

MareIncognito: A Perspective towards Exascale

Jesus Labarta

Barcelona Supercomputing Centre, Spain

Abstract. MareIncognito is a cooperative project between IBM and the Barcelona Supercomputing Center (BSC) targeting the design of relevant technologies on the way towards exascale. The initial challenge of the project was to study the potential design of a system based on a next generation of Cell processors. Even so, the approaches pursued are general purpose, applicable to a wide range of accelerator and homogeneous multicores and holistically addressing a large number of components relevant in the design of such systems.

The programming model is probably the most important issue when facing the multicore era. We need to offer support for asynchronous data flow execution and decouple the way source code looks like and the way the program is executed and its operations (tasks) scheduled. In order to ensure a reasonable migration path for programmers the execution model should be exposed to them through a syntactical and semantic structure that is not very far away from current practice. We are developing the StarSs programming model which we think addresses some the challenges of targeting the future heterogeneous / hierarchical multicore systems at the node level. It also integrates nicely into coarser level programming models such as MPI and what is more important in ways that propagate the asynchronous dataflow execution to the whole application. We are also investigating how some of the features of StarSs can be integrated in OpenMP.

At the architecture level, interconnect and memory subsystem are two key components. We are studying in detail the behavior of current interconnect systems and in particular contention issues. The question is to investigate better ways to use the raw bandwidth that we already have in our systems and can expect to grow in the future. Better understanding of the interactions between the raw transport mechanisms, the communication protocols and synchronization behavior of applications should lead to avoid an exploding need for bandwidth that is often claimed. The use of the asynchronous execution model that StarSs offers can help in this direction as a very high overlap between communication and computation should be possible. A similar effect or reducing sensitivity to latency as well as the actual off chip bandwidth required should be supported by the StarSs model.

The talk will present how we target the above issues, with special details on the StarSs programming model and the underlying idea of the project of how tight cooperation between architecture, run time, programming model, resource management and application are needed in order to achieve in the future the exascale performance.