

Preface

Today, huge amounts of data are being collected in many areas, which create new opportunities to understand phenomena in meteorology, health, finance, and many other sectors. Big Data is considered precious assets of companies, organizations, and even nations. Turning such big data into real treasures requires the support of big data systems and platforms. However, the sheer volume of big data requires significant storage capacity, transmission bandwidth, computation, and power consumption. It is expected that systems with unprecedented scales can resolve the problems caused by varieties of big data with daunting volumes.

The complexity, diversity, frequently changed workloads, and rapid evolution of big data systems raise great challenges in big data benchmarking. Without big data benchmarks, it is very difficult for big data owners to make a decision on which system is best for meeting with their specific requirements. They also face challenges on how to optimize the systems and their solutions for specific or even comprehensive workloads. Meanwhile, researchers are also working on innovative data management systems, hardware architectures, operating systems, and programming systems to improve performance in dealing with big data.

This book includes papers from two workshops, which are the fourth and fifth workshops on Big Data Benchmarks, Performance Optimization, and Emerging Hardware (BPOE-4 and BPOE-5). BPOE-4 (http://prof.ict.ac.cn/bpoe_4_asplos/) is co-located with ASPLOS 2014 (<http://www.cs.utah.edu/asplos14/>), a premier conference on architecture support for operating systems and programming systems. BPOE-5 (http://prof.ict.ac.cn/bpoe_5_vldb/) is co-located with VLDB 2014 (<http://www.vldb.org/2014/>), a premier conference on data management, database and information systems. Both workshops focus on architecture and system support for big data systems, aiming at bringing researchers and practitioners from data management, architecture, and systems research communities together to discuss the research issues at the intersection of these areas.

The call for papers for these two workshops attracted a number of high-quality international submissions. Within a rigorous process, in which each paper was reviewed by at least four experts, we selected 6 papers out of 12 submissions for inclusion in the BPOE-04 and 10 papers out of 18 submissions in the BPOE-05, respectively. In addition, several prestigious keynote speakers were invited, including Prof. Lizy Kurian John at University of Texas at Austin (<http://users.ece.utexas.edu/~ljohn/>) whose topic was “Big Data Workloads: An Architect’s Perspective,” Prof. Dhableswar K. (DK) Panda at Ohio State University (<http://www.cse.ohio-state.edu/~panda/>) whose topic was “Accelerating Big Data Processing with RDMA-Enhanced Apache Hadoop,” Prof. Christos Kozyrakis at Stanford University (<http://csl.stanford.edu/~christos/>) whose topic was “Resource Efficient Cloud Computing,” and Dr. Jeff Stuecheli from IBM (<http://www.linkedin.com/pub/jeff-stuecheli/2/664/a0a>) whose topic was “Power Technology For a Smarter Future.”

We are very grateful to the efforts of all authors related to writing, revising, and presenting their papers at BPOE workshops. Finally, we appreciate the indispensable support of BPOE Program Committees and thank their efforts and contributions in maintaining the high standards of the BPOE workshop.

August 2014

Jianfeng Zhan
Rui Han
Chuliang Weng

Big Data Benchmarks, Performance Optimization, and
Emerging Hardware

4th and 5th Workshops, BPOE 2014, Salt Lake City,
USA, March 1, 2014 and Hangzhou, China, September
5, 2014, Revised Selected Papers

Zhan, J.; Han, R.; Weng, C. (Eds.)

2014, X, 221 p. 100 illus., Softcover

ISBN: 978-3-319-13020-0