Published on *BSC-CNS* (https://www.bsc.es)

Opening Session of the Severo Ochoa Research Seminars

Jeffrey S. Vetter, Oak Ridge National Laboratory and Georgia Institute of Technology, offered the talk entitled "Exploring Emerging Technologies for Extreme Scale HPC Architectures".

Severo Ochoa Excellence Programme
Research Seminar Lectures at BSC

EXCELENCIA
SEVERO
OCHOA

The Opening session of the 3rd season of Severo Ochoa Research Seminar Lecture Series (SORS) on Thursday, **10 September, at 11:00** in **Sala d'Actes FIB** was opened with welcome note by **Josep Casanovas**. The speaker was presented by **Jesus Labarta**.

Speaker: Jeffrey S. Vetter, Oak Ridge National Laboratory and Georgia Institute of Technology, http://ft.ornl.gov/~vetter/

Title: Exploring Emerging Technologies for Extreme Scale HPC Architectures

Date: Thursday, 10 September- 11:00

Venue: Sala d'Actes FIB

Abstract:

While architectures and programming models have remained relatively stable for almost two decades, new architectural features, such as heterogeneous processing, nonvolatile memory, and optical interconnection networks, will demand that software systems and applications be redesigned so that they expose massive amounts of hierarchical parallelism, carefully orchestrate data movement, and balance concerns over performance, power, and resiliency. This instability has led to two inevitable problems: decreased programmer productivity, and difficult performance prediction. In this talk, I will describe two solutions to these problems, respectively: our OpenARC compiler and runtime system, and our Aspen performance modeling language. First, OpenARC is a research compiler that supports OpenACC and OpenMP4, and can generate code in CUDA, OpenCL, and LLVM IR. OpenARC has enabled us to explore how to enable performance portability of applications across diverse architectures. Second, Aspen is a domain specific language for structured analytical modeling of applications and architectures. It is designed to enable rapid exploration of new algorithm and architectures. Once created, Aspen models can then be used for a variety of purposes including predicting performance of future applications, evaluating system architectures, informing runtime scheduling decisions, and identifying system anomalies.

Bio:

Jeffrey Vetter, Ph.D., holds a joint appointment between Oak Ridge National Laboratory (ORNL) and the Georgia Institute of Technology (GT). At ORNL, Vetter is a Distinguished R&D Staff Member, and the founding group leader of the Future Technologies Group in the Computer Science and Mathematics Division. At GT, Vetter is a Joint Professor in the Computational Science and Engineering School, and the Director of the NVIDIA CUDA Center of Excellence. His work has won awards at major conferences: Gordon Bell Prize at SC10, Best Paper Awards at the International Parallel and Distributed Processing Symposium (IPDPS) and EuroPar, Best Student Paper Finalist at SC14, and Best Presentation at EASC 2015. His recent books "Contemporary High Performance Computing (Vols. 1 and 2)" survey the international landscape of HPC. See his website for more information: http://ft.ornl.gov/~vetter/.

www.bsc.es/marenostrum-support-services/hpc-education-and-training/severo-ochoa-research-seminar/exploring

INFORMÀTICA DE BARCELONA


Jeffrey S. Vetter
Barcelona Supercomputing Center
Barcelona
10 Sep 2015
OAK RIDGE NATIONAL LABORATORY
UNIVERSITAT POLITÈ


Exploring Emerging Technologies in the Extreme Scale HPC Co-Design Space
Jeffrey S. Vetter
Barcelona Supercomputing Center
Barcelona
Georgia Tech
http://ft.ornl.gov
vetter@computer.org
FIB
FACULTAT D'INFORMÀTICA DE BARCELONA
UNIVERSITAT POLITÈCNICA DE


Overview
Our community has major challenges in HPC as we move to extreme scale
Power, Performance, Resilience, Productivity
New technologies emerging to address some of these challenges
Heterogeneous computing
Multimode memory systems including nonvolatile memory
Not just HPC: Most uncertainty in at least two decades
Exascale includes even more diversity and uncertainty
We need performance prediction and portability tools now more than ever!
Aspen is a tool for structured design and analysis
Co-design applications and architectures for performance, power, resiliency
Automatic model generation
Scalable to distributed scientific workflows
FIB Alumni
www.fibalumni.net

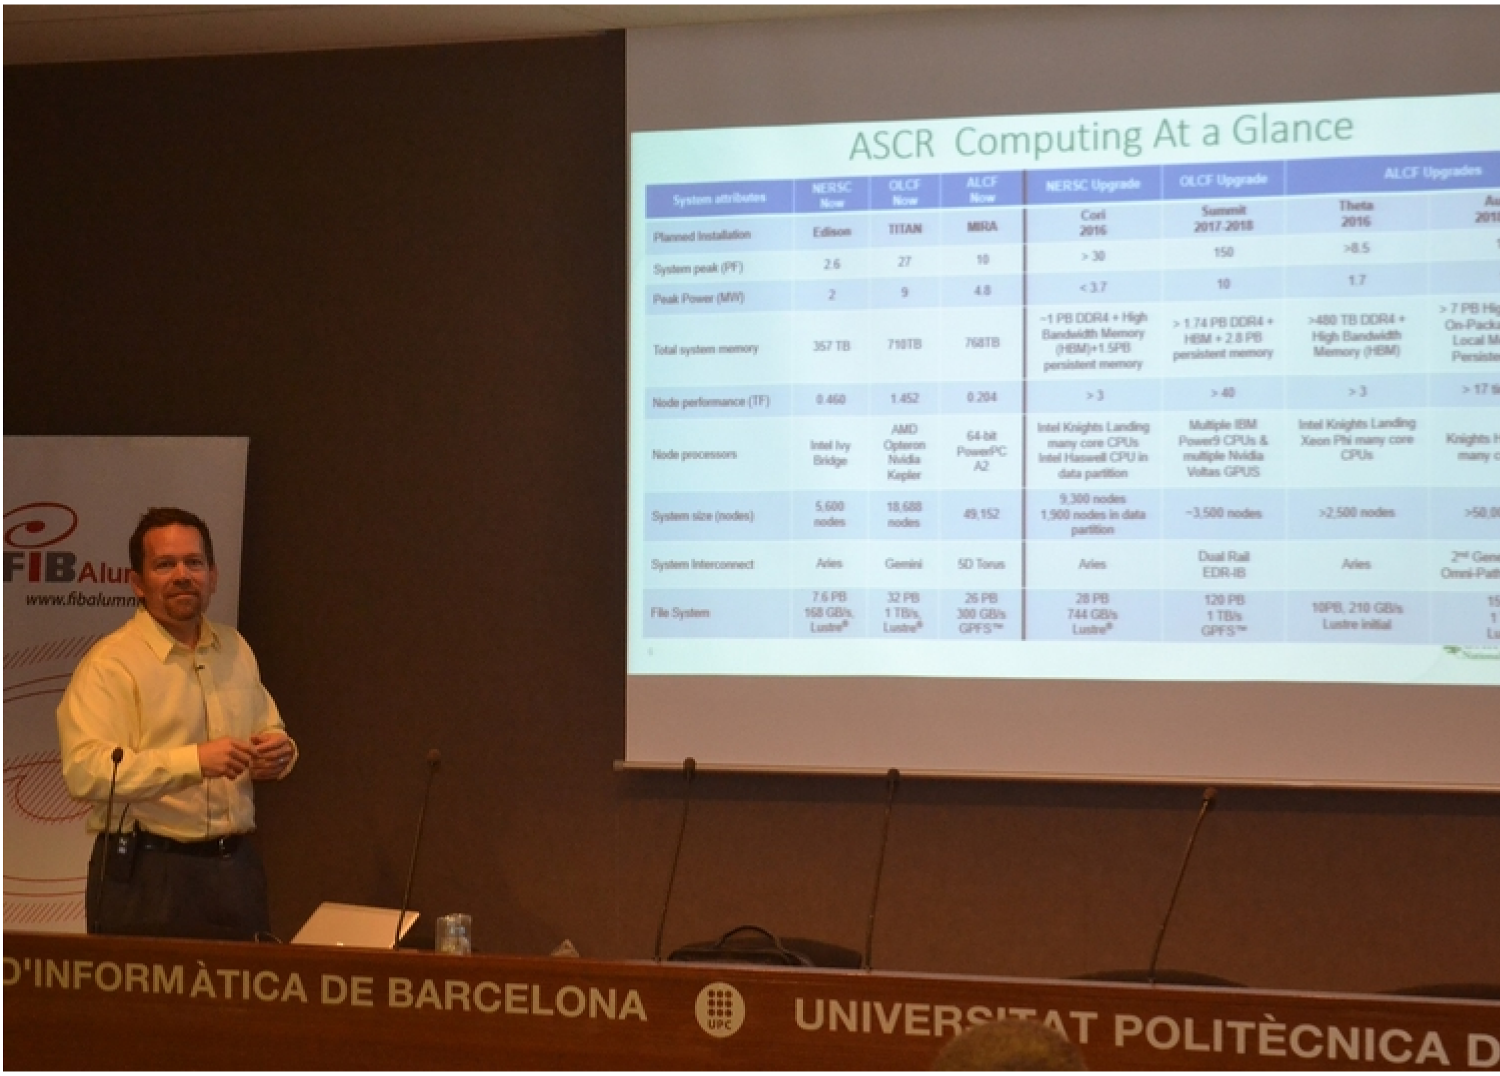

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

---

Source URL (retrieved on *19 Mar 2025 - 23:00*): https://www.bsc.es/es/news/bsc-news/opening-session-the-severo-ochoa-research-seminars