SEMINAIRE EXCEPTIONNEL

Lundi 30 novembre 2015 de 14h à 16h

Salle Renaudeau, Bâtiment Laplace

ENS Cachan Plan d'accès

Avancées sur les modèles de performance pour les nouvelles architectures HPC

Aleksandar ILIC

Professeur à l'Université de Lisbonne et Professeur invité à l'ENS Cachan

Cache-Aware Roofline Model: Performance, Power and Energy-Efficiency

Philippe Thierry

INTEL, Oil& Gas Engineering Manager

High Level Methodologies for Performance Characterization and Prediction

Abstracts page suivante

Cache-Aware Roofline Model: Performance, Power and Energy-Efficiency

As architectures evolve towards more complex multi-core designs, deciding what optimizations provide the best tradeoff between performance and efficiency is becoming a prominent issue. To help in this decision process, a set of fundamental models will be presented in this talk, which allow characterizing the upper-bounds for performance, power, energy and energy-efficiency of multi-core architectures. These models evaluate how key micro-architectural aspects, such as accessing different functional units or different memory hierarchy levels, affect the attainable performance, power and efficiency of the processor (by also considering different power domains).

Based on these concepts, a set of total models will also be presented, which allow describing the realistically attainable limits as sustained by the architecture, namely, performance, energy and energy-efficiency ranges, as well as a complete power envelope for a single frequency. These models are also applied across a range of operational frequencies for an Intel processor to derive a set of supported performance and energy-efficiency ranges and different power envelopes.

The proposed models are rigorously validated on 5 Intel processors supported on different micro-architectures by relying on hardware counters and especially developed highly accurate performance/power monitoring tools. The experimental results show a very high accuracy of the proposed models, and their ability to provide more intuitive and useful guidelines than the state-of-the-art approaches, when characterizing a set of synthetic and 38 standard benchmarks from SPEC CPU2006, PARSEC, SPLASH and MKL.

High Level Methodologies for Performance Characterization and Prediction

Designing a supercomputer to satisfy the needs of future applications and workloads within a given power envelope is considering the rapidly evolving high technology environment not an easy task.

In this frame, the prediction of performance can be used for many different needs from designing a new micro architecture or memory hierarchy to defining the interconnection and storage of the future.

Several tools already exist for analyzing the different aspects of application characterization and performance prediction.

They have, however, so far rarely been connected due to their different precisions and resolutions.

Based on a first approximation of the application behavior mostly involving memory bandwidth and

floating point (FP) demands we can demonstrate that realistic performance predictions can be easily obtained at the application level for single and multiple node configuration.