

# 2015 International Workshop On Data-Intensive Scalable Computing Systems

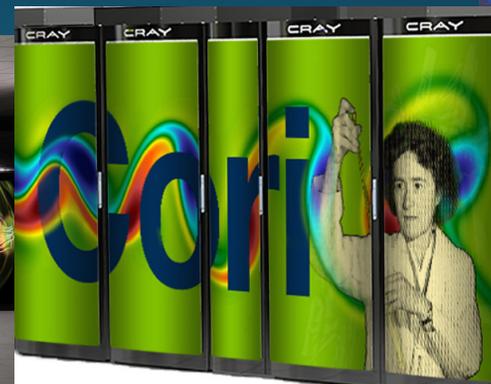
Austin, Texas  
15 November 2015

Held in conjunction with SC15  
Held in cooperation with ACM SIGHPC



# Welcome!

- Forum for researchers and developers to discuss recent results and future challenges of data intensive applications on traditional HPC systems and latest data-centric computing systems
- DISCS-2015 is fourth in a series
- Has been held in conjunction with SC since 2012



# Selecting our Program

- Accepted 9 out of 15 submissions
- All submissions received at least three reviews
- Online discussion among program committee to discuss submissions and reviewer ratings
- Accepted papers distributed into three sessions

# Morning Program

- Opening remarks, 8:45-9:00
- Keynote Address, 9:00-10:00  
*Data Intensive Research Infrastructure to Advance Scientific Discovery*  
Irene Qualters, National Science Foundation
- Morning Break, 10:00-10:30
- Session 1: New Paradigms for Data-Intensive Processing
  - Session chair: Weikuan Yu
  - 10:30-11:00: *SJM: an SCM-based journaling Mechanism with Write Reduction for File Systems*
  - 11:00-11:30: *Efficient Disk-to-Disk Sorting: A Case Study in the Decoupled Execution Paradigm*
  - 11:30-12:00: *A Low-cost Adaptive Data Separation Method for the Flash Translation Layer of Solid State Drives*
  - 12:00-12:30: *Big Data Analytics on Traditional HPC Infrastructure Using Two-Level Storage*
- Lunch Break, 12:30-2:00

# Afternoon Program

- Session 2: Parallel I/O Acceleration
  - Session chair: Michael Kluge
  - 2:00-2:30: *Route-aware Independent MPI I/O on the Blue Gene/Q*
  - 2:30-3:00: *Experimental Evaluation of a Flexible I/O Architecture for Accelerating Workflow Engines in Cloud Environments*
- Afternoon Break, 3:00-3:30
- Session 3: Analytics Frameworks
  - Session chair: Sarp Oral
  - 3:30-4:00: *A Case Study of MapReduce Speculation for Failure Recovery*
  - 4:00-4:30: *Supporting Online Analytics with User-Defined Estimation and Early Termination in a MapReduce-Like Framework*
  - 4:30-5:00: *Performance Evaluation and Tuning of BioPig for Genomic Analysis*
- Closing remarks, 5:00-5:15

# Instructions to Presenters

- 25 minute presentations, 5 minute question and answer period
- Please check in with your session chair at least 5 minutes before the session starts
- Please send your slides to [discs2015@easychair.org](mailto:discs2015@easychair.org) by 11/18 for posting to the workshop web site

# Journal Special Issue

- Special issue of Parallel Computing: Systems & Applications
- Authors of accepted papers are encouraged to submit extended versions of their workshop paper
- Instructions will be sent out by mid-December 2015
- Target submission deadline early Spring 2016

# Organizers

- Steering Committee

- Yong Chen (chair), Texas Tech University
- William D. Gropp, University of Illinois at Urbana-Champaign
- Xian-He Sun, Illinois Institute of Technology
- Rajeev Thakur, Argonne National Laboratory

- General Chair

- Philip C. Roth, Oak Ridge National Laboratory

- Program Committee Chairs

- Weikuan Yu, Florida State University
- Shane Canon, Lawrence Berkeley National Laboratory

- Proceedings Chair

- Michael Kluge, TU-Dresden

- Publicity Chair

- Jianhui Yue, Auburn University

# Program Committee

- Michael J. Brim, Oak Ridge National Laboratory
- Suren Byna, Lawrence Berkeley National Laboratory
- Pietro Cicotti, San Diego Supercomputer Center
- Shuibing He, Illinois Institute of Technology
- Hyun-Wook Jin, Konkuk University
- Mitch Horton, University of Tennessee, Knoxville
- Dries Kimpe, KCG Holdings, Inc.
- Michael Kluge, TU Dresden
- Quincey Koziol, The HDF Group
- Joshua Ladd, Mellanox
- Michael Lang, Los Alamos National Laboratory
- John Leidel, Texas Tech University
- Zhuo Liu, Yahoo
- Jay Lofstead, Sandia National Laboratories
- Sarp Oral, Oak Ridge National Laboratory
- Balaji Palanisamy, University of Pittsburgh
- Seung Woo Son, University of Massachusetts Lowell
- Jian Tan, Alibaba
- Wei Tang, Argonne National Laboratory
- Zhiqi Tao, Intel Corporation
- Douglas Thain, University of Notre Dame
- Yandong Wang, IBM
- Bin Wang, Auburn University
- Belle Xiang, Huawei
- Weijun Xiao, Virginia Commonwealth University
- Pengcheng Xiong, Horton Networks
- Cong Xu, Intel
- Peixiang Zhao, Florida State University
- Fang Zheng, IBM

# Enjoy the Workshop!

- Proceedings already available at <http://dl.acm.org/citation.cfm?id=2831244>
- Have feedback? Find us during a break, during closing remarks, or send email to [discs2015@easychair.org](mailto:discs2015@easychair.org)