Escuela Ingeniería Informática Valladolid
UVa

Charlas de los Miércoles

Curso 2017-18

Cache-Aware Roofline Model: Performance, Power and Energy-Efficiency Limits of Multi-cores

Prof. Dr. Leonel Sousa

Technical University Lisbon, INESC-ID (Portugal)

inesc id lisboa

INSTITUTO SUPERIOR TÉCNICO
Universidade Técnica de Lisboa

In this talk, a set of fundamental Cache-aware Roofline Models (CARMs) are presented, which describe the performance, power, energy and energy-efficiency upper-bounds of parallel architectures, namely multi-core CPU and GPU 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 energy-efficiency. The proposed models are validated on several CPU and GPU architectures with a very high accuracy, achieved by relying on hardware counters and a set of custom performance/power monitoring tools. Besides, the experimental results show the ability of the proposed CARMs to provide more intuitive and useful guidelines than the state-of-the-art approaches, especially when optimizing and characterizing real-world applications from different benchmark suites. The performance CARM was also recently integrated by Intel as a fully supported feature into the Intel Advisor, and it is described as "an incredibly useful diagnosis tool (...) that developers can use to guide them (in the application optimization process), ensuring that they can squeeze the maximum performance out of their code with minimal time and effort''.

9 de mayo de 2018, 12h
Salón de Grados
Escuela de Ingeniería Informática

ESDUVa

acm

La asistencia a estas charlas tiene reconocimiento académico de créditos por actividades culturales.

ETSI Informática
@ETSIIUVa

UVa