Chapter 1

Cluster on Simulation and Modeling Platform

Description of scientific issues. The goal of this cluster is to bring together the community around a common approach for modeling architectures. The motivation for this cluster is twofold: (1) modeling methodologies are facing significant challenges, especially due to the complexity of architectures and the emergence of multi-cores, (2) modeling/simulation tools are largely fragmented within the academic and industrial communities, making it exceedingly difficult to share, reuse or compare research or product achievements. Within this cluster, we will put an emphasis on simulation, like in HiPEAC1, but we also want to investigate analytical/statistical modeling, and FPGA prototyping, as two major alternative approaches.

Within HiPEAC1, we had kick-started this effort, then exclusively focused on simulation, which led to the development of the UNISIM simulation framework. Within HiPEAC2, our first goal is less to promote a single all-encompassing environment than to define methods for the different simulator efforts and tools to interoperate. The second goal of HiPEAC2 is to address a set of urgent research challenges. Some of the key challenges are:

- Simulation speed. It is the #1 challenge as architectures evolve to multi-cores with tens of cores. Numerous solutions are emerging, such as transaction-level modeling, sampling, parallel simulators, native execution, statistical simulation, FPGA prototyping, but they should be rapidly investigated.
- Technology. Numerous technology issues are now creeping up into system design and can no longer be ignored: besides power and area cost issues, they include wire delays, clock domains, temperature,...
- Application complexity. With the advent of smartphones used for 3D games and GPS navigation, embedded applications are getting almost as complex as desktop applications. Moreover, these smartphones are now running Linux or MacOS, so that full-system simulation is a must have, especially with multi-core architectures where the scheduler is playing an increasingly important role.
- Heterogeneous architectures. Embedded architectures and even the new general-purpose architectures are heterogeneous architectures. Modeling/Simulating such architectures raises special challenges. These challenges will be considered, especially through a tight cooperation with the Design cluster.

Summary of activities and main investigated approaches. During the first year, the cluster has focused on two issues: simulation/modeling speed and interoperability.

For *simulation speed*, there is no dominant approach and it is still very unclear which approach will ultimately prevail. So we focused on identifying all the potential approaches existing within the HiPEAC community. As a result, we had multiple presentations on the topic at the 1st and 3rd meetings. These presentations span most known approaches on the topic. The different approaches and talks are listed below.

- Fast functional and timing simulation, Nigel Topham, Edinburgh University, UK
- Hybrid simulation (toggling between native execution and timing simulation), Stefan Kraemer, Aachen University, Germany
- Sampling for multi-cores/MPSoC, Smail Niar, INRIA Lille, France
- Statistical simulation, benchmark generation, analytical modeling and workload characterization, Lieven Eeckhout, Ghent University, Belgium
- Fast analysis of program behavior on architectures, Erik Hagersten, Uppsala University/ACUMEM, Sweden
- Transaction-Level Modeling, Daniel Gracia-Perez, CEA, France
- Sampling for full-system simulation, Ayose Falcon, HP Labs Barcelona, Spain
- Hybrid FPGA/software simulation, and statistical sampling, Babak Falsafi, EPFL, Switzerland

For *interoperability*, the discussions have started at the 2nd meeting and are still under way. The general idea is to promote a *bottom-up* approach where, instead of imposing a single platform, we progressively create APIs allowing the different existing tools to connect together. For instance, if an existing full-system simulation supports a given power modeling API, any API-compliant power modeling tool can be plugged into that simulation. We believe this approach creates a win-win situation. First, no one has to drop his/her effort invested in developing or mastering the utilization of an existing platform or tool, everyone can keep their own tools. Second, anyone who has developed a simulator or tool can potentially have many complementary tools at his/her disposal. For instance, no single group can afford to develop a full-system simulator, a power model and sampling techniques to speed it up. However, using the APIs approach, the full-system simulator developer can leverage existing and complementary power models and simulation speed techniques. Third and conversely, the audience of any given simulator tool/platform becomes much larger because other researchers do not have to drop their own tools/simulators in order to leverage the works of others, they can combine it with their own work. In the next few years, we will evaluate and possibly implement this synergetic approach at building a common set of tools.

There were also several discussions and presentations revolving around *full-system simulation* which is one of the key aforementioned issues:

- Full-system multi-core simulation, Babak Falsafi, EPFL, Switzerland
- A full-system simulator based on the Cell processor, Alex Ramirez, UPC/BSC, Spain (no slide, presentation of SARCSim)
- Virtual platforms, Daniel Gracia-Perez, CEA, France

Defining needs. At the 2nd meeting, a lively discussion erupted on the necessity to clearly define needs before selecting any speed or interoperability approach. While HiPEAC1 was very much focused on micro-architectures, HiPEAC2 spans more deeply into the MPSoC domains with partners like Aachen University, and in data centers with partners like HP Labs Barcelona. As a result, there are multiple different use cases for simulation tools, each with different requirements. Consequently, it was decided to first identify and list these different use cases, and to progressively draw conclusions from these use cases on which approaches we should preferentially push within HiPEAC2.

We identified three applications scenarios so far, which are summarized on the HiPEAC web site.

RAPIDO Workshop

In order to achieve further cross-fertilization on the key topic of simulation speed, the HiPEAC cluster on Simulation and the cluster on Design have together supported the RAPIDO workshop on the topic. This workshop was launched at the initiative of Smail Niar, University of Valenciennes/INRIA, France, and the

first edition will take place in the form of a 1-day workshop at the HiPEAC 2009 conference in Cyprus. The workshop will include 6 invited presentations and 9 selected presentations.

Future plans. The future plans have been largely mentioned before and briefly summarized here:

- Finalize the use case discussion and draw conclusions for the cluster orientations.
- Attempt to implement, possibly through a number of experiments, the interoperability approach.
- From an organizational standpoint, now privilege a few long presentations of tools and simulators, allowing participants to understand how they could connect their own tools and how APIs should be designed.
- Further confront the different simulation speed approaches.