CoreGRID Researcher Exchange Programme Report

Michael Parkin
CoreGRID Industrial Research Fellow
Barcelona Supercomputing Centre
Universitat Politècnica de Catalunya

REP Title	Investigation into creating SLAs for 'Dynamic' Grid Workflows
REP	CR38 UPC - CR32 UOM
Dates of REP	10–21 December 2007 (two weeks)
Hosting Institution	The University of Manchester
Hosting Scientist	Rizos Sakellariou

1 Introduction

The Resource Management and Scheduling (RMS) Institute of the CoreGRID Network has objectives, amongst others, of providing scheduling mechanisms that support future business models for Grid computing. This aspect of resource provision is critical if Grids are to be sustainable in the long term. Fundamental in supporting this objective is the creation of Service Level Agreements (SLAs) in the form of Contracts for Grid resource usage in order that business-level relationships can be formed.

This REP report describes work that was carried out on investigating, creating and describing contracts containing SLAs, carried out as part of the following RMS Institute deliverables:

- Task 6.1: Definition of a Grid scheduling architecture.
- Task 6.3: Workflow scheduling strategies.

This REP proposed contributes to the above deliverables by facilitating work exploring cost function models for some types of workflows; research into creating quality-of-service guarantees or SLAs for workflows executed in a distributed computing environment assumes that the dependency graph between tasks is known before the workflow is executed. However there are situations where, because of the conditional execution of code within each task that may spawn further tasks, the full task graph may not be determined prior to the workflow being executed or during its execution. This type of workflow can be

defined as a *dynamic workflow*. Thus, given a dynamic workflow for execution, the total number of tasks and total job duration may only be determined once the workflow has completed. There is, therefore, a question as to how the original SLA is agreed if the set of resources required are not known prior to execution.

An example of a programming paradigm that produces dynamic workflows is Barcelona Supercomputing Centre's GRID Superscalar environment. At run-time GRID Superscalar converts a sequential application into a workflow which can be executed on different resources within a computational Grid. However, given two different sets of input data or starting parameters, the GRID Superscalar-enabled application may produce different different task graphs for each execution because of conditional execution of code within each task. Thus, the total duration and number of resources required to complete a job can only be determined at the end of the application's execution.

2 Results

This REP allowed us to investigate how a function-based approach could be used to help solve the problems with creating an SLA for a dynamic workflow. During the two weeks the researcher was able to derive and define a set of requirements the cost function must meet. These are that:

1. The user is always charged proportionally for the CPU time used.

2. The user is 'rewarded' for better estimates of the behaviour of their dynamic workflow. In other words, if the estimated and actual job behaviour are similar then this should be reflected in a lower price charged for executing the workflow. However, because a deviation from the estimated time has the knock-on effect of the provider having to reschedule other jobs, the user should also be proportionately penalised for being inaccurate too.

3. The users reward and penalty should also be proportional to their confidence in the estimated behaviour of the workflow. For example, if a user is confident that their estimation is accurate and the actual behaviour proves them right, they should stand to gain more — through a lower price for running the job — than if they were unsure about the behaviour of the workflow. An expression of the users confidence (or uncertainty) in the estimated behaviour of the workflow benefits the provider as confident estimates can be scheduled with greater certainty than those given a low confidence, as a low confidence alerts the provider to the fact that the workflow probably won't behave as estimated and the provider can schedule the workflow with a greater margin of error.

   However, if the user is inaccurate with a confident estimation they should be penalised more than with an inaccurate and unsure estimation as this will be a greater inconvenience to the resource provider. It should also be the case that a highly inaccurate, unsure estimation should result in the final pricing formula tending towards a traditional SLA pricing model.

4. The user not be exposed to an unlimited charge for resource usage. For example, if an error occurs in a computational job and it spawns many more tasks than were

expected the user should be able to state in the SLA that the cost for running the whole job should not go over a certain amount, i.e. they should be able to specify a cut-off value that, when reached, pauses or terminates the computational job.

5. The provider should be able to charge a 'booking fee' proportional to the time reserved by the user to run the workflow. This is fair to the provider as it allows them to charge for the overhead of reserving the resource for the user. It also provides an incentive for the user to actually use the providers resources they have booked.

A cost function that meets the above requirements was then constructed. This function and details of how it operates are being collected as a conference paper to be submitted to the European Conference on Parallel Computing (Euro-Par) 2008. The working title of this paper is 'Function-Based SLA Pricing for Dynamic Workflows'. Notification of acceptance will be given at the start of May.

3 Acknowledgments

This research was made possible through the European Commission's Sixth Framework Programme (FP6) CoreGRID Network of Excellence (Contract IST-2002-004265). The REP researcher would also like to thank ATOS Origin Research & Innovation for their support in part-funding the their CoreGRID Industrial Research Fellowship and John Brooke, co-director of the UK E-Science North West Centre, and Viktor Yarmolenko (University of Manchester) for the time they spent with the researcher during the exchange.