

# ERCIM “Alain Bensoussan” Fellowship Scientific Report

Fellow: **Ján Perháč**  
Visited Location : **NTNU – Norwegian University of Science and Technology**  
Duration of Visit: **12.4.2010 – 11.4.2011**

## I - Scientific activity

During my ERCIM Fellowship my primary research focus were these areas: *HPC* (High Performance Computing), *mobile and distributed GPGPU* (General Purpose GPU Computing) and partially also the *Cloud Computing*. I chose these three areas not only because they were primarily (GPGPU) or partially (HPC) part and interest of my PhD. studies and doctoral thesis, but for many well-known and world-wide respected researchers they represent one of the best possible option for reaching the next level of computational science, the exa-scale. It is well known that heterogeneity already is the key factor for all modern computer designs and architectures in the HPC area.

Although I was able to choose my primary hosting institute from a long list of well-known ERCIM research institutes, I chose NTNU or more precisely the HPC-Lab led by Dr. Anne C. Elster exactly because of the focus of her research group, laboratory and students as well. The High Performance / Heterogeneous and Parallel Computing Lab (HPC-Lab) is part of the Dept. of Computer and Information Science (IDI) at the Norwegian University of Science and Technology (NTNU) in Trondheim, Norway. The HPC-Lab has several custom-made multi-core high-end workstations with the latest-generation of GPGPU-enabled graphics cards donated by NVIDIA through Dr. Elster being part of their Professor Partnership Program. The HPC-Lab laboratory represents the absolute top-of-the-art laboratory when it comes to HPC and heterogeneous computing research. In addition the sum of the computational power all its desktop-class workstations is comparable with a smaller supercomputer, what was absolutely ideal for my research on the edge of possibilities of current state-of-the-art technologies.

Thanks to the amazing initiatives, activities and financial support from Dr. Elster and the HPC-Lab, the results of her group, as well as the results of my ERCIM Fellowship, were presented on several key global HPC-related conferences and exhibitions with the number of 10000+ attendees, including PARA 2010 in Reykjavik, Iceland, ISC 2010 in Hamburg and SC'2011 in New Orleans, the latter including a stay at her home in Austin, Texas and a visit to University of Texas at Austin. Dr. Elster included me as a participant at the NTNU stand at the latter two conferences.

One of the projects that I have worked during ERCIM Fellowship was the incremental work and development of the *Transparent Heterogeneous Open Resource (THOR) framework*, which is a joint work between NTNU, Universidad Complutense de Madrid, Spain and other affiliates. The THOR is a framework for providing seamless access to HPC systems composed of heterogeneous resources, including mobile devices. This framework is developed with the focus on the core module, in particular the policy engine and has been implemented on a

scaled-down heterogeneous cluster within a cloud-based computational environment. The THOR framework is an excellent way for demonstrating the interplay between many architecturally different technologies (CPUs, GPUs, embedded devices) in one heterogeneous environment driven by the cloud technology. The partial results of my contribution on the THOR framework were presented at the High-Performance Computing on Complex Environment Workshop (HPCCE 2010) which Dr. Elster co-organized with fellow Working Group leaders of EU COST Action IC0805: Open European Network for High Performance Computing on Complex Environments at the IEEE Cluster 2010 conference in Heraklion, Greece. This work was published as a joint paper in the IEEE Cluster 2010 proceedings. We also organized a Birds-of-Feather session on this topic at the Supercomputing (SC'10) conference in New Orleans, USA

I also had the opportunity to collaborate with the students from the HPC-Lab, all of which were supervised by Dr. Elster. The collaboration with the students was also a great way of exploring new options and possibilities related with porting codes for / from different platforms (CPUs to GPUs) and also with many other issues related with the GPGPU technology and heterogeneous computing as a whole.

From the international point of view, I had the opportunity to work on the THOR framework with visiting PostDoc Dr. José Luis Vázquez-Poletti from Universidad Complutense de Madrid, Spain and I also met graduate students Krystyna Napierala and Jarek Palczynski from Poznan, Poland that visited HPC-Lab in June 21-23, 2010.

During my fellowship I also reviewed several papers for the HPC-related HiPC 2010 conference.

## **II- Publication(s) during your fellowship**

VAZQUES POLETTI JOSE LUIS, PERHAC JAN, ELSTER ANNE CATHERINE, RYAN JOHN:

### **THOR: A Transparent Heterogeneous Open Resource Framework**

IEEE International Conference on Cluster Computing Workshops and Posters, 2010, 20-24 September 2010, Heraklion, Crete, GREECE,  
**ISBN 978-1-4244-8396-9**

*Abstract*—Heterogeneous computing which includes mixed architectures with multi-core CPUs as well as hardware accelerators such as GPU hardware, is needed to satisfy future computational needs and energy requirements. Cloud computing currently offers users whose computational needs vary greatly over time, a cost-effect way to gain access to resources. While the current form of cloud-based systems is suitable for many scenarios, their evolution into truly heterogeneous computational environments is still not fully developed.

This paper describes THOR (Transparent Heterogeneous Open Resources), our framework for providing seamless access to HPC systems composed of heterogeneous resources. Our work focuses on the core module, in particular the policy engine. To validate our approach, THOR has been implemented on a scaled-down heterogeneous cluster within a cloud-based computational environment. Our testing includes an Open CL encryption/decryption algorithm that was tested for several use cases. The corresponding computational benchmarks are provided to validate our approach and gain valuable knowledge for the policy database.

### **III -Attended Seminars, Workshops, and Conferences**

#### **SUPERCOMPUTING 2010**

International Conference for HPC, Networking, Storage and Analysis, November 13-19, 2010, New Orleans, Louisiana, USA

<http://sc10.supercomputing.org/>

#### **MEMICS 2010**

Workshop on Mathematical and Engineering Methods in Computer Science, October 22-24, 2010, Mikulov, Czech Republic

<http://www.memics.cz/2010>

#### **PARA 2010**

State of the Art in Scientific and Parallel Computing, June 6-9, 2010, Reykjavík, Iceland

<http://www.yourhost.is/para2010/home.html>

#### **INTERNATIONAL SUPERCOMPUTING CONFERENCE (ISC) 2010**

ISC – The HPC Event, May 30 – June 3, 2010 Hamburg, Germany

<http://www.supercomp.de/isc10/>

### **IV – Research Exchange Programme (12 month scheme)**

1) **CRCIM** - 18.10.2010 - 22.10.2010 – Czech Republic, exchange contact: prof. Václav Matyáš, PhD.

During my first ERCIM exchange programme (18.10.2010 - 22.10.2010) I visited the Czech Research Consortium for Informatics and Mathematics (CRCIM), which is also a member of ERCIM. The visited institute was Masaryk University in Brno, more specific The Laboratory for Parallel and Distributed Systems (ParaDiSe). During that week I had two talks in this laboratory, discussion with laboratory members about their current projects and the possibility of the use of GPGPU - enabled devices in their projects and I was invited to the MEMICS'10 conference, which was focused on Mathematical and Engineering Methods in Computer Science.

2) **INRIA** – 7.3.2011 - 11.3.2011 – France, exchange contact: Emmanuel Jeannot, Ph.D.

During my second ERCIM exchange programme (7.3.2011 - 11.3.2011) I visited INRIA (Bordeaux, France), which is also a member of ERCIM. The visited research team is called RUNTIME and it is intend to define, implement and validate a family of generic runtime systems that will provide a flexible and efficient foundation for building programming environments and applications in the intensive parallel computing field. The primary goal of my one-week exchange programme was to look into possibilities of linking the StarPU framework (developed by INRIA) with the THOR framework (developed by the NTNU). During that week I had one talk before the whole RUNTIME team and discussion with some RUNTIME team members about the heterogeneous environment as defined by the StarPU framework and about the possibilities of linking of these two frameworks.