

OGF-Europe on Software development tools for distributed computing OGF-Europe Community Outreach Seminar Wednesday 4 March 2009, Galilei Room EGEE 4th User Forum/OGF25 & OGF-Europe's 2nd International Event, Catania, Italy

The rise of Virtualised and Distributed infrastructures and the emergence of multi-core processing capabilities have led to a new challenge to accomplish the effective use of compute resources through the design and orchestration of distributed applications. This skill was once considered to preserve high performance and parallel computing, but the new high throughput computing infrastructures offer substantial capacity to support user needs for scalability and agility. As legacy, monolithic applications are replaced with service oriented applications, questions arise about the key steps to be taken in architecture and design to maximise the usefulness of the infrastructures and what lesson can be learned from industry leaders in the design of distributed applications. The OGF-Europe second Community Outreach Seminar aimed at bringing some clarity to the challenge of developing distributed applications, a clear understanding of the tools available to support the software designer and the collection of the best practices across European community.

The workshop was led by Barcelona Supercomputing Center and was chaired by Rosa M. Badia and Daniele Lezzi (BSC, Spain), Andre Merzky (CCT/LSU, USA), Krzysztof Kurowski (PSNC, Poland) and Lorenzo Dini (CERN, Switzerland). The workshop included two sessions; each session introduced a topic emphasising the adherence to standards and provided practical examples of use also from industry to strengthen the relevance for a broad audience. The first session, opened by Krzysztof Kurowski, PSNC, covered the DRMAA OGF standard from a use case driven point of view. Pawel Lichocki, PSNC, continued the talk focusing on academic and industrial use cases; the former was based on presenting the Single Program Multiple Data simplex application, while the latter concerned usage of OGF standards (DRMAA, BES and JSDL) with web-services in order to provide a complete solution to HPC computing; this approach is

currently developed in the context of the SMOA project (successor of OpenDSP project) at PSNC which adopts a refined architecture by means of the design of a specific Web Service interface to OGSA-BES (HPC BES Basic Profile) and Advance Reservation APIs (tested with LSF, SGE, PBSPro)

The workshop continued with a talk, introduced by Thilo Kielmann, Vrije Universiteit, Amsterdam, on the SAGA OGF standard; the aim was to discuss in particular the SAGA API as an effective way of developing distributed applications. The salient features of the API and several real examples of distributed applications and tools that have been developed using SAGA has been highlighted by Shantenu Jha (CCT/LSU).



Krzysztof Kurowski, PSNC, introduces the first session
Image courtesy of Daniele Lezzi, BSC

The first session ended with a description of the ETICS Test System from Lorenzo Dini, CERN. The ETICS System is an advanced build and test system designed to facilitate the management of complex distributed software in general and grid software in particular. The ETICS Multi-Node feature allows users to automate the deployment of tests spanning multiple nodes handling their communication and

synchronisation. The presentation described how to exploit this feature to perform standard compliance and interoperability tests of distributed software. Some practical case studies were presented to test the interoperability in job submission among different middleware implementations (ARC, gLite and UNICORE).

The second session analysed the new challenges that the high performance community is facing in the design of new programming models due the emergence of new multicore and GPU architectures.; while for many years Graphics Processing Unit was only used to perform graphical operations, a new trend of using GPU for handling computation traditionally performed by CPU is arising. Similarly, new multicore processor devices had appeared in the recent years, becoming a revolution in the way computation have been done since now. Both approaches, multicore and GPUs, require right now additional effort from programmers. The session covered trends in programming for these two new approaches to computation and showed potential usage in distributed environments. Antonino Tumeo, Politecnico di Milano and NVidia Cuda champion for Italy, gave an overview of the available hardware, the CUDA framework and GPU-specific approach to programming. Rosa M. Badia presented the current efforts at Barcelona Supercomputing Center for the development of CellSs and SMPs programming models

Next Steps. "The single major activity which supports the uptake of (Grid API) standards are implementations, and in particular implementations by (commercial and non-commercial) middleware vendors. DRMAA and SAGA have an excellent track record in that respect, partially caused by the involvement of vendors in the standardisation process, which motivated early implementations by some of those vendors." Andre Merzky, LSU/CCT

It is inherently difficult to interest vendors in standards though. API level standards like DRMAA and SAGA, however, have the unique ability to directly appeal to the end users; end users are using APIs much more directly as they are using other middleware components with, for example, standardised protocols or service interfaces. Users asking a vendor for the implementation of a standard is the single most effective way to interest vendors in implementing that standard.

Spotlight on **DRMAA** and **SAGA**

- ☀ DRMAA develops an API specification for the submission and control of jobs to one or more Distributed Resource Management (DRM) systems. The objective is to facilitate the direct interfacing of applications to today's DRM systems by application's builders, portal builders, and Independent Software Vendors (ISVs).
- ☀ SAGA is a high-level API that directly addresses the need of application developers of a standardized, common interface across various grid middleware systems and their versions. A guiding principle for achieving this simplicity is the 80–20 rule : serve 80 % of the use cases with 20 % of the effort needed for serving 100 % of all possible requirements.

DRMAA and SAGA are OGF Grid Recommendations, respectively GFD.133 and GFD.90; the full specifications can be found at <http://www.ogf.org/documents/>

OGF-Europe is promoting the OGF standards to the wider community to make them useful for its members and to increase the incentive for Grid middleware vendors, such as EGEE, Platform, OMII, Globus, and others, to implement these standards.



Rosa M. Badia, BSC, during her talk on CellSs and SMPs
Image courtesy of Daniele Lezzi, BSC

The Hybrid Supercomputing

As spotlighted in the second session the ongoing trend is the research of new programming models that could effectively scale from embedded systems, through server and desktop systems, to clusters and grids. Another point that could be investigated in future workshops is the heterogeneous (super)computing; the hardware, GPU, FPGA, CELL, etc., has reached a high level of maturity but there is lot of work to do in the development of the tools that could effectively orchestrate such heterogeneous environments. The current investigations of SMPSs hierarchical from BSC and of CUDA to exploit CPU and GPU mixed environments are clear examples of efforts in this direction.

“GPUs have evolved to the point where many real world applications are easily implemented on them and run significantly faster than on multi-core systems.

Future computing architectures will be hybrid systems with parallel-core GPUs working in tandem with multi-core CPUs.”

Jack Dongarra, Professor, University of Tennessee;
Author of “Linpack”

Right now there is no specific group in OGF that covers multicore and/or GPU programming and in fact there is no need for such. However, it is important that this aspect of the new technologies is taken into account, since future computers will need different programming models than the current ones. It will be interesting for OGF-Europe to select a set of OGF groups where a reflection activity of how these new programming trends should be taken into account from the grid perspective.

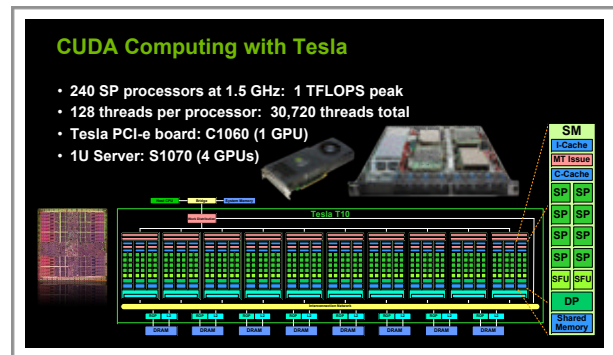


Image courtesy of Antonio Tumeo, Politecnico di Milano

StarSs and OpenMP

StarSs and its implementations (SMPSs, CellSs) aim to provide a programming model characterised by its high productivity and performance portability. From a sequential code with a few annotations, StarSs runtimes generates at runtime a task dependency graph and exploits the inherent concurrency of the application. OpenMP v3.0 considers the tasks, but not data dependencies between them. BSC is actively participating in the **OpenMP standardization committee (ARB)** to include this and related ideas.

Further readings

Programming models group at BSC
http://www.bsc.es/plantillaG.php?cat_id=385

CUDA Zone
http://www.nvidia.co.uk/object/cuda_learn_uk.html