HeteroPar' 2010

Eighth International Workshop on

Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms

---

August 30, 2010
Ischia-Naples, Italy

Frédéric Vivien
Program chair

Organized as part of

Euro-Par 2010

16th European Conference on

Parallel and Distributed Computing

HeteroPar' 2010 Technical Program

Monday, August 30, 2010

10h00 **Opening session**

10h10 **Keynote speaker**

Marco Danelutto
Dept. Computer Science, University of Pisa, Italy
Structured programming models targeting heterogeneous architectures

11h00 *Break*

11h30 **Session 1: General approaches, from compilation to execution middleware**

Tomasz Buchert, Lucas Nussbaum, and Jens Gustedt
Accurate emulation of CPU performance

Ruymán Reyes and Francisco de Sande
Case Studies in Automatic GPGPU code generation with `llc`

Muhammad Aleem, Radu Prodan, and Thomas Fahringer
On the Evaluation of JavaSymphony for Homogeneous and Heterogeneous Multi-core Clusters

13h00 *Lunch*

14h30 **Session 2: Iterative Algorithms**

Raphaël Couturier, David Laiymani, and Sébastien Miquée
MAHEVE: An Efficient Reliable Mapping of Asynchronous Iterative Applications on Volatile and Heterogeneous Environments.

David Clarke, Alexey Lastovetsky and Vladimir Rychkov
Dynamic Load Balancing of Parallel Computational Iterative Routines on Platforms with Memory Heterogeneity

15h30 **Session 3: Case studies**
Ignasi Barri, Josep Rius, Concepció Roig and Francesc Giné
Dealing with Heterogeneity for Mapping MMOFPS in Distributed Systems

16h00 *Break*

16h30 **Session 3 (continued): Case studies**

Brett Becker and Alexey Lastovetsky
Max-Plus Algebra and Discrete Event Simulation on Parallel Hierarchal Heterogeneous Platforms

17h00 **Closing remarks**

Keynote Speaker: Marco Danelutto

August 30, 10h10

Structured programming models targeting heterogeneous architectures

Abstract: More and more heterogeneous architectures are available with increasing computing capabilities. Heterogeneity comes from different factors, including different general purpose CPUs, GP-GPUs, dedicated co-processors as well as programmable FPGA subsystems. This poses new challenges to programmers, which are forced to learn and use a variety of programming frameworks, tools and methodologies to efficiently exploit the available heterogeneous hardware.
Structured programming frameworks represent a viable and efficient solution to support heterogeneous architecture targeting. We discuss how the structured programming approach may be used to address networks of heterogeneous processing elements and how the approach can be extended to target architectures including specialized co-processors. We show preliminary results and we outline the current open challenges.

About the speaker. Marco Danelutto is an Associate Professor at the Department of Computer Science, University of Pisa, since 1998. He received a Ph.D. in Computer Science in 1990. His main research interests are in structured parallel and distributed programming models (algorithmic skeletons, parallel design patterns, macro-data flow implementation models), autonomic computing, software components and semi formal methods and tools supporting parallel and distributed computing. He contributed to the development of GCM (the Grid Component Model) as leader of the Programming model Institute within the CoreGRID EU Network of Excellence and as responsible of the Non Functional Component Features work package in the GridCOMP EU STREP project. In early '90s he has been one of the main designers of P3L and he is currently maintaining Muskel, a full Java skeleton library targeting generic networks of Java enabled workstations. He is author and co-author of about 110 papers appearing in international journals and refereed conferences.

Session 1

General approaches: from compilation to execution middleware

August 30, 11h30 – 13h00

11h30 Tomasz Buchert, Lucas Nussbaum, and Jens Gustedt

Accurate emulation of CPU performance

Abstract – This paper addresses the question of CPU performance emulation, which allows experimenters to evaluate applications under a wide range of reproducible experimental conditions. Specifically, we propose Fracas, a CPU emulator that leverages the Linux Completely Fair Scheduler to achieve performance emulation of homogeneous or heterogeneous multi-core systems. Several benchmarks reproducing different types of workload (CPU-bound, IO-bound) are then used to thoroughly compare Fracas with another CPU emulator and hardware frequency scaling. We show that the design of Fracas results in a more accurate and a less intrusive CPU emulation solution.

12h00 Ruymán Reyes and Francisco de Sande

Case Studies in Automatic GPGPU code generation with `llc`

Abstract – The evolution of high performance computers is progressing toward increasingly heterogeneous systems. These new architectures pose new challenges, particularly in the field of programming languages. New tools and languages are needed if we want to make a full use of the advantages offered by these new architectures. `llc` is a language with a C-like syntax where parallelism is expressed using compiler directives. In this work we focus our attention on the new backend of our prototype compiler for `llc` which generates CUDA code. We evaluate the performance of the target code using three different applications. The preliminary results that we present make us believe that our approach is worth to be explored more deeply.

12h30 Muhammad Aleem, Radu Prodan, and Thomas Fahringer

On the Evaluation of JavaSymphony for Homogeneous and Heterogeneous Multi-core Clusters

Abstract – Programming hybrid heterogeneous multi-core cluster architectures is today an important topic in scientific and mainstream communities. To address this challenge, we developed JavaSymphony providing high-level programming abstraction and a middle-ware that facilitates the development and high-performance execution of Java applications on modern shared and distributed memory architectures. In this paper we present results of programming and executing a three-dimensional ray tracing application on several homogeneous and heterogeneous many-core cluster architectures.

Session 2

Iterative Algorithms

August 30, 14h30 – 15h30

14h30 Raphaël Couturier, David Laiymani, and Sébastien Miquée

MAHEVE: An Efficient Reliable Mapping of Asynchronous Iterative Applications on Volatile and Heterogeneous Environments

Abstract – With the emergence of massive distributed computing resources, such as grids and distributed clusters architectures, parallel programming is used to benefit from them and execute problems of larger sizes. The asynchronous iteration model, called AIAC, has been proven to be an efficient solution for heterogeneous and distributed architectures. An efficient mapping of applications' tasks is essential to reduce their execution time. In this paper we present a new mapping algorithm, called MAHEVE (Mapping Algorithm for HEterogeneous and Volatile Environments) which is efficient on such architectures and integrates a fault tolerance mechanism to resist computing nodes failures. Our experiments show gains on a typical AIAC application execution time of about 55%, executed on distributed clusters architectures containing more than 400 computing cores with the JaceP2P-V2 environment.

15h00 David Clarke, Alexey Lastovetsky and Vladimir Rychkov

Dynamic Load Balancing of Parallel Computational Iterative Routineson Platforms with Memory Heterogeneity

Abstract – Traditional load balancing algorithms for data-intensive iterative routines can successfully load balance relatively small problems. We demonstrate that they may fail for large problem sizes on computational clusters with memory heterogeneity. Traditional algorithms use too simplistic models of processors' performance which cannot reflect many aspects of heterogeneity. This paper presents a new dynamic load balancing algorithm based on the advanced functional performance model. The model consists of speed functions of problem size, which are built adaptively from a history of load measurements. Experimental results demonstrate that our algorithm can successfully balance data-intensive iterative routines on parallel platforms with memory heterogeneity.

Session 3

Case studies

August 30, 15h30 – 17h00

15h30 Ignasi Barri, Josep Rius, Concepció Roig and Francesc Giné

Dealing with Heterogeneity for Mapping MMOFPS in Distributed Systems

Abstract – In this paper, we present a distributed heterogeneous system called *OnDeGaS* (On Demand Game Service), that fits the scalability and latency requirements of MMOFPS networked games. To exploit platform capabilities efficiently, the *OnDeGaS* system performs a mapping mechanism that assigns the game sessions of a MMOFPS, taking advantage of the specific available computational resources of individual nodes. We show through simulation that this mapping mechanism is able to deal with different heterogeneity conditions in the distributed area. It allows the system to grow at any moment according to the existing demand, while latency values are maintained under the acceptable threshold permitted in MMOFPS games.

16h00 **Break**

16h30 Brett Becker and Alexey Lastovetsky

Max-Plus Algebra and Discrete Event Simulation on Parallel Hierarchal Heterogeneous Platforms

Abstract – In this paper we explore computing max-plus algebra operations and discrete event simulations on parallel hierarchal heterogeneous platforms. When performing such tasks on heterogeneous platforms parameters such as the total volume of communication and the top-level data partitioning strategy must be carefully taken into account. Choice of the partitioning strategy is shown to greatly affect the overall performance of these applications due to different volumes of inter-partition communication that various strategies impart on these operations. Max-plus algebra is regularly used in discrete event simulations and many other important computational applications thus warranting the exploration of and improvement upon the running times of basic max-plus operations on parallel platforms which are inherently hierarchal and heterogeneous in nature. The main goal of this paper is to present benefits waiting to be exploited by the use of max-plus algebra operations on these platforms and thus speeding up more complex and quite common computational topic areas such as discrete event simulation.