

# Third International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2 '10)

# Organizing Committee

- Steering Committee



D. K. Panda  
Ohio State University



William Gropp  
University of Illinois,  
Urbana Champaign



Vijay Saraswat  
IBM Research

- Program Co-chairs



Pavan Balaji  
Argonne National Laboratory



Abhinav Vishnu  
Pacific Northwest National Laboratory

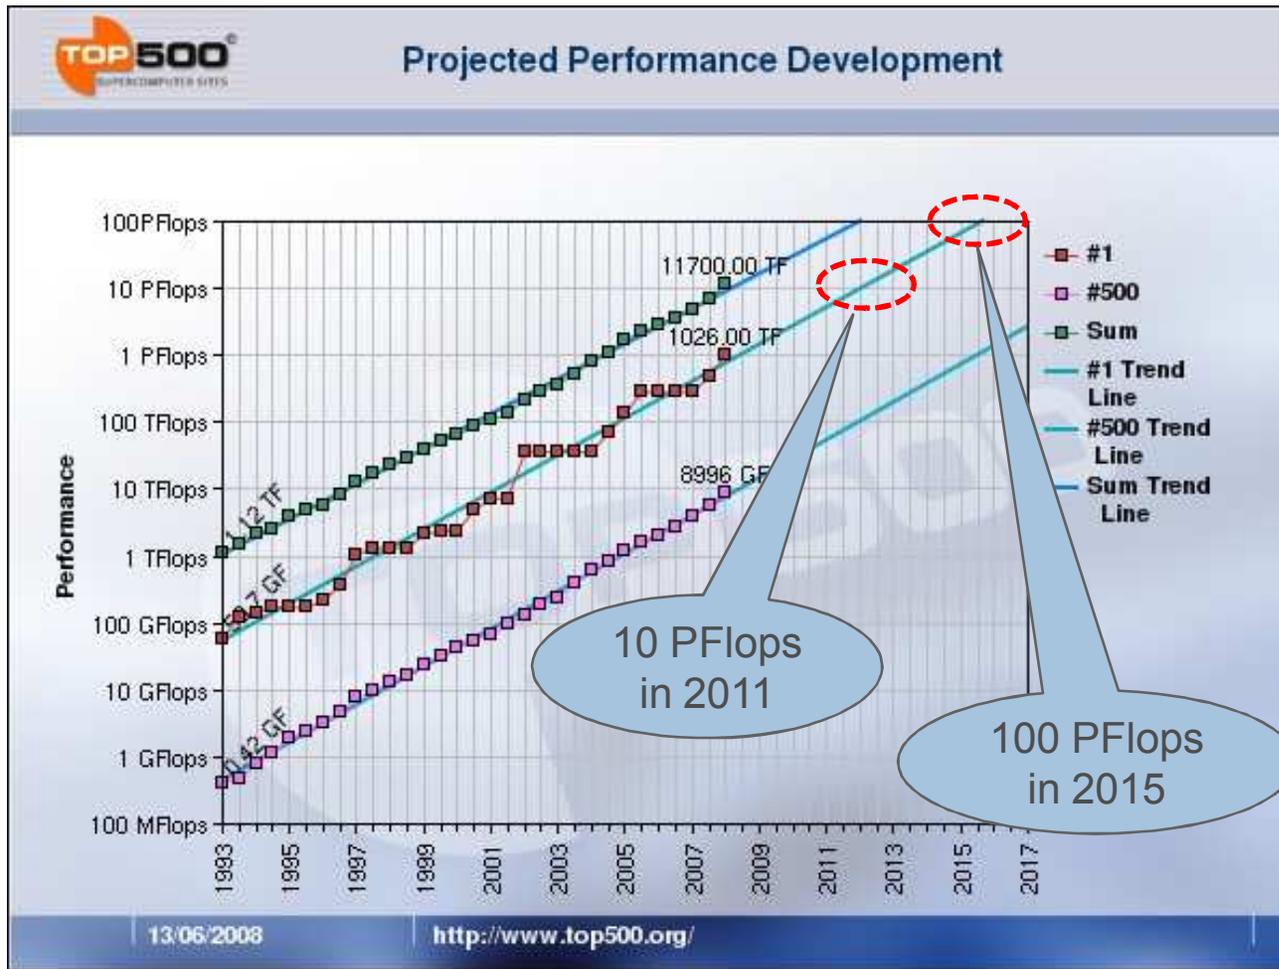


# Workshop Goals

- Primary Goal: Consolidating ideas for programming and managing high-end computing systems
- Large machines almost always mean complex architectures
  - Sometimes it means a larger count of existing machines
  - Oftentimes it means disruptive changes to our current understanding of system architectures
- We hope to serve as a venue for researchers to exchange ideas in how the community can build and use the world's largest systems!



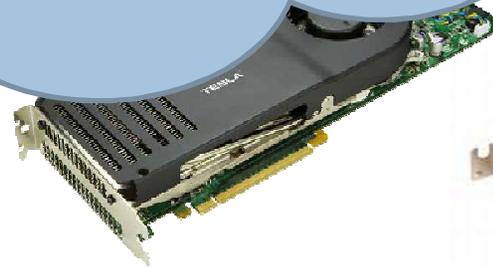
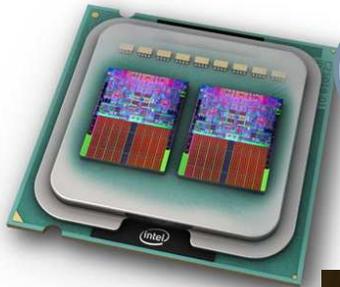
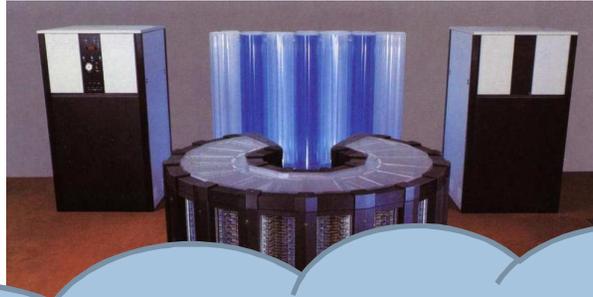
# Ready or Not, Here it Comes!



*Expected to have an Exaflop system in 2018-2019 !*



# Systems Are Already Very Complex



Need to Convert  
System Capacity  
to  
System Capability



# Challenges for Massive Systems

- Scalability
  - Systems with hundreds of thousands of cores exist
    - We will have a million cores soon
  - Small performance problems can get brutal at scale !
- Ease of Use (The “Wizard Gap”)
  - Not everyone is a parallel programming expert
  - ... and yet, everyone wants the best performance
- We live in a heterogeneous jungle
  - Multicore processors, Integrated Accelerator cores, GPUs, FPGAs
- Reliability and Maintainability
  - Failure detection and transparency
  - Ease of administration



# Workshop Areas of Interest

- Parallel Programming Models and their associated Systems Software as applicable to high-end computing systems
- Questions that we are pursuing:
  - How should the largest machines in the world be programmed?
  - What system software is necessary to program and manage these massive systems efficiently and productively?
  - What tools do we need to manage such massive systems?



# Program Committee: 23 members

- **Ahmad Afsahi**, Queen's University
- **George Almasi**, IBM Research
- **Taisuke Boku**, Tsukuba University
- **Ron Brightwell**, Sandia National Laboratory
- **Franck Cappello**, INRIA, France
- **Yong Chen**, Oak Ridge National Laboratory
- **Ada Gavrilovska**, Georgia Tech
- **Torsten Hoefler**, Indiana University
- **Zhiyi Huang**, University of Otago, New Zealand
- **Hyun-wook Jin**, Konkuk University, South Korea
- **Zhiling Lan**, Illinois Institute of Technology
- **Doug Lea**, State University of New York at Oswego
- **Heshan Lin**, Virginia Tech
- **Jiuxing Liu**, IBM Research
- **Guillaume Mercier**, INRIA, France
- **Scott Pakin**, Los Alamos National Laboratory
- **Fabrizio Petrini**, IBM Research
- **Bronis de Supinski**, Lawrence Livermore National Laboratory
- **Sayantana Sur**, Ohio State University
- **Rajeev Thakur**, Argonne National Laboratory
- **Vinod Tipparaju**, Oak Ridge National Laboratory
- **Jesper Traeff**, University of Vienna, Europe
- **Weikuan Yu**, Auburn University

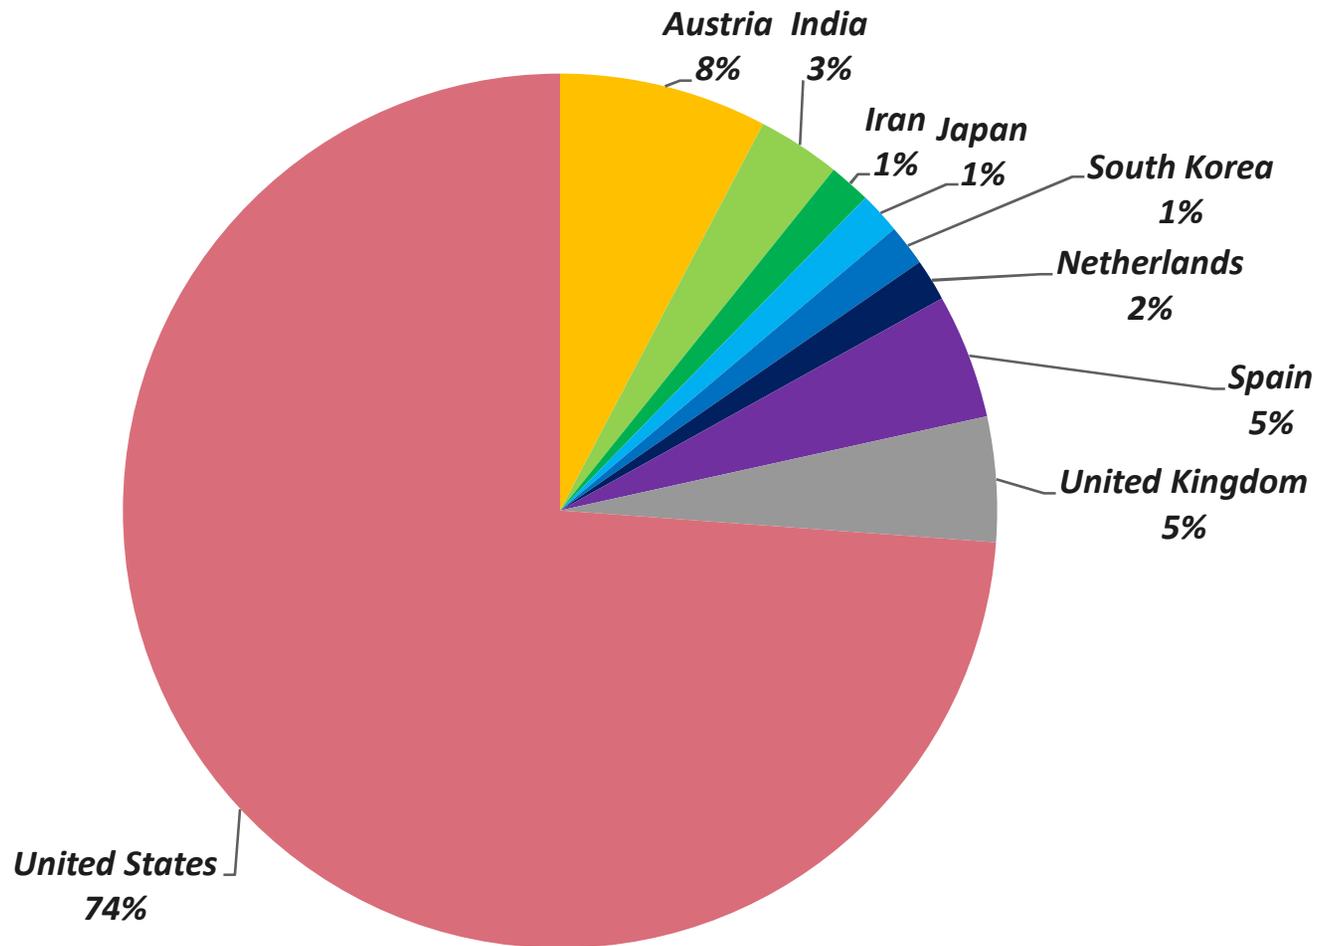


# Paper Reviews

- Totally 20 papers were submitted
  - 4-5 reviews per paper
  - Average of 4.75 reviews per paper
  - Program Committee telemeeting to finalize accepted papers
  - 10 papers were accepted
- Some of the accepted papers have been invited to submit to a special issue of the International Journal of High Performance Computing Applications (IJHPCA) on “Programming Models, Software and Tools for High-End Computing”



# Author Attraction per Country



# Logistics

- Four sessions:
  - Three sessions for regular technical paper presentations
  - One session for a technical panel discussion
- If you are a presenter:
  - Please meet your session chair before your session starts
  - Bios for each presenter needed
  - Session chairs might require presentations to be copied to a common presentation laptop before the session
- Wireless Connectivity
  - SSID: hhonor
  - Passcode: SDGO-5300 (first one is a 'Oh', last two are 'zeroes')



# P2S2 Sessions

## ■ Session1: Communication

[Time: 9:00 am to 10:30 am]

Session Chair: Vinod Tipparaju

- *"Efficient Zero-Copy Noncontiguous I/O for Globus on InfiniBand"*, Weikuan Yu and Jeffrey Vetter
- *"Scaling Linear Algebra Kernels using Remote Memory Access"*, Manojkumar Krishnan, Robert Lewis and Abhinav Vishnu
- *"High Performance Design and Implementation of Nemesis Communication Layer for Two-sided and One-Sided MPI Semantics in MVAPICH2"*, Miao Luo, Sreeram Potluri, Ping Lai, Emilio P. Mancini, Hari Subramoni, Krishna Kandalla, Sayantan Sur and Dhabaleswar K. Panda

## ■ Session 2: Panel: "Is Hybrid Programming a Bad Idea Whose Time has Come?"

[Time: 11:00 am to 12:30 pm]

Moderator: Pavan Balaji

- *Taisuke Boku, Tsukuba University, Japan*
- *Allen Maloney, University of Oregon*
- *Bronis de Supinski, Lawrence Livermore National Laboratory*
- *Vinod Tipparaju, Oak Ridge National Laboratory*



# P2S2 Sessions (contd.)

## ■ Session3: Programming Models and Performance Evaluation

[Time: 1:30 pm to 3:30 pm]

Session Chair: Hui Jin, Illinois Institute of Technology

- *"Performance Modeling for AMD GPUs", Ryan Taylor and Xiaoming Li*
- *"A Hybrid Programming Model for Compressible Gas Dynamics using OpenCL", Ben Bergen, Marcus Daniels and Paul Weber*
- *"Message Driven Programming with S-Net: Methodology and Performance", Frank Penczek, Sven-Bodo Scholz, Alex Shafarenko, Chun-Yi Chen, Nader Bagherzadeh, Clemens Grellck and Jungsook Yang*
- *"Implementation and Performance Evaluation of XcalableMP: A Parallel Programming Language for Distributed Memory Systems", Jinpil Lee and Mitsuhsa Sato*

## ■ Session 4: Scheduling and Cache Management

[Time: 4:00 pm to 5:30 pm]

Moderator: Weikuan Yu, Auburn University

- *"Scheduling a ~100,000 core Supercomputer for maximum utilization and capability", Phil Andrews, Patricia Kovatch, Victor Hazlewood and Troy Baer*
- *"Improving the Effectiveness of Context-based Prefetching with Multi-order Analysis", Yong Chen, Huaiyu Zhu, Hui Jin and Xian-He Sun*
- *"Hierarchical Load Balancing for Large Scale Supercomputers", Gengbin Zheng, Esteban Meneses, Abhinav Bhatele and Laxmikant V. Kale*



# Plans for P2S2 2011

- Steering Committee:
  - D. K. Panda (Ohio State University)
  - William Gropp (University of Illinois, Urbana Champaign)
  - Vijay Saraswat (IBM Research)
- Program Co-chairs:
  - Pavan Balaji (Argonne National Laboratory)
  - Abhinav Vishnu (Pacific Northwest National Laboratory)
  - Yong Chen (Oak Ridge National Laboratory)
- Host conference to be decided (maybe ICPP 2011)
- Like this year, we again plan to have special issue with a journal
  - Possibly the Journal of High Performance Computing Applications (JHPCA)
- Get Involved !
  - Paper submissions and workshop attendance
  - Advertizing

