Individual Grant Review Report

Workflow Optimisation Services for e-Science Applications
GR/S24886/01

David W Walker[1], Robert J. Allan[2], and John Darlington[3]

1. Background and Context

The Grid computing community is now converging on a service-oriented architecture in which applications are composed from geographically dispersed, interacting Web services, and expressed in a workflow description language, typically based on XML. Workflow techniques generally enable a collection of services to be combined dynamically. However, although there is broad consensus on the overall architecture of the Grid there are many unresolved issues that are still active research areas and for which implementations are not publicly available.

Grid computing presents many challenges to both middleware developers and application scientists. Amongst these is the discovery of resources to perform a particular task or application, and where multiple resources perform the same function, selecting the optimal one with respect to a set of user-specified criteria. The importance of such optimisation issues is highlighted in service-rich environments in which an application may be able to select from a number of semantically equivalent services that may be characterised by differing performance, cost, and quality of solution. By definition, such optimisation needs to be undertaken in a dynamic environment, within which resource properties can change. Optimisation therefore must utilise a number of heuristic techniques to enable service selection.

The Workflow Optimisation Services for e-Science Applications (WOSE) project has investigated techniques for optimising applications composed of Web services, and developed prototype systems that demonstrate the findings of the research.

This report describes the research contributions made by the WOSE project in the areas of service discovery and selection, the restructuring of workflows, and portals for application integration. Tradeoffs between static and dynamic service discovery are addressed in terms of when tasks are bound to specific services. Interoperability between workflow description languages is also considered. Features of the framework are illustrated with an implementation that supports dynamic invocation of Web services that are coordinated using the Business Process Execution Language (BPEL) workflow description language.

As described in the original research proposal, this project involved collaboration between Cardiff University, Imperial College London, and Daresbury Laboratory, and is divided into four main work packages.

1. Service composition and discovery.
2. Application integration.
3. A workflow optimization service.
4. Dissemination.

Tasks 1 and 4 were led by Cardiff University. Task 2 was performed by Daresbury Laboratory, and task 3 by Imperial College.

2. Key Advances and Supporting Methodology

Our review of the research achievements of the WOSE project is broken down by work package.

2.1. Service composition and discovery

Task 1 involved the design, implementation, evaluation, and use of service-oriented middleware for dynamic service selection. The aim was to bind selected abstract services in a workflow to concrete services immediately before the service is to be run. A set of semantically equivalent services that satisfy the abstract

[1] School of Computer Science, Cardiff University, 5 The Parade, Roath, Cardiff CF24 3AA.
[2] Daresbury Laboratory, Daresbury, Warrington, Cheshire WA4 4AD
[3] Department of Computing, 180 Queen's Gate, Imperial College, London SW7 2AZ

service specification are discovered through a service registry. Out of these, the service that best satisfies the user-specified selection criteria is chosen and invoked. The selection criteria can be quite general – for example, cost, numerical accuracy, or service response time. In the evaluation work performed, minimum service response time was used as the selection criterion. In the prototype middleware developed in the WOSE project, a proxy service was developed that, given an input abstract service specification, discovered a set of conforming Web services from one or more UDDI registries, selected one of these services based on runtime information gathered from the service hosts or on historical performance information stored in a database, and then invoked the selected service. The proxy service, the discovery service, and the selection service are the main components of the prototype WOSE architecture for dynamic service selection. The prototype WOSE architecture is described in [1,2]. Experiments were carried out using bioinformatics and molecular dynamics workflows, and used runtime information gathered from service hosts using the Ganglia distributed monitoring system[4]. BPEL was used as the workflow language, and ActiveBPEL as the workflow engine[5]. These experiments are described in [3,4].

In addition to the main focus of this work package on the prototype system for dynamic service selection, two related pieces of work were performed. The first was a web-based tool, based on XSLT, for the automatic conversion of workflows written using a subset of Xscufl to BPEL. This allowed workflows from the Taverna project[6] to be run using the ActiveBPEL workflow engine. This tool is described in [5]. The second piece of work was a statistical performance model, based on the first order statistic, that quantitatively relates the performance benefits of selecting from multiple service instances to the characteristics of the probability density function describing the service response time for a particular service instance. This work quantifies the expectation that there is a greater potential for performance benefits when the probability distribution is wide. This work will be described in a paper currently in preparation that will be submitted to a special issue of the journal *Scientific Programming* on dynamic workflows (see Section 2.4).

The main research contribution of the work performed in this Work Package was the demonstration, using both experimental results and a statistical performance model, of the impact of dynamic service selection on service response times. This work shows that dynamic service selection will be effective in reducing service response times when there is sufficient variation of workload on the host machines, and when the workload varies slowly compared with the timescale for service execution.

2.2. Application integration

Work carried out at Daresbury has focussed on issues of usability and integration with several existing e-Science applications. The applications have been from the following projects:

- HPCPortal and NGS Portal (EPSRC): Multi-functional Java portal (based on uPortal and Sakai) showing Grid, collaboration, data and information management services. Currently a system for managing and using JSDL documents for application submission has been released. We plan to release a similar tool for management of workflow documents, targetting distributed enactment engines as identified from a UDDI registry.
- Integrative Biology (EPSRC): from chemical reaction to whole-organ modelling of processes in cancer and heart disease. We have supported IB and IBVRE in their investigations of the use of workflows in workbench and portal environments.
- e-HTPX (BBSRC/ DTI): High-throughput Protein Crystallography applications on synchrotron facilities at SRS Daresbury, ESRF Grenoble and Diamond Light Source Oxford. This has been the focus for most of our work in the second half of the WOSE project, see below.
- e-Minerals (NERC): Modelling of physical and chemical processes on mineral surfaces responsible for the transport and handling of pollutants. Uses workflows of multiple applications and visualisation driven from scripts. The e-Minerals My_Condor_Submit system uses Condor-G and includes data input/ output via an SRB service, automatic meta-data collection via an Rcommands service, and linking of applications and semantic tools (RDF/ OWL) using an AgentX service within the DAG-based workflow.
- NW-GRID (NWDA) and Diamond e-Infrastructure: A Web services interface called RMCS has been released recently and the system is being made available to NW-GRID users and for the e-Infrastructure of the new Diamond Light Source facility.

[4] http://ganglia.sourceforge.net/
[5] http://www.activebpel.org/
[6] http://taverna.sourceforge.net/

- ESRC e-Infrastructure (ESRC): a new project funded from 1st October 2006 by ESRC. It will include an evaluation of Kepler and Taverna workflow tools for quantitative social science applications.

Our most recent work on the e-HTPX application pipeline is described in a paper for the Sixth IEEE International Symposium on Cluster Computing and the Grid [6]. This paper investigates the use of the Business Process Execution Language for Web services (BPEL4WS/BPEL) for managing scientific workflows. The complexity, unpredictability and inter-dependency of the components in a scientific workflow often demand great flexibility in a workflow-language in order to support; 1) exception handling, 2) recovery from uncertain situations, 3) user interactions to facilitate interactive steering and monitoring, 4) dynamism to adapt to the changing environment, 5) compensation handling to reverse the effects of previous activities that have been abandoned, and 6) flexibility to support dynamic selection of services at runtime and to support changing data requirements. These requirements are illustrated with examples taken from a real scientific workflow; the e-HTPX project for high throughput protein crystallography. In the context of addressing these requirements, the features of the BPEL4WS specification are discussed, which is widely regarded as the *de-facto* standard for orchestrating Web services for business workflows. An investigation and proposal for the use of the Web services Invocation Framework (WSIF) to extend BPEL is also provided. In summary, by extending BPEL where necessary (in conjunction with standard BPEL features), workflows can be successfully adapted to fulfill all the requirements of complex scientific workflows.

The e-HTPX pipeline, which concerns the analysis and determination of a 3D structure of a large protein molecule from X-ray diffraction images, relies on the user being able to collect and process data via a number of applications which are contained in the CCP4 suite. A portal interface is used to enable users to launch workflows to evaluate which processing method is best and then to run a full analysis with the chosen technique.

In carrying out this work it was necessary to develop Web service interfaces to a number of components in languages such as Java, Perl and C. The consequences of this, which may be considered to be research outcomes, are noted as follows:

- Need to make the choice and execution of workflows simple for users, e.g. from pharmaceutical companies, implied that we have exposed a number of pre-configured workflows via the e-HTPX portal interface.
- The data analysis workflow is actually part of a larger “pipeline” for protein synthesis, crystallisation, diffraction, analysis and structure deposition in the Protein Data Bank at EBI. We did not wish to automate all these steps, so have to cater for human intervention and the ability to swap home laboratory services for facility services
- We identified a need to provide “best practice” guides so that services could be handled correctly. This has identified document-based service invocation based on WSDL to be preferable to the simpler RPC style. Several papers have been written on this.
- We identified the use of industry standards, such as BPEL and WS-RF to be our preferred route. We have suggested how standards such as WS-Addressing and WS-Notification can enhance the use of BPEL and provide more optimised workflows. Tutorials on WS-RF have been written based on our own experiences and are proving very popular (high Google rating)

We have identified the requirement for well-defined data models or semantic tools to support data interoperability between applications linked in workflows. This is almost a pre-requisite. We currently have a preference for flexible data models using semantic technology (RDF/OWL) to map vocabularies rather than trying to impose a large, complex and fixed XML model.

2.3. A workflow optimization service

Within the WOSE architecture, as developed by Cardiff University in [1], the basic optimisation service has been further enhanced by Imperial College with the aim of providing quality-of-service (QoS) support for end-user workflows. This work focused on satisfying QoS requirements of workflows whilst providing a sufficient QoS guarantee. Experimental simulation has been performed using SimJava, a discrete-event simulation package in Java. This is because suitable computing nodes weren’t available to host many different web services. Details of the simulation can be found in [7], [8] and [9]. Different workflows, including one from the particle physics domain, were chosen to validate the claims and analysis of [7], [8] and [9]. Different approaches have been utilised to tackle the issue of reducing workflow failures while providing QoS guarantee. This work can be classified broadly into two categories.

Local Level Scheduling/Workload Allocation: The work presented at the 2006 UK e-Science All Hands Meeting [7] aims at Local Level Scheduling meaning a specific web service is selected for a specific workflow task individually. A stochastic optimisation framework has been developed that provides a sufficient QoS guarantee for the entire duration of workflows. The work aims at reducing workflow failures (workflows failing to meet their QoS requirements). We have developed a 2-stage stochastic programming approach to workflow scheduling using an ILP formulation of QoS constraints, workflow structure, performance models of Grid services and the state of the Grid. The approach gives a considerable improvement over other traditional schemes. This is because the Sample Average Approximation (SAA) approach obtains epsilon-optimal solutions minimised and approximated over uncertain conditions while providing QoS guarantee over the workflow time period. The developed approach performs considerably better particularly when the coefficient of variation of execution times and the workflow complexity are high. At both low and high arrival rates, the developed approach comfortably outperforms the traditional schemes.

Global Level Scheduling/Workload Allocation: The works presented at the MWS workshop [8] in Hong Kong and the ICDCIT conference [9] in India focus on modeling web services through queuing theory and mapping a fraction of workload (collection of workflow tasks) to web services instead of scheduling each and every workflow task one by one as in [7]. This work also focuses on reducing workflow failures.

In [8], a web service is modelled using a G/G/k queue (general inter-arrival time and general service time distributions with k processing threads). Failures (QoS requirement violation) of jobs are minimised by solving a mixed-integer non-linear program (MINLP). This novel approach was evaluated through an experimental simulation and the results confirm that the proposed workload allocation strategy performs considerably better in terms of satisfying QoS requirements of Grid workflows than scheduling algorithms that do not employ such workload allocation techniques. The effectiveness of the workload allocation strategy has been evaluated through experimental simulation. Results confirm that the workload allocation strategy performs considerably better than the algorithms that do not use these strategies. When the arrival rates are low, the workload allocation technique performs similarly to scheduling algorithms based on real time performance information. Workflow and workload nature also do not change the performance of the scheme notably. Moreover execution time variability does not change the performance of the workload allocation strategy significantly for both low are high arrival rates.

In [9], a web service is modelled as an M/M/k (Markovian inter-arrival and service times and k processing threads) queue and a numerical solution for missed deadlines (failures) of Grid workflow tasks is obtained. Again results similar to those in [7] are obtained.

2.4. Dissemination

The research results from the WOSE projects were distributed through international and UK conferences and journal publications. In addition, a publicly-accessible project webpage was maintained throughout the project[7], as well as an internal Twiki for use by project personnel. The prototype system for dynamic service selection was demonstrated at the UK e-Science All Hands Meeting in September 2005. Another major dissemination activity was the organization, with WOSE project funding, of a workshop on *Workflow Optimisation in Distributed Environments*, which was chaired by David Walker and held 19-20 October 2006 at the e-Science Institute (eSI) in Edinburgh[8]. This workshop also received £2600 in funding from the eSI. The workshop featured invited talks from the UK, Europe, and the USA[9]. As a result of the workshop two follow-on activities have been instigated. The first is a special issue of the journal Scientific Programming entitled *Dynamic Computational Workflows: Discovery, Optimization, and Scheduling*, co-edited by Dr. Ewa Deelman (USC Information Sciences Institute, USA) and David Walker (Cardiff University, UK). The second is a planned proposal to the eSI Themes Programme to establish a workflow theme. In addition, Dr Omer Rana provided a tutorial at the Sixth IEEE International Symposium on Cluster Computing and the Grid in Singapore on "Grid-based Scientific Workflow Management". The tutorial was attended by 25 participants – attracting people from the US, Europe and Asia-Pacific region. 15% of the material contained in this four-hour tutorial was based on results from the WOSE project.

[7] http://www.wesc.ac.uk/projects/wose/
[8] http://www.nesc.ac.uk/esi/events/702/
[9] See workshop report at http://users.cs.cf.ac.uk/David.W.Walker/WOSE/WorkflowWorkshopReport.pdf

A complete list of the research outputs from the WOSE project is available online at http://users.cs.cf.ac.uk/David.W.Walker/WOSE/ResearchOutputs.pdf.

3. Project Plan Review

The project funded a post-doctoral research associate at Cardiff University (CU), and various staff and PhD students at Imperial College (IC) and Daresbury Laboratory (DL), in keeping with their matrix management approach to research projects. All personnel funded by the WOSE project are listed on the pro-forma. The project officially began 18 May 2004 when Dr Lican Huang started work as the RA at CU. However, there were delays in appointing appropriate staff at IC and DL. This was compounded by Dr Steven Newhouse, who was one of the co-investigators and the main contact person at IC, leaving IC to work at OMII-UK. The project therefore was granted a six-month no-cost extension, so the end date became 17 November 2006. Dr Newhouse withdrew from the project, and was replaced at Imperial College by Dr A. S. McGough.

Other related work at Daresbury which has been funded separately but has contributed to the objectives of the WOSE project is as follows:

- e-CCP: Looking at semantic models to support data-interoperabiilty between scientific applications developed by the UK Collaborative Computational Projects which can be linked into workflows. This project has produced the AgentX toolkit, written in C for use with existing applications (e.g. written in C and Fortran). AgentX has been successfully used to ensure data interoperability between applications in a number of small workflow scenarios in computational chemistry and materials modelling. See http://www.datarepresentation.org
- JISC e-Framework for Education and Research: Dr RJ Allan has represented the e-Research community at the e-FER Working Group meetings and events (see the paper on e-Research Service Classification which includes workflow management http://www.grids.ac.uk/Papers/Classes/classes.html). Dr Allan has also presented an overview of workflows and work of the WOSE project at the JISC Seminar 13/2/07.

Florian Urmetzer: Florian completed an M.Sc. Thesis at University of Reading in 2005 entitled Evaluating Visual Workflow Management Tools for Distributed Grid Computing. This focussed on usability of the interfaces to workflow management systems. Florian has subsequently contributed to a usability study of portals for the e-Social Science community (funded by ESRC's CQeSS project).

4. Research Impact and Benefits to Society

The vanilla Web service based approach adopted in the WOSE project in fully compatible with service-oriented, standards-based middleware being developed and deployed under the auspices of organisations such as OMII-UK, EGEE, and the OGF. The workflow optimization and service selection prototypes developed and evaluated in the WOSE project are of direct relevance to sophisticated scheduling and Grid management systems that are essential to the efficient use of distributed service-oriented environments, and which will form an important part of the next generation of Grid middleware. The main research impact of the WOSE project has been to demonstrate the potential for future distributed scheduling systems to exploit dynamic service selection and workflow restructuring techniques.

The e-HTPX services and portal are currently being re-factored and deployed on the MX Beamlines of the new Diamond Light Source Synchrotron. We note that DLS is the largest single investment in science in the UK for 30 years. The DLS e-Infrastructure includes data management, metadata collection and harvesting and remote analysis and computation using some of the tools noted above. Asif Akram is now employed on the e-HTPX project.

The full SimJava simulation is expected to be mapped in the near future into the scheduling framework of ICENI II (based on web services), a proprietary Grid middleware at IC. ICENI II has seen many early adopters such as the Immunology Grid project, the DTI markets project.

5. Explanation of Expenditure

Project expenditure was essentially as set out in the proposal budget, except there was an underspend on travel. This arose because the cost of organizing the eSI workshop (see Section 2.4) was less than expected, partly due to the support given by eSI. After careful consideration it was decided to use the travel underspend to support Peter Burnap as an RA at CU for the period May-August 2006. This virement of funds was

justified by the need to carry out further work on the evaluation of the experimental results from the prototype dynamic service selection system, and to develop the statistical model referred to above in Section 2.1.

6. Further Research or Dissemination Activities

The project led to several "spin-off" activities:

1. A special issue of the journal Scientific Programming entitled *Dynamic Computational Workflows: Discovery, Optimization, and Scheduling*, co-edited by Dr. Ewa Deelman (USC Information Sciences Institute, USA) and David Walker (Cardiff University, UK). The call for papers can be viewed at: http://users.cs.cf.ac.uk/David.W.Walker/WOSE/CFP.pdf
2. A proposal is being developed for submission to the eSI Themes Programme to establish a workflow theme.

Experience from the WOSE project is feeding into a number of e-Science projects such as the Diamond Light Source Synchrotron e-infrastructure and the ERSC e-infrastructure mentioned in Section 4.

References

[1] "Workflow Optimisation for e-Science Applications," DW Walker, OF Rana, Y Huang, and L Huang, in Proceedings of the 27th International Conference on Information Technology Interfaces, pages 86-91, June 2005. ISSN 1330-1012.

[2] "Dynamic Web Service Selection for Workflow Optimisation," L Huang, DW Walker, Y Huang, and OF Rana, in Proceedings of the UK e-Science All Hands Meeting, September 2005. ISBN 1-904425-53-4.

[3] "Dynamic Workflow Management Using Performance Data," L Huang, DW Walker, OF Rana, and Y Huang, in Proceedings of the Sixth IEEE Symposium on Cluster Computing and the Grid, pages 154-7, May 2006. ISBN 0-7695-2585-7.

[4] "Workflow Optimisation Through Service Selection," DW Walker, L Huang, OF Rana, and Y Huang, Scientific Programming (submitted February 2007).

[5] "A Workflow Portal Supporting Multi-Language Interoperation and Optimisation," L Huang, A Akram, DW Walker, RJ Allan, OF Rana, and Y Huang, Concurrency and Computation: Practice and Experience (submitted December 2005).

[6] "Application of Business Process Execution Language to Scientific Workflows." A Akram, D. Meredith and R.J. Allan, in Proceedings of the Sixth IEEE International Symposium on Cluster Computing and the Grid, pages 269-274, May 2006.

[7] "Grid Workflow Scheduling in WOSE," Y Patel, A S McGough and J Darlington, Proceedings of the All Hands Meeting Conference, Nottingham, September 2006.

[8] "Average-Based Workload Allocation Strategy for QoS-Constrained Jobs In A Web Service-Oriented Grid", Y Patel and J Darlington, Proceedings of the MWS workshop at EDOC Conference, Hong Kong, October 2006.

[9] "Allocating QOS-Constrained Applications in a Web Service-Oriented Grid", Y Patel and J Darlington, Proceedings of the Third International Conference on Distributed Computing and Internet Technology, Bhubaneshwar, India, December 2006, pages 278-290.