Future Generation Computer Systems 25 (2009) 893–894

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Editorial

# Special section on networks for grid applications

Piero Spinnato [a,*], Pascale Vicat-Blanc Primet [b], Chris Edwards [c], Michael Welzl [d]

[a] LNGS-INFN, Italy
[b] INRIA, France
[c] Lancaster University, UK
[d] University of Innsbruck, Austria

ARTICLE INFO

*Article history:*
Received 9 January 2009
Received in revised form
25 March 2009
Accepted 26 March 2009
Available online 5 April 2009

## 1. Preface

Grid computing has been defined as "coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organisations" (Ian Foster: http://www-fp.mcs.anl.gov/~foster/); see also [1] for a recent overview of the current status and perceived future directions of the Grid. In practice, the above definition implies that a computation intensive application is distributed across a number of computers for the sake of gathering more computing power (or of locating the processing near remote data repositories to limit data movement…). This step from local clusters to massively distributed applications that use the Internet as its underlying means of communication entails a requirement for efficient network usage, but the Internet is a general-purpose network that was not built to provide maximum efficiency to the Grid.

Grid developers and practitioners are increasingly realising the importance of efficient network support. Entire classes of applications would greatly benefit from a network-aware Grid middleware, able to effectively manage the network resource in terms of scheduling, access and use. Conversely, the peculiar requirements of Grid applications provide stimulating drivers for new challenging research towards the development of Grid-aware networks. Cooperation between Grid middleware and network infrastructure driven by a common control plane is a key factor to effectively empower the global Grid platform for the execution of network-intensive applications, requiring massive data transfers, very fast and low-latency connections, and stable and guaranteed transmission rates. See [2] for a review of an illustrative example of an infrastructure of this kind.

When considering support for such applications in the network, there are many issues that require consideration. For example, is the network infrastructure dedicated to a particular (set of) Grid applications, or will the applications be sharing the network resource with other more traditional internet traffic? Will support for applications come from techniques that operate at the network layer, and thus impact on the way individual network routers treat packets, or from the transport layer where congestion control algorithms are adapted to suite the nature of these traditionally high-bandwidth applications. There is also the question of whether the applications require use of the network resources immediately, or are effectively able to specify resource requirements ahead of time, for potential reservation. Beyond the consideration of the network alone, there are the issues surrounding the relationship between the network and the applications, and the mapping of one set of requirements onto the other.

In this special section, the papers (described in more detail in Section 2) all have different approaches to answering some of these fundamental questions. First, "Evaluation of Flow-Aware Networking (FAN) Architectures Under GridFTP Traffic" [Future Generation Computer Systems, this issue], deals with network prioritisation only, and is based on the availability of techniques that are able to identify *flows* of Grid traffic. Second, "Flow Scheduling and Endpoint Rate Control in Grid Networks" [Future Generation Computer Systems, this issue] provides an end system congestion and rate control mechanism, which leads to the definition of network bandwidth allocation profiles, for either dedicated networks, or isolated classes. Third, "Multi-Cost Job Routing and Scheduling in Grid Networks" [Future Generation

* Corresponding address: Servizio Calcolo e Reti, LNGS, INFN, S.S. 17/bis km 18+910, 67010 Assergi (AQ), Italy.

*E-mail addresses:* piero.spinnato@lngs.infn.it (P. Spinnato), pascale.primet@inria.fr (P. Vicat-Blanc Primet), ce@comp.lancs.ac.uk (C. Edwards), michael.welzl@uibk.ac.at (M. Welzl).

0167-739X/$ – see front matter 
doi:10.1016/j.future.2009.03.006

Computer Systems, this issue] provides both immediate and advanced resource reservation for network and end system components. The final paper, entitled "A Secure Credit-Based Cooperation Stimulating Mechanism for MANETs Using Hash Chains" [Future Generation Computer Systems, this issue], moves away from a "traditional" perspective of Grid networking, looking at packet forwarding in a mobile ad-hoc network.

## 2. Papers presented in this special section

In "Evaluation of Flow-Aware Networking (FAN) Architectures Under GridFTP Traffic" [Future Generation Computer Systems, this issue], the authors consider the applicability of Flow-Aware Networking (FAN) in a Grid context. FAN employs per-flow admission control and implicit flow differentiation to avoid the overhead of per-packet prioritisation. The performance of the 2GFAN architecture is compared against DiffServ, using GridFTP sessions as the traffic source. The work is based on simulation, using ns2, and focuses on using throughput and delay parameters. The results show that, used in a Grid scenario, FAN architectures perform better than DiffServ, and that FAN based on Priority Fair Queuing (PFQ) scheduling gives the strongest performance, even though it is a more computationally complex approach.

The paper "Flow Scheduling and Endpoint Rate Control in Grid Networks" [Future Generation Computer Systems, this issue] looks at providing support for the effective movement of massive data sets across either a dedicated network or an isolated traffic class, in order to allow the efficient scheduling of distributed Grid resources. By formulating bulk data transfer scheduling, the paper attempts to provide a model for minimising network congestion, by defining a set of multi-interval bandwidth allocation profiles. Using a real testbed, the paper describes and compares a number of practical solutions for enforcing such allocation profiles within an end point. This includes assessing the complexities of implementing the mechanism in the Linux kernel. Parameters considered within this work include bandwidth, round trip time, and local OS-based timers, and the work is fundamentally based on the ability to describe the volume, active time window and path. The implementation work was performed on part of the Grid5000 testbed, with high speed data rates, i.e. 1 Gbit/s. Results show instability with TCP when the round trip time is large, and the need for traffic shaping to avoid burstiness.

The paper "Multi-Cost Job Routing and Scheduling in Grid Networks" [Future Generation Computer Systems, this issue] looks at the area of resource utilisation in the Grid domain. The authors look at scenarios where resources can be reserved immediately, or in advance. Based on simulation work, the authors propose to register clusters and network utilisation with bit vectors, showing how multi-cost algorithms can be used to solve this complex area. Parameters used include data size and computational complexity, and from the network, path availability and delay. The results show that it is indeed important to consider both processing and computational resources. They also show that the difference in being able to support immediate against advanced reservation is effectively a trade-off between the likelihood of blocking, end to end delay, and the complexity of the deployed algorithm.

As described earlier, Grid computing has been defined as "coordinated resource sharing in virtual teams". This definition may seem to be at odds with the concept of a "wireless Grid", where the fact that nodes can move can drastically limit their availability. One could question whether users of nodes in such an unstable environment could really be regarded as forming a "virtual team", and then consequently question whether the whole notion of a wireless Grid is meaningful. As its title suggests, the paper "A Secure Credit-Based Cooperation Stimulating Mechanism for MANETs Using Hash Chains" [Future Generation Computer Systems, this issue] adds meaning to such potentially unstable environments, rendering the idea of a wireless Grid more plausible. This paper proposes a method to provide incentives for nodes to cooperate in forwarding network packets using credit-based mechanisms. The work in the paper is evaluated using game theory-based analysis, and shows that any level of cooperation by a node will be attainable if the mechanism makes appropriate payments.

## Acknowledgements

The guest editors would like to thank both the authors and the reviewers for their timely contributions, which have collectively helped to make this special edition a success.

## References

[1] George A. Gravvanis, John P. Morrison, Heinz Stockinger, Special section: Defining the grid, experiences and future trends, Future Generation Computer Systems 25 (2009) 399–400 and ref. therein.
[2] Larry Smarr, Maxine Brown, Cees de Laat, Special section: OptIPlanet — The OptIPuter global collaboratory, Future Generation Computer Systems 25 (2009) 109–113 and ref. therein.

**Piero Spinnato** works at the Computing and Network Service of the National Laboratories of Gran Sasso of the Italian National Institute of Nuclear Physics. He graduated in physics at the University of Palermo in 1996, and received his Ph.D. in computer science at the University of Amsterdam in 2003. He has also been at the Edinburgh Parallel Computing Centre, and had appointments from the institute for cosmic physics and applied informatics of the Italian National Research Council, and from a research centre in telecommunications funded by the autonomous province of Trento. His scientific interests include high performance scientific computing and networking, algorithms and technologies for distributed computing, and modelling and simulation of complex systems.

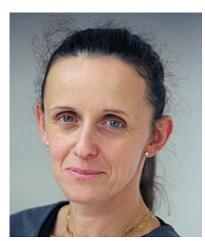

**Pascale Vicat-Blanc Primet** is a senior researcher at the National Institute of Research in Computer Science (INRIA). She leads the INRIA RESO team within the LIP laboratory of the Ecole Normale Superieure de Lyon. She is member of the Grid5000's – French Computer Science's Grid initiative – steering committee, of the scientific "Networks and Telecoms" expert committee of CNRS (National Research and Science Center) and of the French National Research Agency. Her interests include Grid computing and Grid networking, network and Internet protocols, network architecture, quality of service, network measurement, high-speed and low-latency nets, and programmable networks. She has co-chaired the GGF Data Transport Research Group, and participates in a variety of International, European and National Grid projects including: DataGRID, DataTAG, eToile, GRID5000, GdX, EC-GIN, HIPCAL, CARRIOCAS. She a member of IEEE and is active in the OGF (Open Grid Forum) and ETSI standardisation bodies.

**Chris Edwards** has been a Lecturer within the Computing Department at Lancaster University since 2001, where he teaches in the field of telecommunications, networking and operating systems. He has been involved in a number of EU funded projects, most recently u2010 and EC-GIN. On-going industrial sponsorship has come from Cisco Systems in the area of network mobility. In addition, he is currently working in the areas of network infrastructure support for disaster recovery and rescue, network support for GRID applications, and networked multimedia caching. He has authored or co-authored over 30 peer reviewed published research papers in the fields described above, and is a technical programme committee member for a number of international conferences and workshops.

**Michael Welzl** passed his Ph.D. defence at the University of Darmstadt/Germany with distinction in November 2002, and received his habilitation from the same University in June 2007. He spent two years as a research assistant at the Telecooperation department, University of Linz/Austria, before joining the faculty of the newly founded Institute of Computer Science at the University of Innsbruck, Austria, in November 2001, where he now heads a research team on Network Support for Grid Computing in the context of the Distributed and Parallel Systems group.