Topic 17
Metacomputing and Grid Computing

Alexander Reinefeld, Omer F. Rana, Jarek Nabrzyski, and David W. Walker

Topic Chairpersons

The growing number of computers accessible through a network, such as the Internet, has meant that these computers can be collectively employed to solve complex problems. Since the networks that connect such machines are distributed, many attempts have been made to create an infrastructure for distributing large scale applications across regional and national boundaries, to enable nation-scale or continent-scale computing to be realized. The components within such an infrastructure can also range in complexity from workstation clusters to dedicated parallel machines and data repositories.

New paradigms have been introduced to support the programming and deployment of such computational resources, generally as extensions to existing parallel computing models (such as SPMD) and distributed computing techniques based on distributed object technologies. The primary objective in many such systems is to achieve uniformity in programming and use, whilst supporting heterogeneity and transparency in architectures, operating systems and environments. The use of Web technologies, such as HTML, Perl/CGI and various scripting languages, have provided useful middleware tools for integrating information resources to support such large scale applications. These tools have limitations, and are often hard to extend beyond a given application. Java and CORBA have also played an important role in such developments, providing an infrastructure which facilitates integration across platform and programming language boundaries – albeit at a performance cost.

Emerging network applications in areas such as multi-disciplinary information analysis, distributed collaborative environments, post-genomics, and distributed visualization, require the coordinated use of geographically distributed resources. The development of such applications can be significantly simplified if specialized operations were supported in the infrastructure – such as the provision of caching, authentication, resource discovery, resource scheduling, albeit for heterogeneous environments.

A Grid can be thought of as a collection of resources, and infrastructure services, supporting the development, deployment and operational support for applications that are multi-disciplinary and require complex management. The Euro-Par Topic "Metacomputing and Grid Computing" explores such infrastructure technologies which can be successfully employed in the context of heterogeneous, distributed computing, to enable integration and management of application across organizational and national boundaries.

The first session focuses on the experiences made with existing Grid testbeds. The first two papers discuss the use of the Cactus Grid environment which is in daily use in the astrophysics community. Cactus is a representative for a

R. Sakellariou et al. (Eds.): Euro-Par 2001, LNCS 2150, pp. 805–806, 2001.

whole class of scientific applications with the following characteristics: tightly coupled, regular space decomposition, huge memory and processor time requirements. The first paper, written by Matei Ripeanu, Adriana Iamnitchi, and Ian Foster, presents performance predictions of Cactus running on several thousand nodes in the Internet, while the second paper, authored by Gabrielle Allen *et al.*, presents useful tools for remote steering, remote monitoring and remote visualization that have been developed around the Cactus code. Also the third paper, written by Dietmar Erwin and David Snelling, describes a practical Grid testbed, the UNICORE environment. The paper gives an overview on the UNICORE architecture, including the data model, the abstract job objects and the security model.

The second session starts with two papers on the use of CORBA in distributed environments. Alexandre Denis, Christian Pérez, and Thierry Priol describe a parallel CORBA model for interconnecting different MPI code modules without degradation in performance. Thereafter, Diego Sevilla, José García and Antonio Gómez present a lightweight CORBA components model named CORBA-LC which allows to build distributed applications by assembling independent binary components. The third paper in this session by Nathalie Furmento, Steven Newhouse and John Darlington presents a framework on building computational communities with Java and Jini by federating resources from different organizations.

The third session starts with a paper of Karan Bhatia, Keith Marzullo, and Lorenzo Alvisi on causal message logging protocols which employ a hierarchy of shared logging sites (proxies) to improve the fault tolerance and availability in grid environments. Thereafter, Zoltán Balaton, Péter Kacsuk, Norbert Podhorszki, and Ferenc Vajda propose to adapt the GRM and PROVE tools of the P-GRADE graphical program development environment so that they can be used in the monitoring subsystem of the EU Datagrid project. Finally, Junwei Cao, Darren Kerbyson, and Graham Nudd present an agent-based hierarchical model for resource management in grid environments.