SPECIAL ISSUE EDITORIAL
GRID PERFORMANCE AND DEPENDABILITY

Nigel Thomas

School of Computing Science
University of Newcastle upon Tyne
Newcastle upon Tyne, United Kingdom
nigel.thomas@ncl.ac.uk

Grid computing is a rapidly growing area of study that will potentially make significant changes to the way in which large scale computation is performed. The main drivers for this development originally came from the physical sciences, in particular genomics and particle physics. However, other areas of application have also had a major influence on its development, notably medicine and engineering design. In general any problem that generates large amounts of data requiring analysis will potentially benefit from grid research. More recently the commercial benefits of grid computing have begun to be recognised, both by generating business from providing supporting infrastructure, but also by supporting business decision support directly.

This special issue concerns performance and dependability evaluation of grid based services. To date much of the research effort in grid computing has been directed towards functional issues in getting a working infrastructure deployed. However, more recent efforts have emphasised support for existing and emerging scientific efforts through the provision of high level services. With the involvement of real end users with real problems to solve there has been a far greater awareness of the need to satisfy non-functional requirements. Amongst these, performance and dependability are key discriminators for successful services.

The papers included in this special issue represent a broad cross-section of current activity in the area. The first paper, by Jarvis *et al*, presents an overview of some of the emerging issues. The focus in this paper is on *performability*, that is, the joint consideration of both performance and dependability. The authors argue that these issues should be considered jointly if they are going to have the desired impact on user satisfaction. Seven topics are presented in the paper which the authors consider to be key areas of work that need to be addressed if grid computing is to successfully deliver a satisfactory level of service quality.

The second paper, by Al-Ali *et al*, considers the management of service quality in more detail. Their objective is to improve upon the *best effort* approach that has become the norm amongst Internet based applications. To achieve this they present a quality of service management framework based on *G-QoSM*. Their emphasis in this paper is on network quality of service, employing the *Differentiated Services Architecture* and *Bandwidth Broker*. Network quality of service is considered by many to be the most important and most difficult quality of service management problem in grid computing. The authors show that by employing network reservation a significant increase in performance and predictability can be made.

The next paper is by Looker, Munro and Xu. Here the focus of the experimentation is on the reliability of communication between remote services, although the fault injection approach they describe can be applied to many aspects of quality of service. Fault injection is a powerful technique that is widely used to simulate specific errors in systems in order to test the effectiveness of any fault tolerance mechanisms that are employed. The authors have developed a tool, WS-FIT, to inject faults into systems built from web services. In this case the faults that are injected include delays on packets, thus simulating network congestion. The paper also details how test cases are automatically generated from WSDL specifications. Experimental results are presented for an ex-

ample web service based system of an electronically controlled heater unit.

The final paper in this special issue concerns the performance of a system constructed from unreliable services. The paper, by Martin and Mitrani, details modelling experiments on a system of alternative services, each of which may independently fail. Failures in this case are not considered to be catastrophic, but merely interruptions in the service, which can be resumed once repair has taken place. The crucial task is to determine an optimal distribution of workload across the services in order to minimise the response time. This is a computationally costly, but the authors propose a heuristic policy which they show to be near optimal.

There is a growing community of researchers involved in the performance and dependability evaluation of grid services. In recognition of this there have been a number of new academic workshops initiated in the area. This rise in activity is evidence of the importance of this topic and the number of interesting problems that are to be solved. It is hoped that this special issue will further stimulate interest and prompt further quality research in the area.

BIOGRAPHY

Dr Thomas has been a lecturer in the School of Computing Science at the University of Newcastle since January 2004. Prior to this he was a lecturer at the University of Durham for five years. Most of his reasearch interests lie within the field of performance engineering. His principal activity involves in Markov modelling using queueing theory and stochastic process algebra. The areas of application that interest him most at present are security (including e-voting) and the grid (principally scheduling and hosting problems). He is also working with models of massively parallel systems in computer science, sociology and behavioural biology, using stochastic process algebra.