

This article was downloaded by:

On: 14 January 2011

Access details: *Access Details: Free Access*

Publisher *Taylor & Francis*

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



## **Molecular Simulation**

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713644482>

### **Workflow issues in atomistic simulations**

C. Chapman<sup>a</sup>; J. Wakelin<sup>b</sup>; E. Artacho<sup>b</sup>; M. T. Dove<sup>b</sup>; M. Calleja; R. Bruin<sup>b</sup>; W. Emmerich<sup>a</sup>

<sup>a</sup> Department of Computer Science, University College London, London, UK <sup>b</sup> Department of Earth Sciences, Cambridge, UK

**To cite this Article** Chapman, C. , Wakelin, J. , Artacho, E. , Dove, M. T. , Calleja, M. , Bruin, R. and Emmerich, W.(2005) 'Workflow issues in atomistic simulations', *Molecular Simulation*, 31: 5, 323 — 328

**To link to this Article:** DOI: 10.1080/08927020500065223

**URL:** <http://dx.doi.org/10.1080/08927020500065223>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

# Workflow issues in atomistic simulations

C. CHAPMAN<sup>†\*</sup>, J. WAKELIN<sup>‡</sup>, E. ARTACHO<sup>‡</sup>, M. T. DOVE<sup>‡</sup>, M. CALLEJA, R. BRUIN<sup>‡</sup> and W. EMMERICH<sup>†</sup>

<sup>†</sup>Department of Computer Science, University College London, Gower Street, London WC1E 6BT, UK

<sup>‡</sup>Department of Earth Sciences, Downing Street, Cambridge CB2 3EQ, UK

(Received July 2004; in final form October 2004)

This article describes the techniques and mechanisms that have been used to tackle workflow problems encountered in the *eMinerals* project. We examine how established tools and technologies can be brought together to specify and deploy a *computational process*, consisting of a set of jobs and tasks, on our production level mini-grid infrastructure, with respect to a specific problem—the distribution of calculations required to determine, in a systematic way, the mechanisms by which pollutant molecules such as DDT, dioxins and biphenyls, become bound to soil minerals. We also briefly discuss the use of data standards such as the Chemical Mark-up Language (CML) and web-service based grid standards as a means to facilitate workflow specification.

**Keywords:** Workflow; Grid; Combinatorial problems; CML

## 1. Introduction

One of the aims of the *e-minerals* project [1,2] is the systematic study of adsorption characteristics of pollutant molecules on soil mineral surfaces. The problem is very complex from multiple points of view such as:

- i) In order to establish significant trends and possibly, correlations between adsorption energies and phenomenological results, the study of a large enough set of systems is required. The targeted family of organic pollutants of interest has over 200 different molecules, which, together with the several relevant soil minerals, their likely surfaces, possible edges, steps etc. make a huge combinatorial problem with lots of (fairly independent) calculations. We are particularly interested in the chlorinated dibenzodioxins, dibenzofurans, biphenyl, DDT and PCB families. Some of them are of industrial origin, while others are from agricultural practices. These compounds are particularly persistent in the environment, and represent a pressing problem even if some of them are no longer used.
- ii) Each separate case requires many calculations (defining by “calculation”, a geometry relaxation to a close local energy minimum), in as much as many possible locations and orientations are *a priori*

possible for a given molecule on a given surface.

- iii) Different situations will require different levels of theory (from empirical force fields to *ab initio* calculations).

These characteristics demand special computational treatments beyond which is nowadays customary for condensed matter simulations. It requires efficient automation of the calculations, including inter-code operability, in addition to making simultaneous use of large amounts of computational resources.

These type of workflow problems are inherent to grid computing. Most scientific problems cannot be solved using only a single application but usually require the definition of a *computational process* or *workflow process*, consisting of a set of tasks or jobs, which must be executed as a whole for significant results to be obtained. However, the grid is a relatively complex environment to deploy such computational processes. It consists of a number of resources independently managed, each with its own scheduling and operating system, and each offering very different execution environments: one must be able to coordinate the remote submission of jobs across these heterogeneous and autonomous resources, transfer any required input and output data, monitor the job executions, handle potential faults and most importantly, be capable of responding to the *dynamic* nature of

\*Corresponding author. Department of Computer Science, University College London, Gower Street, London WC1E 6BT, UK.

the grid—where the availability of specific resources might vary greatly from one moment to the next.

From a user perspective the problem is relatively simple to define: how can one specify a computational process and simply ‘throw’ as many resources as possible at it in order to be completed in a timely fashion and with minimum intervention? But from a computing perspective, no single tool or solution exists to tackle the workflow problems that we have encountered so far in the *e*-minerals project.

Here we will discuss the ways in which we have brought together existing tools and technologies to allow users to specify computational processes and deploy them on our *e*-minerals mini-grid infrastructure—as well as non-dedicated *e*-science resources such as HPCx [3]—with respect to the scientific problems that we aim to solve. We will also discuss how the use of standards such as the Chemical Mark-up Language (CML) will facilitate these processes, and how the adoption of web services by the grid community provides us with more flexibility in specifying these processes in the near future.

## 2. A real world problem: pollutants in soils

The workflow problem that we introduce here as our reference use case is the management and distribution of calculations required to determine, in a systematic way, the mechanisms by which pollutant molecules such as DDT, dioxins and biphenyls, become bound to soil minerals.

Though the scientific details are outside the scope of this article (these are described more extensively in [4]), this has all the traits of a complex simulation scheme: there are large numbers of relevant molecules, minerals and mineral surfaces that need to be taken into account, and for each combination we need to systematically investigate the potential energy surface of the system, which in itself requires a large number of calculations.

The computational process is shown in figure 1. This process, consisting of SIESTA jobs, calculations dealing with energy minimisations, is a particularly good example of the type of workflow requirements that can be encountered in a grid environment. A vast majority of the tasks can be executed in parallel, but some are inter-dependent, where the output of one job is processed and fed into another, and after each calculation, a small post-processing step is required before the next simulation. Recursion may be required, and we must also take into account the large number of jobs, and the potentially large amount of data that will be produced. This neither makes it possible for the computational process to be handled manually nor can it be computed by a single machine in a reasonable amount of time.

Hence, from a computational perspective, the structure of this process is a particularly complete example of the type of workflow problems that may be deployed in a

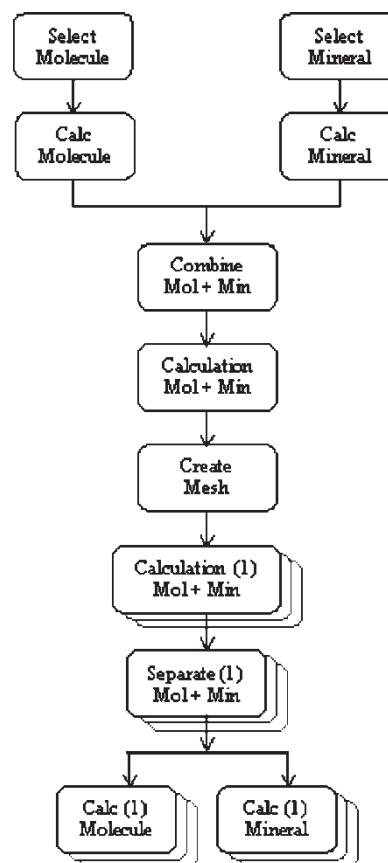


Figure 1. Sequence of simulations required to model the interactions between pollutants and soil minerals.

grid environment. It is a combination of sequential and parallel executions and further more, it is not static in nature. The outcome of a particular subset of the process (1) in figure 1 could modify the course of action to be taken next. For example, a particular result may or may not require further processing to take place before moving on to other tasks in the process.

Coupled to job executions is the issue of data management: input data needs to be transferred alongside jobs to the machines responsible for their execution, output data needs to be stored and later retrieved from data stores, data might need to be parsed and converted to the format expected by a particular application. These data management tasks must be fully integrated into our computational process.

The difficulty in managing all of these tasks manually, particularly when jobs may take considerable amount of time, in a scale of days or weeks, or may simply be too numerous to handle, means we need to *automate* the workflow process.

## 3. The *e*minerals mini-grid

The environment that we use to deploy our use case process is the *e*minerals mini-grid—our production level grid infrastructure. As described in [5–7], it encapsulates

a number of contributed and shared resources such as:

- 3 PBS-managed 16-node clusters;
- The UCL Condor pool (~1000 nodes) and another Condor pool at Cambridge (26 nodes);
- An IBM pSeries parallel computer;
- A Sunfire 880;
- UK e-science level 2 grid resources such as HPCx and National Grid Service (NGS) systems.

It is worth briefly describing some aspects of our mini-grid architecture that must be taken into consideration when deploying a workflow.

All resources are accessed via a globus resource allocation manager (GRAM) interface—either 2.4 or 3.2 [8,9]. The GRAM is meant to provide a secure and *standard* interface for job submission and monitoring, which will hence conceal from the users many differences that exist between the different scheduling systems run by the underlying resources.

Regarding data management, a number of storage resource broker (SRB) [10,11] vaults were set up on our dedicated resources to support the high data storage requirements of our users. The SRB essentially brings together a number of separate data stores (vaults) in order to make them appear to the user as a single unified virtual file system. It also provides functionality such as data replication in order to allow increased robustness.

Though our mini-grid is in itself a powerful infrastructure [7], the fact that these resources are independently managed proves to be a relatively big obstacle to overcome when dealing with workflows. The GRAM may provide us with a standardised way of submitting jobs to a resource, but it is still merely a job submission interface and does not provide us with the means to manage and queue multiple submissions to a single or multiple independent resources. There are, for example, no simple and centralized mechanisms to determine resource availability; some resources may be heavily used, whilst others may be left mostly idle, and this may vary during the lifetime of the computational process. We must also find means of linking GRAM submission to SRB storage functions to ensure that required input data is retrieved from the SRB and transferred to the resource alongside the job for execution and output data stored back into the SRB upon job completion.

It is hence left to the client (i.e. the machine from which the user is initiating the computational process) to manage the deployment and correct execution of the computational process across these resources. This task would consist of:

- Selecting suitable and available resources to run jobs;
- Handling individual job submissions to resources according to the workflow requirements;
- Manage the transfer of executable binaries and input data retrieved from the SRB stores;

- Monitoring the correct execution of the jobs;
- Rescheduling jobs elsewhere if, for example, errors occur or new and better resources become available;
- Transfer jobs back into the SRB for storage.

## 4. Tools and usage techniques

### 4.1 Job submission management

Our approach in the e-minerals project has been to build upon established tools and technologies, such as Condor, Globus and the SRB [6,7], and we have maintained this principle when dealing with workflow problems. The tools and technologies that we have adopted to handle the computing aspects (job submission) are all part of the widely adopted Condor system [12–14]. Though primarily a high throughput computing resource management system for machine pools and clusters, it also provides us with the following client tools to manage the submission of jobs to grid resources:

- *Condor-G*: Condor-G is a *client-side job scheduler*, used to submit jobs to resources presenting a Globus GRAM interface, such as our mini-grid resources. Essentially, it provides means for users to “queue” the jobs to be submitted. It allows a large number of jobs and their requirements to be specified very simply, whereupon it will manage their individual submissions, handle the transfer of input and output files between client and resource and monitor the correct execution of the job. It is also much more flexible and convenient to use than Globus client tools, relying on the traditional Condor submission interface.
- *Condor-G with advertisements*: One of the limitations of using Condor-G is that users must specify explicitly the resource on which they want a job to be run, and this does not really allow users to take advantage of an environment consisting of many different resources with different capabilities. To partially overcome this limitation, Condor-G can provide limited *matchmaking* capabilities: users can define advertisement files describing the characteristics of available resources, such as the operating system, memory available, maximum amount of jobs that can be run, etc, which will be used by Condor-G to determine which site is most suitable for a specific job to be run.
- *Condor glide-in*: Glide-in is the “step above” Condor-G, providing us with means to respond to the dynamic nature of the grid. It essentially allows users to build a *personal Condor pool* out of grid resources. It achieves this by submitting server-side components (daemons) of Condor to selected sites as jobs via the GRAM. Once these daemons have been allocated resources and started, they will report back to the client and notify it of resource characteristics and availability.

Condor-G can then submit jobs directly to these daemons, which will be responsible for monitoring their execution locally. This approach provides us with additional functionality, such as checkpointing—where the state of a job is saved on a regular basis so that it can be restarted from where it was stopped last in case of failure—or migration of jobs—the execution of a job can be stopped and continued on another resource. From a workflow perspective, this addresses many of our requirements: allocations can be requested from all the grid resources, and Condor will then dynamically assign queued jobs to suitable resources as they become available. Multiple resources can be exploited in parallel in this manner.

– *Directed Acyclic Graph Manager (DAGMan)*: DAGMan is used in conjunction with the Condor-G scheduler (and hence glide-in) and allows us to specify dependencies that may exist between jobs. Following a specified sequence, DAGMan will submit jobs to the scheduler and monitor the queue for job terminations, upon which it will submit the next set of jobs in the sequence. The process is specified in the form of an acyclic graph, with jobs represented as parents and children nodes in the graph. It also provides us with means to run pre- and post execution scripts on the client machine, which can be used, for example, to rename files or to simplify cleaning tasks.

An appealing aspect of these Condor tools is the fact that they require only the minimum to be set up on the resource. A user in a grid setting generally has very little influence over the management of the resources, and hence cannot make any assumption as to what would be available on the server side. The only requirement of the Condor middleware is a correctly configured and, in the case of glide-in, a complete (i.e. alongside the GRAM, GridFTP, and other Globus tools are required) Globus installation—Globus being a *de-facto* standard in grid computing.

## 4.2 Data management

There are three aspects of data management that must be taken into consideration: data storage, data transfer—between client, execution and storage sites—and data transformation—which deals with the parsing and conversion of data from one application format to another.

As described previously, the SRB is used to handle our storage needs. The most obvious and perhaps, simplest way of managing interactions with the SRB is for the client to retrieve input files and binaries from the SRB before job submission and store the output data into the SRB upon job completion; the actual transfer of files to and from the resource is being handled by Condor-G's file transfer mechanisms. However, the client might quickly become a system bottleneck. Indeed, it must have sufficient storage capacity and transfer capabilities to handle the data produced and required by all jobs in a

process, which may become highly problematic when these increase in number.

A much preferable approach is to ensure that data is transferred *directly* between resources and the SRB, with the client only being responsible for orchestrating the transfers remotely. This can be achieved by using DAGMan to perform the following sequence of actions when submitting a job to a specific resource:

- Before submitting a job, a script is sent to be run on the selected resource, which will set up the necessary environment, retrieve required input files and binaries from the SRB and store them in a temporary location.
- The job is then submitted. Rather than relying on file transfer mechanisms, job requirements will specify the location of the required files on the remote resource, and a location as to where the output files should be stored on the remote resource upon completion.
- A final script is sent over to the resource, which, when run, will retrieve the output files from the specified location and store them into the SRB. The remote environment is then cleaned up and the script exits.

DAGMan will ensure that these three tasks are performed in order and that each task is completed before the next begins. Appropriate applications were written to simplify the specification of these tasks, so that data management tasks can be specified within a Condor-G submission file, without having to actually deal directly with the DAGMan scripts. These applications enable users to specify their jobs as they would normally do with Condor-G, but with an additional set of parameters allowing them to specify which files should be retrieved and stored into the SRB from the actual resource. The actual three-step process described above would hence be abstracted from the user by the application.

It is also possible to use DAGMan for data transformation. Before and after submission of each job, a small script can be run to parse and convert the data into and from the format expected by the executable program. However, the use of data standards can considerably ease this task and this is discussed in a later section.

## 5. Applying these mechanisms to our workflow process

Though the tools and mechanisms described above allow us to meet many of our workflow requirements, there are still many issues that need to be addressed. One of the primary difficulties that we have encountered in applying these mechanisms is that they result in a potentially large collection of scripts, submission files and DAG-Man specification files being required to specify the process. In order for such a process to be manageable, it is necessary to adopt a unified approach to process specification. We have achieved this by writing a case specific front-end system using Perl, which will generate all the required files and submit these to the Condor



middleware according to the user's requirements and input parameters.

If we consider individual tools and technologies, though Condor-G and DAG-Man worked in most cases "out of the box", the requirements of Condor glide-in, with respect to infrastructure layout, have proven to be difficult to meet. In order to use glide-in with a cluster resource, it needs to be able to access individual nodes in order to communicate with the daemons it has submitted as jobs. This unfortunately is often difficult: individual execution nodes are often behind a firewall or on private local area network (LAN) and hence cannot be accessed externally. There are means to circumvent this problem (open firewall ports—add client to private LAN), but we can only achieve this with resources over which we have administrative control (i.e. dedicated e-minerals resources). Hence, we have built our initial workflow system limiting ourselves to using mainly Condor-G, advertisements and DAG-Man/SRB links into our application.

## 6. Adopting standards to facilitate workflow management

### 6.1 The Chemical Mark-up Language (CML)

Chemical mark-up language (CML) is an extensible mark-up language (XML) based data formatting standard for chemical data [4,15–17]. It defines a standardised way of representing this data, concentrating specifically on molecules, and supporting amongst other things a hierarchy for compound molecules, reaction and macro-molecular structures/sequences. Being based on XML, it can be easily extended to incorporate additional information if required. CML has been described in [6].

With regard to workflow and combinatorial problems, adopting a data presentation standard can solve the problem of incompatibility between application specific-formats and remove the need for data parsing and conversion. Indeed, each application used in a workflow may have its own specific ways of representing similar information and it is often necessary to rely on parsers to transform data from one representation to the next. Of course, different parsers may be required for different types of conversion.

CML-enabled versions of applications allow increased compatibility and facilitate workflow integration. The e-minerals team has been involved in a number of CML related projects. We have, for example, participated in the development of a CML-enabled version of SIESTA [2] for use in the case described above and are currently in the process of developing a CML-enabled version of DL-POLY, modifying individual functions of the application for these to be capable of processing or producing CML data.

### 6.2 Web services for grid computing

The recent adoption of web services by the grid computing community, initiated with the open grid services

architecture (OGSA) [18] and later continued with the web service resource framework (WSRF) [19], has opened up access to a wide range of web service technologies and tools, and of particular relevance here, are web service orchestration technologies, such as the business process execution language (BPEL) [20].

These differ from traditional grid workflow specification tools by specifying a workflow in terms of service invocations rather than jobs. This approach allows all the various aspects of our workflow—execution management, data management and general service management aspects such as fault handling—to be specified as a single integrated process rather than handled separately. The BPEL language also allows us to specify conditional statements (where we can specify behaviour according to different outcomes), specify exceptions (where one can state what actions to take when an error occurs), and manipulate XML-based data (such as CML) within the process rather than have modifications of input files being performed via independent scripts.

Adopting BPEL would allow our workflows to be specified in a unified and manageable manner. The only drawback that we have currently identified would be that specifying such processes require an intricate knowledge of the inner workings of web services and an interesting area of research would be the abstraction of such a language to a scientific level (e.g. using CML as the data standard). We hope that current work with WSRF [21] will enable us in the future to further investigate the use of BPEL as a workflow specification language.

It should be noted that many grid technologies have been or currently are being brought in line with web service standards, including the Globus toolkit, the main building block of our mini-grid and Condor, development of which is partially handled by e-minerals project members. These standardisations of tools will considerably simplify interactions with a wider range of technologies at the infrastructure level.

## 7. Conclusion

We have described here some of the techniques that we are currently using in the e-minerals project to deal with workflow specification and management, and some of the future directions that are being taken. Our approach so far has been to focus on "out of the box" technologies, in order to reach a production level infrastructure within a reasonable amount of time, and also to demonstrate that much can be actually achieved with existing tools and technologies. Many of the problems that we have faced are mainly related to the integration of technologies, and most development that has been done has gone into simplifying the means by which we specify workflows, and reduce the need for users to actually deal with the complexity of specifying every interaction.

This workflow process was an opportunity to identify repeating patterns, such as the need to transfer data from the SRB to the actual resource at job execution time as described in section “Data management”, and develop applications that would operate at a higher level of abstraction.

There is nevertheless still much to do to attain a higher level of flexibility—specifying conditional statements and exceptions, dynamically modifying input data, etc.—whilst still maintaining a certain degree of abstraction in order to make complex workflow specification relatively accessible for users without important computational skills. Hopefully, the incorporation of web services, CML and BPEL will go some way towards assisting us at the infrastructure level. At the very least, this will provide with a sound basis to develop further higher-level services such as the DAGMan/SRB application (section “Data management”) to assist the users in achieving their goals.

## Acknowledgements

We are grateful to NERC for funding of the *e*-Science project “Environment from the Molecular Level”.

## References

- [1] M.T. Dove, M. Calleja, J. Wakelin, K. Trachenko, G. Ferlat, P. Murray-Rust, N.H. de Leeuw, Z. Du, G.D. Price, P.B. Wilson, J.P. Brodholt, M. Alfredsson, A. Marmier, R.P. Tyer, L.J. Blanshard, R.J. Allan, K. Kleese van Dam, I.T. Todorov, W. Smith, V.N. Alexandrov, G.J. Lewis, A. Thandavan, S.M. Hasan. Environment from the molecular level: an *e*science testbed project. *Proceedings of UK e-Science All Hands Meeting 2003*, pp. 302–305, (EPSRC, ISBN 1-904425-11-9).
- [2] S. Wells, M. Alfredsson, J. Bowe, J. Brodholt, R. Bruin, M. Calleja, R. Catlow, D. Cooke, M. Dove, Z. Du, S. Kerisit, N. de Leeuw, A. Marmier, S. Parker, D. Price, B. Smith, H. Spohr, I. Todorov, K. Trachenko, J. Wakelin, K. Wright. Science outcomes from the use of grid tools in the *e*minerals project. *Proc. Of the All Hands Meeting 2004*, Nottingham (2004).
- [3] HPCxm capability computing, documentation, <http://www.hpcx.ac.uk/support/documentation/index.html>
- [4] A. Garcia, P. Murray-Rust, J. Wakelin. The use of XML and CML in computational chemistry and physics programs. *Proceedings of the UK e-Science All Hands Meeting 2004*, pp. 1111–1114 (2004), (ISBN 1-904425-21-6).
- [5] R.P. Tyer, L.J. Blanshard, K. Kleese van Dam, R.J. Allan, A.J. Richards, M.T. Dove. Building the *e*Minerals minigrid. *Proceedings of UK e-Science All Hands Meeting 2003*, pp. 483–487, (EPSRC, ISBN 1-904425-11-9).
- [6] M. Calleja, L. Blanshard, R. Bruin, C. Chapman, A. Thandavan, R. Tyer, P. Wilson, V. Alexandrov, R.J. Allen, J. Brodholt, M.T. Dove, W. Emmerich, K. Kleese van Dam. Grid tool integration within the *e*Minerals project. *Proceedings of the UK e-Science All Hands Meeting 2004*, pp. 812–817 (2004), (ISBN 1-904425-21-6).
- [7] M. Calleja, R. Bruin, M.G. Tucker, M.T. Dove, R.P. Tyer, L.J. Blanshard, K. Kleese van Dam, R.J. Allan, C. Chapman, W. Emmerich, P.B. Wilson, J.P. Brodholt, N. Thandavan, V.N. Alexandrov. Collaborative grid infrastructure for molecular simulations: The *e*Minerals minigrid as a prototype integrated compute and data grid. *Mol. Simul.*, (in press).
- [8] The globus project. (2003), <http://www.globus.org/toolkit>
- [9] I. Foster, C. Kesselman. Globus: a meta-computing infrastructure toolkit. *Int. J. Supercomput. Appl.*, (1997).
- [10] R.W. Moore, C. Baru. Virtualization services for data grids, Chapter 11. In *Grid Computing: Making The Global Infrastructure a Reality*, F. Berman, A.J.G. Hey, G. Fox (Eds), John Wiley (2003).
- [11] M. Doherty, K. Kleese, M. Williams, B. Strong, L. Blanshard, R. Tyer, T. Folkes, D. Corney, A. Sansum, M. Bly, N. White. SRB in action. *Proceedings of UK e-Science All Hands Meeting 2003*, pp. 51–58, (EPSRC, ISBN 1-904425-11-9).
- [12] M. Livny, K. Miller, T. Tannenbaum, D. Wright. Condor—a distributed job scheduler. In *Beowulf Cluster Computing with Linux*, Thomas Sterling (Ed.), The MIT Press, (2002).
- [13] Condor Team, Condor version 6.7.0 manual. University of Wisconsin-Madison. (2003), <http://www.cs.wisc.edu/condor/manual/v6.7/>
- [14] P. Wilson, W. Emmerich, J. Brodholt. Leveraging HTC for UK *e*Science with very large condor pools: demand for transforming untapped power into results. *Proceedings of the UK e-Science All Hands Meeting 2004*, pp. 308–315 (2004), (ISBN 1-904425-21-6).
- [15] P. Murray-Rust, H. Rzepa. Chemical markup, XML and the world-wide web. 2. Information objects and the CMLDOM. *J. Chem. Inf. Comput. Sci.*, (2001).
- [16] Y. Zhang, P. Murray-Rust, M.T. Dove, R.C. Glen, H.S. Rzepa, J.A. Townsend, S. Tyrrell, J. Wakelin, E.G. Willighagen. JUMBO—an XML infrastructure for *e*Science. *Proceedings of the UK e-Science All Hands Meeting 2004*, pp. 930–934 (2004), (ISBN 1-904425-21-6).
- [17] J. Wakelin, P. Murray-Rust, S. Tyrrell, Y. Zhang, H.S. Rzepa, A. García. CML tools and information flow in atomic scale simulations. *Mol. Simul.*, (in press).
- [18] I. Foster, C. Kesselman, J. Nick, S. Tuecke. The physiology of the grid: an open grid services architecture for distributed systems integration. Globus Project (2002)
- [19] I. Foster, J. Frey, S. Graham, S. Tuecke, K. Czajkowski, D. Ferguson, F. Leymann, M. Nally, I. Sedukhin, D. Snelling, T. Storey, W. Vambenepe, S. Weerawarana. Modelling stateful resources with web services. Version 1.1, IBM (2004)
- [20] T. Andrews, F. Curbera, H. Dholakia, Y. Golland, J. Klein, F. Leymann, K. Liu, D. Roller, D. Smith, S. Thatte, I. Trickovic, S. Weerawarana. Business process execution language for web services. (2003), Specification v1.1, <ftp://www6.software.ibm.com/software/developer/library/ws-bpel.pdf>
- [21] C. Chapman, P. Wilson, T. Tannenbaum, M. Farrellee, M. Livny, J. Brodholt, W. Emmerich. Condor services for the global grid: interoperability between OGSA and condor. *Proc. Of the All Hands Meeting 2004*, Nottingham (2004).