Super resource-managment for Grid Computing

Adrian Li Mow Ching†, Dr Lionel Sacks† and Paul McKee‡

† Dept of E&EE, UCL, ‡ BTExact Technologies

Abstract: Grid computing has emerged as the new paradigm in distributed computing and has undergone significant developments over recent years. To realise the long term goal of the 'Global Grid' requires a convergence between commercial and technical Grid applications, from large scale computationally intensive applications to real-time services. More importantly the need arises for an infrastructure capable of supporting an 'always on' Grid service. In this context the provision of services across such large, distributed and heterogeneous platforms requires a distributed resource management mechanism with capabilities beyond traditional centralised management. Coupled with the challenges of scale are the complexities associated with the heterogeneity of shared platforms and tiered administrative domains. Resources in a fully Grid enabled environment will be administered locally and constrained by local policies, therefore higher level resource management systems have to accommodate for this dynamic availability. In this paper we discuss the challenges facing higher level resource management within a Grid environment, looking in particular at a distributed, self organised, approach to Grid resource management in a global context.

1. Introduction

The need for high performance distributed computing has driven the development of a set of Grid computing technologies that allows resource sharing across multiple domains. Grid computing is defined as secure, controlled resource sharing across heterogeneous platforms within multiple administration domains creating 'virtual organisations' [1]. Many organisations now require platforms upon which they can run computationally intensive applications. From financial services companies trying to predict the behaviour of stock markets to aerospace companies modelling and developing a new aircraft, the needs for a high capacity computing fabric are apparent. The rise of Peer to Peer computing and the new practises in e-business have created the need for a set of technologies that allow inter-organisational resource sharing. Grid computing technologies has emerged as a solution to the problem. It comprises of a set of technologies that provides the low-level underlying structures capable of supporting resource sharing across distributed and heterogeneous environments [2].

The requirements for a Grid resource management platform can be derived from the following aspects of Grid computing. The Open Grid Services Architecture (OGSA) [3] direction of Grid computing will transform Grid technologies in to a service oriented technology that offers Grid applications as web services. This coupled with the increasing number of Grids that are coming 'online' contribute to the long term goal of achieving a 'Global Grid' that offers an 'always on' Grid service. Although current uses of Grid networks primarily involve large scale simulations, as the Grid gains wider acceptance beyond the realms of research and academia, we will see 'real-time' commercial Grid applications emerge, with possible consumer Grid services in the long term. These require certain quality of service guarantees capable of sustaining a real-time service. Alongside the challenges that concern applications are the complexities of the underlying resources that make up the Grid infrastructure. Resources are geographically distributed, heterogeneous, and administered within different domains, therefore availability is dynamic and conditional upon local constraints. There are multiple tiers to Grid resources, from an end user workstation to clusters of machines. The restrictions set on these machines at the various hierarchical levels bring various scaling issues to the resource management problem.

The Globus toolkit [4], (the de facto standard for Grid computing software) provides an integrated set of basic Grid services that allows interoperability and communication between systems. Resource management components of the toolkit simply create the underlying fabric that allows resources to be harnessed by local resource managers; schedulers, queuing mechanisms and priority policies. Higher level resource management techniques are needed that that can provide guaranteed quality of service, access to resources and guaranteed job completion. These will enable the support for dependable and consistent future Grid services, including but not limited to, high throughput applications and real-time services.

2. Active Networks and ANDROID

ALAN (Application Layer Active Networks) [5] concepts have been proposed to overcome the limitations of existing networks to cope with the future demand of telecommunication services. It provides an environment in which users can load code in to the network, which can then be dynamically loaded on demand and run within the network as a 3rd party service. Allowing new services to be introduced to the network on a more rapid basis eliminates the need for adding existing components to the network through large scale deployment and standardisation. The on going work for the IST project ANDROID (Active Network DistRibuted Open Infrastructure Development) [6] aims to develop a lightweight, scalable management infrastructure for ALAN type networks. The characteristics of an ANDROID network are similar to a Grid network, and so the challenges of resource management are similar in both cases. In an ANDROID network resources are distributed across a large number of nodes. The resource management system has to distribute work loads across the nodes as well as provide certain levels of QOS which are required by the network services. In an ANDROID network event driven resource management policies have been implemented to provide hard QOS guarantees. We hope to utilise the resource management experience [7, 8] gained from the ANDROID project and apply these ideas to resource management in the context of Grids.

3. Resource management

Resources in a Grid environment are heterogeneous, they include processing power, storage, data and specialised equipment. The consumers of the resources are the users who have specific requirements which are expressed in terms of CPU, memory etc. Resources are administered at different tiers within multiple domains. At the lowest level, a machine's CPU and memory is managed by the operating system scheduler, ensuring that the applications have their fair share of processing power. Local scheduling policies of the OS deal with the competing demands for resources. On the local area network or cluster level, network resources are controlled by a local network administrator who imposes policies on all machines within that cluster. At the highest level of the hierarchy, the Internet level, which has is no centralised management structure as it is an aggregation of lower level local networks / clusters / end systems. It is essentially a collection of different people's Grids. Centralised management techniques try to optimise system wide performance but do not scale well. Resource mechanisms that scale in relation to the hierarchical nature of the different domains as well as with regards to the size of the number of nodes are required to manage such a large set of distributed nodes. Resources are controlled under different administrative domains and each restricted by individual policies. Resource availability varies dynamically and is dependant upon the owner's resource requirements and the user's requirements. Therefore we identify that a global resource management infrastructure also has to be self-organised and self-healing. Due to the dynamic nature of resources, it is possible that a resource becomes unavailable, in such a scenario the resource management mechanisms need be able to autonomously find an alternative resource and distribute the job to that resource accordingly.

The current efforts of the Grid community are focussing on local cluster resource management issues including standardisation for APIs, resource description languages and resource allocation protocols. Resource management mechanisms are very well suited for single Grid clusters incorporating optimised scheduling algorithms and load sharing mechanisms, but have not been proven to work well between multiple Grids. The next stage in the evolution of operations Grids is inter-working between the different communities of Grids and underlying mechanisms that can provide hard quality of service.

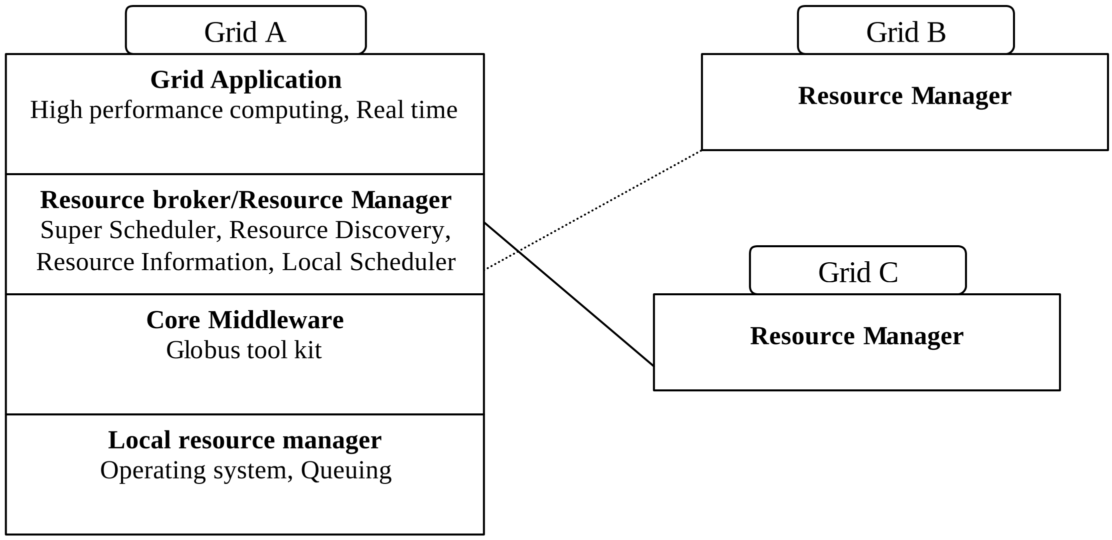

Figure 1. Grid resource management scenario for resource sharing between individual Grids

Figure 1 depicts the structure of components in a distributed resource management Grid scenario. An important function of Grid connectivity and resource sharing is the ability to seamlessly interact with other Grids, utilising their resources and being able to harness the resources within such a dynamic environment. The resource manager in Grid A acts as a super resource manager and apart from using resources from its own Grid, it requires resources from Grid B. However since Grid B's resources are controlled by its own local resource manager resources cannot be guaranteed over long periods of time. The resource manager in Grid A will have to be able harness the aggregate resources available through Grids B and C. The long term realisation of the 'Global Grid' requires decentralised mechanisms in place that can cope with such dynamic relationships to sustain a full end to end Grid service.

4. A self organising approach to resource management

A self organised approach to resource management will provide the infrastructure for a reliable and consistent Grid. The three main components are resource discovery, scheduling, and dynamic resource information. The self-organised function of a resource management system consists of self configuring, self optimising and self healing properties. Self configuring arises from being able to discover resources, and schedule jobs and work load autonomously. In order to provide high throughput, the system should be self-optimising. This can include allocation of work loads to specific resources based on policies that ensure hard quality of service guarantees and preferential attachment to resources. In the event that a resource is no longer available, the system should self-heal. In such a scenario alternative resources that are capable of completing the jobs can be used, the system will adapt and reschedule accordingly.

As an initial piece of work we will investigate the use of small world techniques [9] for resource management. A small world is defined as a highly clustered network with small characteristic path length. It means that individual entities are well acquainted in the vicinity, but yet, a far unit is merely a few degrees of separation away. Small world networks scale better than scale free networks which are dependant on a few larger nodes that can be points of failure within the network. They are suitable for resource discovery across a system due to the low degree of separation between nodes. As an initial study we intend to implement a discovery protocol based on small world principles, capable of discovering the best available resources within a distributed set of Grid nodes. Its performance will be compared to the Meta computing Directory Service (MDS) [10] component of the Globus toolkit which provides information discovery and access to system configuration.

5. Conclusions and further work

This paper has discussed global resource management for Grid computing in the context of providing an 'always on' Grid service. We have identified the need for global resource management techniques that function across the various platforms of the Grid. The scale of the Grid dictates that resource management is distributed with no central pool of information. The lack of control over the underlying resources available, thus providing only dynamic availability, requires management mechanisms that function effectively at the local cluster Grid level as well as between inter-connected Grids. We propose that a self-organised approach to the resource management techniques is taken, capable of self configuring, self optimising and self healing the aggregate resources which it manages. Our initial investigations will involve the implementation of small world techniques for the discovery and monitoring of resources.

6. Acknowledgments

Adrian Li Mow Ching would like to acknowledge the EPSRC for supporting his doctorate studies and to thanks those responsible for his attendance at the recent 5th Global Grid Forum meeting.

7. References

[1] I. Foster, C. Kesselman. and S. Tuecke. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. International Journal of High Performance Computing Applications, 15 (3). 200-222. 2001.

[2] I. Foster and C. Kesselman, (editors). The Grid: Blueprint for a New Computing Infrastructure. Morgan Kaufmann, 1999.

[3] I. Foster, C. Kesselman, S. Tuecke, and J. Nick. The Physiology of the Grid: An Open Grid Service Architecture for Distributed Systems Integration. 2002

[4] The Globus Toolkit. www.globus.org

[5] M. Fry and A. Ghosh, Application level active networking. Computer Networks, 31 (7) (1999) pp. 655-667

[6] M. Fisher, D5: Android Active Networking Architecture, 2001.

[7] O. Prnjat, I. Liabotis, T. Olukemi, L. Sacks, M. Fisher, P. McKee, K. Carlberg, G. Martinez. Policy-based Management for ALAN-Enabled Networks, IEEE 3rd International Workshop on Policies for Distributed Systems and Networks - Policy 2002, Monterey, CA, USA, June 2002.

[8] I. Liabotis, O. Prnjat, L. Sacks. Policy-Based Resource Management for Application Level Active Networks, Latin American Network Operation and Management Symposium, LANOMS 2001, Belo Horizonte, Brazil, August 2001.

[9] D J Watts. Small Worlds: The Dynamics of Networks between Order and Randomness. Princeton University Press, New Jersey, USA. 1999

[10] K Czajkowski, S. Fitzgerald, I. Foster. and C. Kesselman. Grid Information Services for Distributed Resource Sharing. 10th IEEE International Symposium on High Performance Distributed Computing, (2001), IEEE Press, 181-184