

	Poster title ( <i>Authors</i> )
11	<b>Orbe: Scalable Causal Consistency Using Dependency Matrices and Physical Clocks</b> <i>Jiaqing Du (EPFL), Sameh Elnikety (Microsoft Research), and Amitabha Roy and Willy Zwaenepoel (EPFL)</i>
12	<b>Leveraging Sharding in the Design of Scalable Replication Protocols</b> <i>Hussam Abu-Libdeh and Robbert van Renesse (Cornell University) and Ymir Vigfusson (Reykjavik University)</i>
13	<b>Understanding and Mitigating the Impact of Load Imbalance in the Memory Caching Tier</b> <i>Yu-Ju Hong and Mithuna Thottethodi (Purdue University)</i>
14	<b>Limplock: Understanding the Impact of Limpware on Scale-Out Cloud Systems</b> <i>Thanh Do (University of Wisconsin-Madison) and Mingzhe Hao, Tanakorn Leesatapornwongsa, Tirat Patana-anake, and Haryadi S. Gunawi (University of Chicago)</i>
15	<b>When the Network Crumbles: An Empirical Study of Cloud Network Failures and their Impact on Services</b> <i>Rahul Potharaju (Purdue University) and Navendu Jain (Microsoft)</i>
16	<b>vTube: Efficient Streaming of Virtual Appliances Over Last-Mile Networks</b> <i>Yoshihisa Abe (Carnegie Mellon University), Roxana Geambasu (Columbia University), Kaustubh Joshi (AT&amp;T Research), H. Andres Lagar-Cavilla (GridCentric), and Mahadev Satyanarayanan (Carnegie Mellon University)</i>
17	<b>Scalable Lineage Capture for Debugging DISC Analytics</b> <i>Dionysios Logothetis (Telefonics, I+D), Soumyarupa De (Microsoft, Inc.), Vicky Papavasileiou (U.C. San Diego), Zhaomo Yang (U.C. San Diego), and Kenneth Yocum (U.C. San Diego and Illumina, Inc.)</i>
18	<b>Memory-Efficient GroupBy-Aggregate using Compressed Buffer Trees</b> <i>Hrishikesh Amur (Georgia Tech.), Wolfgang Richter and David G. Andersen (Carnegie Mellon University), Michael Kaminsky (Intel Labs), Karsten Schwan (Georgia Tech.), and Athula Balachandran and Erik Zawadzki (Carnegie Mellon University)</i>
19	<b>Memory Footprint Matters: Efficient Equi-Join Algorithms for Main Memory Data Processing</b> <i>Spyros Blanas and Jignesh M. Patel (University of Wisconsin-Madison)</i>
20	<b>Scale-up vs Scale-out for Hadoop: Time to rethink?</b> <i>Raja Appuswamy (Vrije Universiteit) and Christos Gkantsidis, Dushyanth Narayanan, Orion Hodson, and Antony Rowstron (Microsoft Research)</i>
21	<b>EventWave: Programming Model and Runtime Support for Tightly-Coupled Elastic Cloud Applications</b> <i>Wei-Chiu Chuang (Department of Computer Science, Purdue University), Bo Sang, Sunghwan Yoo, and Rui Gu (Department of Computer Science Purdue University), Charles Killian (Google), and Milind Kulkarni (School of Electrical and Computer Engineering, Purdue University)</i>
22	<b>Tolerating Business Failures in Hosted Applications</b> <i>Jean-Sebastien Legare, Dutch T. Meyer, Mark Spear, Alexandru Totolici, Sara Bainbridge, Kalan MacRow, Robert Sumi, Quinlan Jung, Dennis Tjandra, David Williams-King, William Aiello, and Andrew Warfield (University of British Columbia)</i>
23	<b>Consistency Without Borders</b> <i>Peter Alvaro, Peter Bailis, Neil Conway, and Joseph M. Hellerstein (UC Berkeley)</i>
44	<b>High Performance Clustering of Social Images in a Map-Collective Programming Model</b> <i>Bingjing Zhang (Indiana University); Judy Qiu (Indiana University)</i>
45	<b>Does RDMA-based Enhanced Hadoop MapReduce Need a New Performance Model?</b> <i>Md. Wasi-ur-Rahman; Xiaoyi Lu; Nusrat Sharmin Islam; Dhableswar K. (DK) Panda (The Ohio State University)</i>
46	<b>Syndicate: Democratizing Cloud Storage and Caching through Service Composition</b> <i>Jude Nelson (Princeton University); Larry Peterson (Princeton University)</i>
47	<b>Coloring the Cloud for Predictable Performance</b> <i>Alberto Scolari, Filippo Sironi, Davide B. Bartolini, Donatella Sciuto, and Marco D. Santambrogio (Politecnico di Milano)</i>
48	<b>Decentralized Privacy Protection Strategies for Location-based Services</b> <i>Chih-Chun Chen (National Chung Cheng University); Yu-Ling Hsueh (National Chung Cheng University)</i>
49	<b>Recommending Just Enough Memory for Analytics</b> <i>Charles Reiss (UC Berkeley); Randy H. Katz (UC Berkeley)</i>
50	<b>Process-Oriented Recovery for Operations on Cloud Applications</b> <i>Min Fu; Liming Zhu; Anna Liu; Xiwei Xu; Len Bass (University of New South Wales, NICTA)</i>
51	<b>CloudSSI: Revisiting SSI in cloud era</b> <i>Mansoor Alicherry (Vizury Interactive); Ashok Anand (Instart Logic); Shoban Preeth Chandrabose (University of Wisconsin); Theophilus Benson (Duke University)</i>
52	<b>Firewall Placement in Cloud Data Centers</b> <i>Seungjoon Lee (AT&amp;T Labs - Research), Manish Purohit (University of Maryland), and Barna Saha (AT&amp;T Labs - Research)</i>
53	<b>Harmony: Coordinating Network, Compute, and Storage in Software-Defined Clouds</b> <i>Robert Grandl (University of Wisconsin, Madison); Yizheng Chen (VM Ware); Junaid Khalid (University of Wisconsin, Madison); Suli Yang (University of Wisconsin, Madison); Ashok Anand (Instart Logic); Theophilus Benson (Duke University); Aditya Akella (University of Wisconsin, Madison)</i>
54	<b>Pregelix: Dataflow-Based Big Graph Analytics</b> <i>Yingyi Bu (University of California, Irvine)</i>
55	<b>Wide-area Streaming Analytics: Distributing the Data Cube</b> <i>Benjamin Heintz (University of Minnesota); Abhishek Chandra (University of Minnesota); Ramesh K. Sitaraman (University of Massachusetts)</i>
56	<b>High Performance In-memory Caching through Flexible Fine-grained Services</b> <i>Yue Cheng (Virginia Tech); Aayush Gupta (IBM Almaden Research Center); Anna Povzner (IBM Almaden Research Center); Ali R. Butt (Virginia Tech)</i>
57	<b>Compiling Machine Learning Algorithms with SystemML</b> <i>Matthias Boehm; Douglas Burdick; Alexandre Evfimievski; Berthold Reinwald; Prithviraj Sen; Shirish Tatikonda; Yuanyuan Tian (IBM Research - Almaden)</i>
58	<b>SuperCloud: Economical Cloud Service on Multiple Vendors</b> <i>Qin Jia; Robbert Van Renesse; Hakim Weatherspoon (Cornell University)</i>
59	<b>Dynamic Performance Profiling of Cloud Caches</b> <i>Hjortur Bjornsson (University of Iceland), Gregory Chockler (University of London, Royal Holloway), and Trausti Saemundsson and Ymir Vigfusson (Reykjavik University)</i>
60	<b>DEDIS: Distributed Exact Deduplication for Primary Storage Infrastructures</b> <i>João Paulo, José Pereira (INESC TEC &amp; Universidade do Minho)</i>

# Wednesday, October 2nd at the Locatelli Center

