



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/1819/>

Book Section:

Pham, T.V., Lau, L.M.S., Dew, P.M. et al. (2005) A collaborative e-Science architecture towards a virtual research environment. In: Cox, S.J. and Walker, D.W., (eds.) Challenges of Large Applications in Distributed Environments, 2005. CLADE 2005. Proceedings. EPSRC, pp. 13-22. ISBN: ISBN 1-904425-53-4.

Reuse

See Attached

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

A Collaborative e-Science Architecture towards a Virtual Research Environment

Tran Vu Pham^α, Lydia MS Lau^α, Peter M Dew^β and Michael J Pilling^ζ

^α*School of Computing, University of Leeds, Leeds LS2 9JT*

^β*Informatics Institute, School of Computing, University of Leeds, Leeds LS2 9JT*

^ζ*School of Chemistry, University of Leeds, Leeds LS2 9JT*

Abstract

This paper presents a novel Collaborative e-Science Architecture (CeSA) to address two challenging issues in e-Science that have arisen from the management of heterogeneous distributed environments. By combining the capabilities of peer-to-peer and Grid computing, the architecture provides an environment for scientific collaborations within distributed, loosely coupled research communities and brings computation and data intensive resources to the desktops of the scientists in these communities. The Reaction Kinetics research community had been used as a case study to capture realistic requirements. A prototype based on the architecture was developed for user experiment and evaluation. The results of these experiments were promising. It has provided further motivation to evolve CeSA towards a Virtual Research Environment.

1. Introduction

This paper addresses two challenging issues of an e-Science community: firstly, how to satisfy the need for easy access to large-scale computational and data resources, and secondly, how to provide a collaborative research environment to support research activities from heterogeneous distributed communities. Such requirements were noted during a study on the Reaction Kinetics Research Community. Scientists from this community strongly needed access to computation and data intensive resources for their time consuming simulations on chemical reaction mechanisms. They also required support for exchanging different types of datasets directly with other scientists in the field.

Clearly, a Virtual Research Environment needs to satisfy all the above requirements. In the "Roadmap for a UK Virtual Research Environment" produced for the JISC in 2004 [1], grids were identified as the most promising infrastructure to provide computationally extensive resources and storages for large-scale datasets. A Service Oriented Architecture was proposed for the virtual research environment, which links the resources on grids to the scientists via web-based portals.

Using the web-based portals, the scientists can virtually logon the virtual environment from anywhere with a browser and an Internet connection. However, with the portal approach, scientists of a research community are required to pre-register at a centralised administration of the

environment in order to utilise the available resources or collaborate with other scientists. This centralised management does not seem to fit well with the decentralised, or loosely coupled, nature of research activities, where collaborations across different 'communities' can also happen (e.g. a scientist may want to share a dataset with some other scientists from another related research discipline, but these scientists are not members of the same research community covered by the portal).

Peer-to-peer, on the other hand, is a decentralised computing model. It has been very successful in many commercial file-sharing applications at the user end, such as Napster¹ and currently Kazza² or eMule³. This computing model has the potential to bridge the gap between the current centralised approach for the collaborative environments and the loosely coupled research activities amongst the scientists.

In this paper, a novel Collaborative e-Science Architecture (CeSA) is proposed. The aim is to bring together services from the grids beyond boundaries of web-based portals into a collaborative peer-to-peer environment, leveraging advantages of both Grid and peer-to-peer technologies. The architecture focuses on the support for general collaborative activities for scientists as well as for the sharing of computational capability and storages (e.g. for simulations and analyses) required by the

¹ <http://www.napster.com>

² <http://www.kazza.com>

³ eMule project: <http://www.emule-project.net>

community. This paper also reports on the requirements of the Reaction Kinetics research community. A prototype was developed to evaluate the architecture in a realistic environment. The result has shown the feasibility of this approach and the potential benefits to the scientists. Some related projects and future work will also be discussed at the end of the paper.

2. The Collaborative e-Science Architecture (CeSA)

The architecture described in this paper had been extended from an earlier version of the CeSA as presented at CLADE 2005 [2]. The new version was driven by a deeper understanding of the work practices in the user community involved in the prototyping and new features had been added to bring it a step closer towards an architecture for a virtual research environment.

2.1 Architecture Overview

The CeSA is composed of a peer-to-peer collaborative environment on top of a collection of grids (Fig. 1). The management within the grids is often complex and requires tight security control, whereas, the peer-to-peer environment is in a public domain and self-managed by individual users. The design of this architecture is aimed at separating the complex management within the grids from day-to-day collaborations in the peer-to-peer environment. This is made possible by adopting a flexible service oriented architecture.

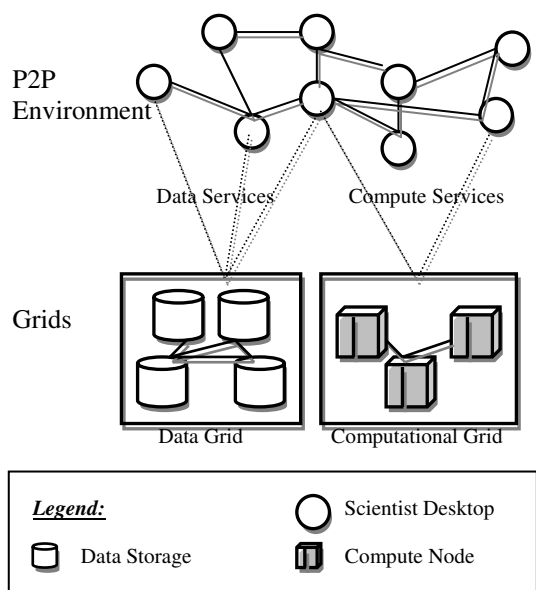


Fig. 1: Physical view of the Collaborative e-Science Architecture

In the peer-to-peer environment, each scientist has a copy of the desktop peer-to-peer application. This desktop application is designed as a window to an integrated research environment, from which a scientist can make queries and get access to computational resources and data provided by the grids. The desktop application also allows the scientists to collaborate with each other in a peer-to-peer fashion, such as forming common interest groups, direct sharing of working data and information available on grids and so on.

In the architecture, the grids are main providers of computational intensive resources and research datasets for the research community in the peer-to-peer environment.

2.2 Grids

The Grid environment can have more than one computational or data grids acting as resource providers for the scientist community in a peer-to-peer environment.

Each of these grids has a service container, called High Level Service Container (Fig. 2), which provides high level services to the scientist community. These high level services can be built by wrapping resources on the grid or they can be 'composite services' composed of other services or workflows available on the grids. The difference between these high level services and the basic Grid services, such as Grid job submission services, is that the former aim to deliver final output to the users, such as results from a simulation on chemical data or a chemical dataset, whereas the latter usually deliver intermediate output to another computer application or service.

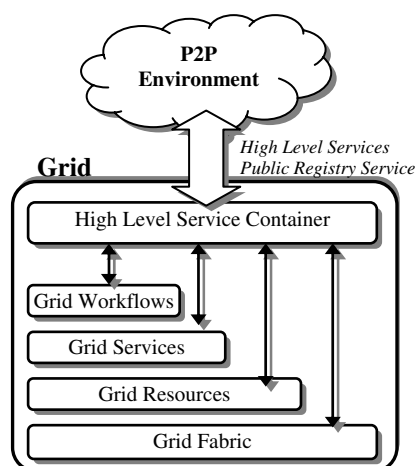


Fig.2: Relationships between a High Level Service Container and other grid components and with peer-to-peer environment

In addition to the high level services, a High Level Service Container also contains a Public Registry Service, which allows the peer users (i.e. scientists) to make queries on services available on the grid. The information about the services can then be published within the peer-to-peer environment so that other peer users can locate and use these services on the grids.

2.3 Peer-to-Peer Environment

The peer-to-peer environment is composed of a network of peer-to-peer applications. Each of the peer-to-peer applications (Fig. 3) consists of a User Interface, a Service Client, a Service Publication and Discovery Agent, a set of Community Services and a set of Peer Core Components.

User Interface is the main gateway for a user to access the functionalities of the peer-to-peer application.

The Service Client interfaces with high level services and registry services provided by the grids. It allows a user to browse, via the user interface, services available in the grid environment.

The Service Publication and Discovery Agent has a very important role in the collaborative architecture. It provides two basic functions: publishing information about services available on grids to the peer-to-peer environment and discovering information about services previously published by other peers. Information about a service published in peer-to-peer environment consists of information such as service name, service URI (Universal Resource Indicator, the identifier that identifies a particular service from a service container), its required input data and its output data. The published information also includes information about the grid that provides the service. The structure of a discovery query is similar to the structure of published information about services. A matching algorithm between the published information and the discovery query will be used during a service discovery process. Because of its important role, the Service Publication and Discovery mechanism needs to be as effective as possible.

Community Services consist of components to support the day-to-day collaborations within a community. Examples include components for file sharing, community/group formation and instant messaging. Through these components, a scientist can set up ad hoc interest groups or a research community. The scientist can then establish sharing resources directly with other scientists in his/her working group or community.




Fig. 3: Components of a Peer-to-peer Application

Peer Core Components include components for communication with other peers, peer identification and peer resources discovery. Service Publication and Discovery Agent and Community Services rely on these core components to communicate with other peers in a peer-to-peer environment.

2.4 Specification of High Level Services for Peer-to-Peer Applications

The design of the CeSA aims to make it independent of implementation technologies. Therefore, the concept of 'high level service' (or sometimes referred as 'service' in this paper) in the CeSA is not coupled with any particular type of services. It can be implemented as a Web Service or a Grid Service.

As discussed earlier in section 2.2, high-level services provided by the grids enable peer users to perform tasks such as simulation or analysis. However, to use such services is often rather challenging for the users. It is necessary to have a simple user interface for users to easily execute services from grids on their desktops. For example, it is not feasible to give them a service handler (e.g. an identification of a Grid service) and ask them to build a service client application to execute the service. This task is difficult even for programmers who have no knowledge about Web Services or Grid Services.

Fortunately, there are commonly two main categories of services in a scientific environment: computational services and data services. The first category is often associated with simulations and analyses used in scientific communities. These services share common characteristics such as using files as input and output. Hence, a simple unified service interface can be developed which consists of the following operations:

- *List input required*: allows service client to query in advance the input files required.
- *Load input*: to upload input files required to the service server prior to execution.
- *Execute*: to run the services after all required input are uploaded.
- *List output*: to query the number of output files produced by the service.
- *Transfer output*: to send the output to the user.

The second category is about discovery and delivery of research data held in the data grids. At a high level, a unified discovery service interface can be developed for all data discovery services. Similarly, a unified service interface can also be built for all data delivery services.

Once all the services conform to a standard set of interfaces service clients can be developed in a way that they can interact with virtually any service.

3. Case Study: Reaction Kinetics Research Community

Having established a generic Collaborative e-Science Architecture, a case study was conducted to test the architecture. The Reaction Kinetics research community was chosen and a prototype was built. Reaction Kinetics researchers study the elementary and complex chemical reactions and their applications. This is a multidisciplinary research subject, spanning across Combustion, Atmospheric and Environmental Studies. The main activities are shown in Fig. 4.

3.1 Requirements of Reaction Kinetics Research Community

A central object in an applied reaction kinetics problem is the chemical reaction mechanisms. Different research activities are taking place on different aspects of reaction mechanisms. A research group in the community may undertake one or more of these activities. The allocation of activities in the groups often depends on the expertise and the resources at the institutions where the groups are based.

As research activities in Reaction Kinetics are inter-related and tied together by the common interests in chemical reaction mechanisms, as depicted in Fig. 4, research groups that undertake these activities also need to collaborate with each other.

There is a very close working relationship amongst the research groups in Reaction Kinetics research. Research in one area of Reaction Kinetics depends on data produced by several

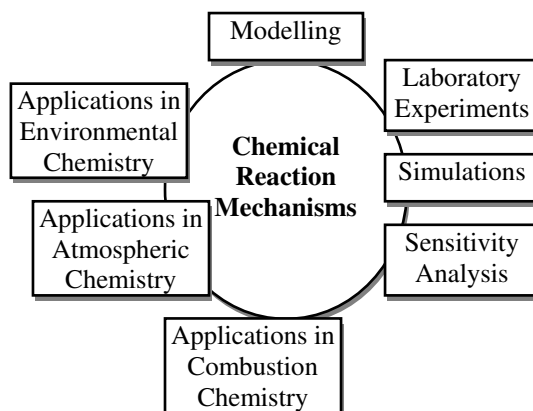


Fig. 4: An Example of Research Activities in Reaction Kinetics Research Community

related areas. Data exchange within the Reaction Kinetics community is an iterative process. A delay caused by one individual or one research group might affect several other people and groups and eventually the whole community. Collaboration amongst individual scientists as well as research groups in Reaction Kinetics from many locations worldwide is very important. A range of computational intensive resources is required for running simulations and analysis. Storages of huge volume of experimental and field data also exist all over the world.

Through the collaboration with chemists in the Reaction Kinetics research group at the School of Chemistry, the University of Leeds, the following requirements for a collaborative research environment were identified:

- Allow scientists who are working on the same or similar research activities to dynamically form working groups.
- Provide efficient support for timely collaborations within and across working groups in the community for sharing expert knowledge, day-to-day working data, such as experimental data, chemical reaction mechanisms and related input data for reaction modelling.
- Provide easy access to computational intensive resources for time and resource consuming simulations and analyses and for archiving of large amount of experimental data.

3.2 A Potential Application of the CeSA

This section will describe how the CeSA can be employed to support a collaborative research environment for the Reaction Kinetics research community.

The requirements (i) and (ii), as identified in section 3.1, can be met by the peer-to-peer collaborative environment of the architecture. Using the peer-to-peer applications, with functionalities described in section 2.3, chemists in Reaction Kinetics community will be able to form ad hoc interest groups. Once the relationship is established, chemists will be able to share their working documents or chemical datasets (e.g. a reaction mechanism under review) with other chemists in a peer-to-peer fashion. If a web-based community were used, all participants will have to register with a central website before any community activities can take place. Centralised rules would need to be followed for uploading or downloading. However, in a peer-to-peer environment, that relationship can be formed straight away by participants, without bothering any third party. In terms of sharing working documents or datasets, sharing in a peer-to-peer environment is more direct and convenient than via a centrally controlled website.

With regard to requirement (iii), access to computational resources for simulations of reaction mechanisms can be done through high-level services provided by the computational grids. Archives of stable versions of chemical reactions mechanisms and other type of data in Reaction Kinetics such as reaction rate coefficients can also be stored on the data grids via data services. Grids are well known with the ability to support computationally intensive calculations and storages for large-scale datasets.

4. Prototype

The purpose of prototyping was to provide an insight into the technical challenges as well as testing the applicability of the CeSA for the user community. In this version of the prototype, only services dealing with computational resources were integrated to peer-to-peer environment.

The high level services developed for the prototype were Grid Services and the High Level Service Container was based on Grid Service Container of OGSA [2].

The peer-to-peer prototype, based on JXTA [4], consisted of generic tools for collaborations, such as file sharing, chatting, group formation, publishing and discovering of information about high level services available on grids. It also provided a graphical interface for the chemists to execute computational services.

4.1 Conversion of Existing Applications into Grid Services

A few applications used in the community for simulations and analyses of chemical reaction mechanisms were wrapped into Grid services. These applications use files as input and output, and could also produce console output. When wrapping these programs into Grid services using Java, input and output (including console output) were mapped to the input and output parameters of Grid services. These Grid services conformed to the unified service interface as specified in section 2.4.

After being wrapped into Grid services using Java, these new services were deployed in a Grid Service Container provided in Globus Toolkit version 3.0.2 (GT3). This Grid Service Container ran on one machine, played the role of a computational grid providing simulation and analysis services for the peer-to-peer community.

4.2 CeSA Prototype System

The CeSA prototype system is basically a peer-to-peer application for the Reaction Kinetics Research Community. The following explain briefly how the main components were built using JXTA and how the GT3 services were integrated to the application.

Service Client: The Service Client consisted of two parts: a service browser and a service executor.

The service browser interacts with GT3 Container Registry Service to retrieve a list of service handlers of the services provided by the container. The result will be displayed to users via the User Interface.

The service executor could invoke a service from a service handler. Globus Toolkit 3.0.2 had two types of service handlers for a Grid service. One type of handlers was for the Factory Service and the other type was for the Grid service instance. The service executor can be used to generate new instances of a service from a factory service handler or to execute the service using the handler of a service instance. In this prototype, the service executor could interact with any of the services, as described in section 4.1, for chemical reaction simulations as they all conformed to the unified service interface.

Service Publication and Discovery Agent: The method used for publication and discovery of service information was based on JXTA protocols [5].

All information about a service, such as service name, service provider, input, output, and so on, was enclosed in a JXTA advertisement. The advertisement about the service was then published in JXTA peer-to-peer network using JXTA discovery protocol.

The discovery of information about services, however, was not based on JXTA discovery, but using JXTA resolver protocol, because the default discovery mechanism provided with JXTA discovery protocol [6] was not flexible enough to deal with complex query requirements. In the prototype, users are able to query using any information about the services or a combination of them. If the JXTA default discovery protocol were used, only a few indexed attributes could be searched. With JXTA resolver protocol, a query could be distributed to the other peers in the environment. On receiving the query, a peer could then search through its cache for service advertisements that matched the criterion specified in the query, such as service name, service provider, etc. The results would be sent back directly to the query issuer.

Community Services: Community Services consisted of a number of components for collaboration within a peer-to-peer environment.

The function for the formation of working groups or communities was developed upon JXTA Peer and Peer Group concepts. An individual peer user was mapped to a peer in the JXTA peer-to-peer network. Similarly, a working group or community is corresponding to a JXTA Peer Group.

Tools for file sharing and instant messaging were built using JXTA pipe. To send a message or a file from one peer to another peer, a pipe between the two peers would be set up first. Then, the message or file would be sent over the pipe. Each peer had a unique pipe advertisement, which was the necessary information for setting up a pipe connection. This information would be published to the peer-to-peer environment when the peer entered the network.

There were also additional components for managing the 'share relationships' amongst the peers and for the searching of shared resources. Shared resources available on a peer, such as a working data file, were more dynamically managed than a shared service. A file might be set to share to a group only at a particular time. Therefore, the approach to resource discovery, more specifically the file search, was different from the method used for service publication and discovery. There was no publication of these resources. As only the owner of the resources could say whether he/she had resources being

shared for the query issuer, the query message had to be distributed to every potential resource owner. Hence, the prototype used broadcasting method to distribute query messages. Scope of a query could also be limited to particular working groups of interest, where the answers most likely could be found.

5. Experiment and Evaluation

5.1 Objectives

The main aim of this experiment was to collect feedback from a number of researchers from the Reaction Kinetics research community on the functionalities of the prototype based on the CeSA. User comments were invited to compare this new way of working with their current practice. More specifically, the objectives of this experiment were:

- To evaluate the effect of using peer-to-peer collaborative functionalities provided by CeSA prototype system for collaborations within the user community.
- To assess how users can benefit from the access to remote simulations and analyses in Reaction Kinetics using remote services.
- To capture user's general attitude to the new collaborative infrastructure.

5.2 The Experiment Process

The experiment was conducted at a Reaction Kinetics research laboratory at The University of Leeds. There were three chemists from Reaction Kinetics research community participating in the initial experiment.

A few copies of CeSA prototype system were installed in the laboratory for the experiment. During the experiment, participants used the prototype to carry out a sample of day-to-day collaborative activities and to run simulations using Grid Services, which were provided by a grid service container located remotely.

A questionnaire was issued to each participant to record his/her feedback on the CeSA prototype system. The questionnaire consisted of a mixture of open and closed questions and was divided into sections which corresponded to the experiment objectives. To assess the effect of using the prototype, there were questions about the comparison between users' current way of working and the method facilitated by the functionalities of the prototype. Any potential benefit of using the prototype was also recorded. General feedback and impression on the CeSA prototype system was collected at the end of the

questionnaire. There were discussions during the experiment to clarify the questions in the questionnaire.

5.3 Results and Analysis

The following is a summary of the findings in relation to the objectives of this evaluation.

Firstly, on the use of peer-to-peer collaborative functionalities, all three participants had a need for sharing working data with their colleagues. All participants recognised the benefits of using the CeSA prototype system for sharing working data because of the following reasons:

- A file can be shared directly from the user machine. Therefore, there is no need to move a file around for sharing.
- Spontaneous sharing of file-in-progress is possible. This allows other users to be able to copy the latest version of the file.
- Users have control over shared data. They can choose to share with a group or a specific person. They can also easily revoke a file from sharing.

The participant recognised the advantages of using peer-to-peer for file sharing, over their current way, such via email or a common website, especially when the shared files were being updated frequently.

One participant expressed his concern about security of peer-to-peer environment. When the access controls are decentralised to individual machines, it will be important to provide a security mechanism to protect user's own computing resources and personal data.

There was also a worry about the stability of the peer-to-peer environment. As JXTA is a decentralised peer-to-peer environment, it takes some time for the community on JXTA network to reach a stable condition. This is actually a trade-off for the dynamics and flexibility of peer-to-peer computing.

A suggestion on documenting and tracking changes was made, as there are times when different users might be working collaboratively on the same set of data.

Secondly, on the use of Grid Services for accessing remote simulations and analyses, the participants recognised the benefit of freeing their desktops for other work, especially for jobs that required long execution time. They all agreed that it would be very useful for them to run their simulations in this way.

They also agreed that when using remote services in the way that provided in this prototype, they could search and use the service straight

away. They could also easily share information about useful services to their colleagues. This way was more convenient than sharing traditional software.

Finally, on the general feedback, the functionalities of the CeSA prototype system had made a good impression on all participants. The participating chemists believed that their research community would benefit from a system like this. Although there were concerns about security issues and the stability of peer-to-peer network, they would be happy to use it if the system was fully implemented and widely accepted by their community. This "acceptance by the community" is a vital condition that any peer-to-peer system has to satisfy. A user-friendly graphical interface was also identified as an important characteristic to encourage the users to use the system.

6. Related Work

There are a number of projects related to this work due to the similarity either in the problem domain or in the architectural issues.

The NERC DataGrid project [7][8][9] has been working on metadata infrastructure to facilitate the discovery and delivery of environmental datasets, which are held in a loosely coupled federation of distributed locations. The main user interface of NERC DataGrid was planned to be web-based portal. The NERC DataGrid can be plugged into the CeSA as one of data grids, providing data services for the user community.

In e-Science programme, myGrid [10][11] and Smart Tea [12], a sub project of Comb-e-chem [13], also aimed at supporting scientists directly. myGrid project focused on the middleware components and services that enabled collaborations on Web Services and composition of these services into workflows. The goal of Smart Tea was to produce a paperless eLab, which enabled publication of experimental data at source using pervasive computing devices. Smart Tea is being extended to interact with myGrid to support researchers in bio-informatics in the myTea project [14]. Using the peer-to-peer approach of the CeSA can enhance the support to individual researchers in some of their collaborative activities.

On the architectural issue, projects such as Chinook [15] and Triana[16] also used Service Oriented Architecture and peer-to-peer computing model. The goal of Chinook was to convert command-line applications in bio-informatics into web services and publish them into a virtual peer-to-peer environment. In Triana, each node in the peer-to-peer environment provides services to other nodes [17]. The difference between the

CeSA and these projects is that CeSA emphasises strongly on supporting direct user collaborations within scientific community.

7. Conclusions and Future Work

This paper has introduced the CeSA as a potential architecture for a collaborative research environment. The case study from Reaction Kinetics research community was used to illustrate the applicability of the CeSA in a realistic environment. The experiment and evaluation on the CeSA prototype conducted by a group of Reaction Kinetics researchers have also been reported. The notable results were that the ability to provide an environment for direct user collaborations and access to remote computational resources of the CeSA was recognised as an advantage. This initially confirmed the correctness of the approach. As the Reaction Kinetics research community is typical of e-Science communities, the architecture therefore has the potential of wider deployment.

The CeSA is being considered to integrate the portal approach within the peer-to-peer application. The aim is to allow 'plug-in remote portlets' from other portals. With this ability, the CeSA will be able to work inter-operably with web-based portal approach.

On the implementation side, prototype of data services for delivering datasets from data grids to peer-to-peer environment is being developed. These data services will later be used in further experiments.

Acknowledgement

Many thanks to Dr. Andrew Rickard, Dr. Lisa Whalley, Jenny Stanton and other researchers working in Reaction Kinetics group at School of Chemistry, University of Leeds, who have been offering us great help on requirement collection, experiments and evaluation of this work.

References

- [1] Allan, R. and et al., Roadmap for a UK Virtual Research Environment, http://www.jisc.ac.uk/uploaded_documents/VRE%20roadmap%20v4.pdf (last accessed June, 2005).
- [2] Pham, T.V., Lau, L.M.S., Dew, P.M and Pilling, M.J., Collaborative e-Science Architecture for Reaction Kinetics Research Community, in the *Proceedings of the Challenges of Large Applications in Distributed Environments 2005 (CLADE2005)*, Research Triangle Park, NC, USA, July 2005.
- [3] Globus Toolkit, <http://www-unix.globus.org/toolkit/>, (2005).
- [4] Project JXTA: <http://jxta.org>, (2005).
- [5] Project JXTA: Java Programmer's guide, http://www.jxta.org/docs/JxtaProgGuide_v2.3.pdf, (2005).
- [6] Traversat, B., Abdelaziz, M., and Pouyoul, E., "A Loosely-Consistent DHT Rendezvous Walker", Sun Microsystems Inc., March 2003.
- [7] NERC DataGrid, <http://ndg.badc.rl.ac.uk/> (2005).
- [8] Woolf A., et al., "Enterprise Specification of the NERC DataGrid", *Proceedings of the U.K. e-science All Hands Meeting*, 2004. S.J.Cox(Ed).
- [9] Lawrence, B., et al., "Googling secure data", *Proceedings of the U.K. e-science All Hands Meeting*, 2004. S.J.Cox(Ed).
- [10] myGrid, <http://www.mygrid.org.uk>, (2005).
- [11] Goble, C.A., Pettifer, S., Stevens, R., and Greenhalgh, C., "Knowledge Integration: In silico Experiments in Bioinformatics", *The Grid: Blueprint for a New Computing Infrastructure*, Second Edition eds. Ian Foster and Carl Kesselman, 2003, Morgan Kaufman, November 2003.
- [12] The Smart Tea Project, <http://www.smarttea.org/> (2005).
- [13] Combechem, <http://www.combechem.org/>, (2005).
- [14] the myTea project, <http://mytea.org.uk/> (2005).
- [15] GSC-Chinook, <http://www.bcgsc.bc.ca/chinook/> (2005).
- [16] The Triana Project, <http://www.trianacode.org/index.html> (2005).
- [17] Wang, I., Shields, M., Taylor, I. and Philp, R. Distributed P2P Computing within Triana: A Galaxy Visualization Test Case, *Proceedings of 17th IPDPS 2003*, IEEE Computer Society, 2003.