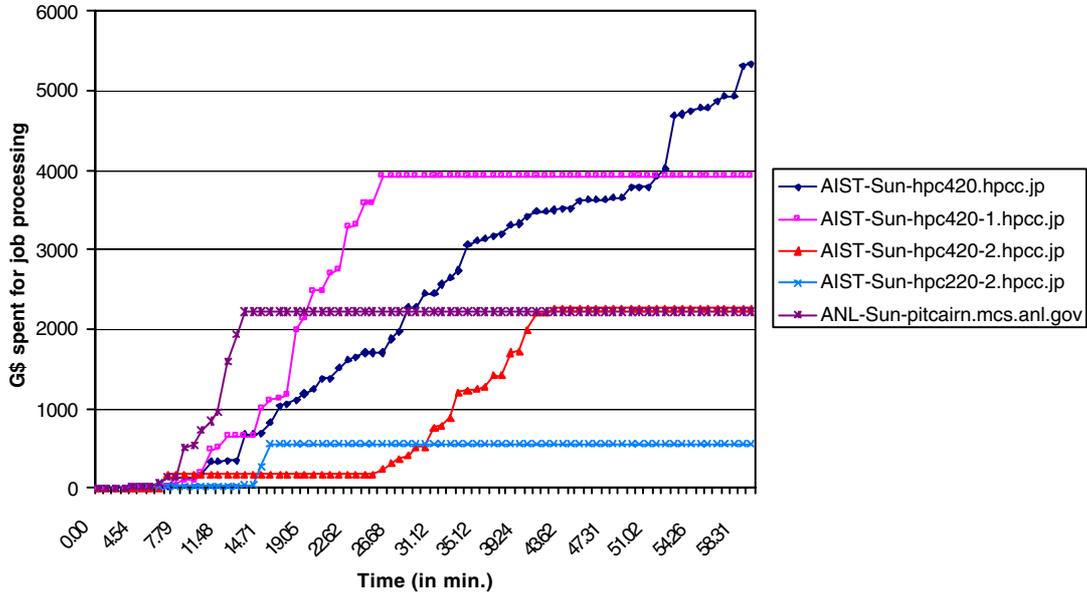


Figure 19. The amount spent on resources during DBC cost optimization scheduling.

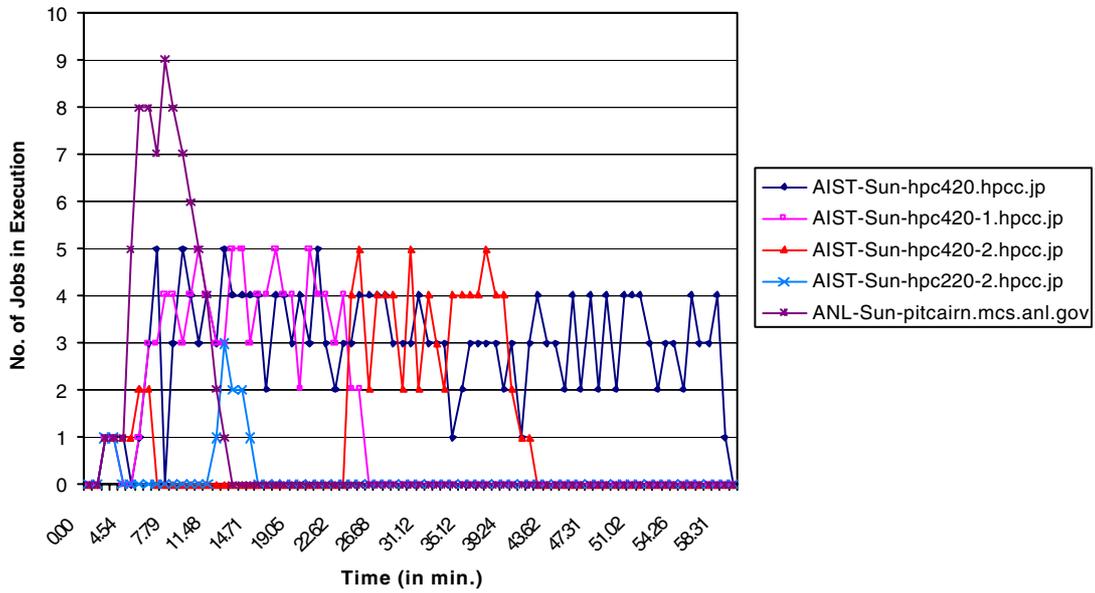


Figure 20. Number of jobs in execution on Grid resources during DBC cost optimization scheduling.



The second experiment, *Optimize for cost* scheduling, was performed on 4 November, 2001 at 00:08:00, AEST, with a 60 min deadline and finished on 4 November, 2001 by 01:07:30. A snapshot of the Nimrod-G monitoring and steering client taken a few minutes (~ 5 min.) before the completion of application processing is shown in Figure 14. This experiment took almost 59.30 min to finish the processing of all jobs using resources available at that time with an expense of 14 277 G\$. It is interesting to note that the second experiment took an extra 25.30 min and saved 3425 G\$ in the process. Figure 18 shows the number of jobs processed on different resources selected depending on their cost and availability. Figure 19 shows the corresponding expenses of processing on resources. Figure 20 shows the number of jobs in execution on resources at different times. From the graphs it can be observed that the broker selected the cheapest resources to ensure that the experiment was completed with minimum expenses, but before the deadline limit. In the beginning expensive resources are used to ensure that the deadline can be met. If for any reason cheapest resources are unable to deliver expected performance, then the broker seeks the help of expensive resources to meet the deadline.

6. RELATED WORK

Although many researchers have explored the use of parallel computing techniques in molecular docking for drug design [23,24], there are only a few efforts that use the world-wide distributed computers for processing docking jobs in parallel. One of the most related efforts is the FightAIDS@Home project [25], which is based on the Entropia's distributed computing network and the Scripps Research Institute's docking application. Here, volunteers need to download Entropia's screen saver program that runs in the background on the volunteer computer. The volunteer PC contacts the Entropia server to download the data to perform docking. When docking on an assigned data is completed, it uploads the results to the server. This execution model is different from our model where the scheduler (Nimrod-G) assigns the work to computers that are available and initiates the execution.

Most of the efforts explicitly develop docking application as a parallel application using a special purpose, legacy or standard, parallel programming languages and interfaces such as PVM and MPI, which requires extra development effort and time. The scalability of such applications and runtime systems is limited to resources available in a single domain and they need powerful computers and networks for faster processing.

Our techniques are novel in many ways. To perform parallel and distributed docking using our tools, there is no need to develop docking application as a parallel application. Our framework supports the composition of the existing molecular docking application as a parameter sweep application without making any changes to it. Our runtime machinery, the Nimrod-G resource broker, creates independent docking jobs automatically and launches their parallel execution on world-wide distributed computers. It hides all the complexity associated with scheduling jobs, shipping appropriated input files, starting and monitoring their execution and shipping results back to the user. Our scheduler also supports the deadline and budget based scheduling, which prioritizes the processing depending on the user requirements—how quickly they need results, how much they want to spend and which one to optimize.

7. SUMMARY AND CONCLUSION

Computational Grids enable the sharing and aggregation of geographically distributed resources for solving large-scale, resource and data-intensive problems faster and cheaper. However, application



development, resource management and scheduling in these environments is a complex undertaking. We have developed a Virtual Laboratory environment and tools for formulating molecular docking for drug design as a parameter sweep application, chemical database management and scheduling docking jobs for processing on a wide area distributed resources by leveraging existing Grid technologies. The new tools developed include a chemical database indexer, CDB server for providing access to molecules in chemical databases as a network service and clients for accessing CDB services from a selected CDB service. We have used the Nimrod-G parameter specification language for composing an existing docking application as a parameter sweep application and the Nimrod-G grid resource broker for processing molecular docking jobs on distributed resources.

We have conducted deadline and budget constrained scheduling experiments for parallel processing of docking jobs on the world-wide grid testbed under two different optimization scenarios. The results of this molecular docking application scheduling on a large-scale distributed resources demonstrate the potential of the Virtual Laboratory tools for service-oriented computing. They prove the effectiveness of computational economy and quality of services (QoS) driven scheduling as an efficient mechanism for the management of supply-and-demand for resources depending on the value delivered to the user. The economy driven service oriented computing encourages the users to utilize resources effectively by trading off between the deadline and budget depending on their QoS requirements.

Our experience with developing a prototype Virtual Laboratory environment for distributed drug design shows the potential and applicability of the Nimrod-G tools for data intensive computing. We are extending the current system to support adaptive mechanisms for the selection of the best CDB service depending on access speed and cost. We are also looking into applying the experience gained in this work to develop a virtual laboratory environment for enabling high-energy physics events processing on distributed resources on a larger scale.

ACKNOWLEDGEMENTS

We would like to thank John Crossley (Monash University) and Rob Gray (DSTC) for proofreading earlier drafts. They offered constructive comments and suggestions for enhancing the quality of the paper. We appreciate the efforts of Heinz Stockinger (CERN) for reviewing the chemical database server implementation. We thank the National Institute of Advanced Industrial Science and Technology, Tokyo, Japan and the Argonne National Laboratory, Chicago, USA for providing access to the resources used in the experimentation. We thank Geoffrey Fox for handling the reviews in record time. We express our special thanks to all six anonymous reviewers for their critical comments and excellent advice on improving the paper.

REFERENCES

1. Foster I, Kesselman C (eds.). *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1999.
2. Lunney E. Computing in drug discovery: The design phase. *IEEE Computing in Science and Engineering Magazine*. <http://computer.org/cise/homepage/2001/05Ind/05ind.htm>.
3. Buyya R (ed.). *High Performance Cluster Computing: Architectures and Systems*, vols. 1 and 2. Prentice-Hall: Englewood Cliffs, NJ, 1999.
4. Sullivan W, Werthimer D, Bowyer S, Cobb J, Gedye D, Anderson D. A new major SETI project based on Project Serendip data and 100 000 personal computers. *Astronomical and Biochemical Origins and the Search for Life in the Universe—Proceedings of the Fifth International Conference on Bioastronomy*, 1997.
5. Buyya R. World Wide Grid (WWG). <http://www.csse.monash.edu.au/~raj कुमार/ecogrid/wwg/>.



6. Abramson D, Giddy J, Kotler L. High performance parametric modelling with Nimrod/G: Killer application for the Global Grid? *Proceedings International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE Computer Society Press: Los Alamitos, CA, 2000.
7. Buyya R, Abramson D, Giddy J. Nimrod/G: An architecture for a resource management and scheduling system in a global computational grid. *Proceedings of the 4th International Conference on High Performance Computing in Asia-Pacific Region (HPC Asia 2000)*, China, 2000.
8. Abramson D, Susic R, Giddy J, Hall B. Nimrod: A tool for performing parameterized simulations using distributed workstations. *Proceedings 4th IEEE Symposium on High Performance Distributed Computing*, Virginia, August 1995.
9. Shoichet B, Bodian D, Kuntz I. Molecular docking using shape descriptors. *Journal of Computational Chemistry* 1992; **13**(3):380–397.
10. TurboLinux. EnFuzion Manual. <http://www.turbolinux.com/downloads/enf/man/enfuzion.htm>.
11. Buyya R, Abramson D, Giddy J. Economy driven resource management architecture for computational power grids. *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 2000)*, U.S.A., 2000.
12. Foster I, Kesselman C. Globus: A metacomputing infrastructure toolkit. *International Journal of Supercomputer Applications* 1997; **11**(2):115–128.
13. Ewing A (ed.). DOCK Version 4.0 Reference Manual. University of California at San Francisco (UCSF), U.S.A., 1998. <http://www.cmpchem.ucsf.edu/kuntz/dock.html>.
14. Kuntz I, Blaney J, Oatley S, Langridge R, Ferrin T. A geometric approach to macromolecule-ligand interactions. *Journal of Molecular Biology* 1982; **161**:269–288.
15. Meng E, Gschwend D, Blaney J, Kuntz I. Orientational sampling and rigid-body minimization in molecular docking. *Proteins* 1993; **17**(3):266–278.
16. Buyya R, Giddy J, Abramson D. An evaluation of economy-based resource trading and scheduling on computational power grids for parameter sweep applications. *Proceedings of the Second Workshop on Active Middleware Services (AMS 2000)*, In conjunction with HPDC 2001, Pittsburgh, U.S.A., 1 August 2000. Kluwer Academic Press, 2000.
17. Buyya R. Economic-based distributed resource management and scheduling for Grid computing. *PhD Thesis*, Monash University, Melbourne, Australia, 12 April 2002. <http://www.buyya.com/thesis/thesis.pdf>.
18. Tripos, Inc.. SYBYL Mol2 File Format, U.S.A. 7 October 2001. <http://www.tripos.com/services/mol2/>.
19. Tripos, Inc. SYBYL Toolkit for Molecular Design and Analysis, U.S.A. 7 October 2001. <http://www.tripos.com/software/sybyl.html>.
20. Buyya R. The Virtual Lab Tools for Chemical Database (CDB) Management and Access. <http://www.buyya.com/vlab/>.
21. Brooke J, Foster M, Pickles S, Taylor K, Hewitt T. Mini-Grids: Effective test-beds for Grid application. *Proceedings of the First IEEE/ACM International Workshop on Grid Computing (GRID 2000)*, Bangalore, India, 17 December 2000. Springer: Berlin, 2000.
22. Buyya R, Abramson D, Giddy J, Stockinger H. Economic models for resource management and scheduling in Grid computing. *Concurrency and Computation: Practice and Experience* 2002; **15**(13–15):1507–1542.
23. Morris G, Goodsell D, Huey R, Olson A. Distributed automated docking of flexible ligands to proteins: Parallel applications of AutoDock 2.4. *Journal of Computer-Aided Molecular Design* 1996; **10**:293–304.
24. Schneider R. Drug Design Using Parallel Processor. http://www.cordis.lu/esprit/src/results/res_area/hpnc/hpnc15.htm.
25. Entropia. Fight AIDS at Home Project: A joint effort of Entropia and Scripps Research Institute. <http://www.fightAIDSatHome.org/>.