

2014 International Workshop on Data-Intensive Scalable Computing Systems

(DISCS 2014)

**New Orleans, Louisiana, USA
16 November 2014**



**IEEE Catalog Number: CFP14A57-POD
ISBN: 978-1-4799-7160-2**

2014 International Workshop on Data Intensive Scalable Computing Systems

DISCS 2014

Table of Contents

Workshop Organization	v
-----------------------------	---

Programming Models and Frameworks

Par-BF: A Parallel Partitioned Bloom Filter for Dynamic Data Sets	1
<i>Yi Liu, Xiongzi Ge, David H.C. Du, and Xiaoxia Huang</i>	
dispel4py: A Python Framework for Data-Intensive Scientific Computing	9
<i>Rosa Filguiera, Iraklis Klampanos, Amrey Krause, Mario David, Alexander Moreno, and Malcolm Atkinson</i>	
Efficient, Failure Resilient Transactions for Parallel and Distributed Computing	17
<i>Jay Lofstead, Jai Dayal, Ivo Jimenez, and Carlos Maltzahn</i>	

I/O and Storage

BPAR: A Bundle-Based Parallel Aggregation Framework for Decoupled I/O Execution	25
<i>Teng Wang, Kevin Vasko, Zhuo Liu, Hui Chen, and Weikuan Yu</i>	
Rethinking Key-Value Store for Parallel I/O Optimization	33
<i>Yanlong Yin, Antonios Kougkas, Kun Feng, Hassan Eslami, Yin Lu, Xian-He Sun, Rajeev Thakur, and William Gropp</i>	
PSA: A Performance and Space-Aware Data Layout Scheme for Hybrid Parallel File Systems	41
<i>Shuibing He, Yan Liu, and Xian-He Sun</i>	

Algorithms and Operations

Distributed Multipath Routing Algorithm for Data Center Networks	49
<i>Eun-Sung Jung, Venkatram Vishwanath, and Rajkumar Kettimuthu</i>	
CULZSS-Bit: A Bit-Vector Algorithm for Lossless Data Compression on GPGPUs	57
<i>Adnan Ozsoy</i>	

A Caching Approach to Reduce Communication in Graph Search Algorithms	65
<i>Pietro Cicotti and Laura Carrington</i>	
Mapping of RAID Controller Performance Data to the Job History on Large Computing Systems	73
<i>Marc Hartung and Michael Kluge</i>	
Author Index	81