

# **2020 IEEE/ACM Fifth International Parallel Data Systems Workshop (PDSW 2020)**

**Atlanta, Georgia, USA  
12 November 2020**



**IEEE Catalog Number: CFP20W50-POD  
ISBN: 978-1-6654-1595-8**

**Copyright © 2020 by the Institute of Electrical and Electronics Engineers, Inc.  
All Rights Reserved**

*Copyright and Reprint Permissions:* Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

***\*\*\* This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP20W50-POD
ISBN (Print-On-Demand):	978-1-6654-1595-8
ISBN (Online):	978-1-6654-1594-1

**Additional Copies of This Publication Are Available From:**

Curran Associates, Inc  
57 Morehouse Lane  
Red Hook, NY 12571 USA  
Phone: (845) 758-0400  
Fax: (845) 758-2633  
E-mail: [curran@proceedings.com](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# 2020 IEEE/ACM Fifth International Parallel Data Systems Workshop (PDSW) **PDSW 2020**

## Table of Contents

Message from the Workshop Chairs .v.....  
Workshop Organization .vi.....

### Session 1

Keeping It Real: Why HPC Data Services Don't Achieve I/O Microbenchmark Performance .1.....  
*Philip Carns (Argonne National Laboratory, USA), Kevin Harms (Argonne National Laboratory, USA), Bradley W. Settlemyer (Los Alamos National Laboratory, USA), Brian Atkinson (Los Alamos National Laboratory, USA), and Robert B. Ross (Argonne National Laboratory, USA)*

Towards On-Demand I/O Forwarding in HPC Platforms .7.....  
*Jean Luca Bez (Federal University of Rio Grande do Sul), Francieli Zanon Boito (University of Bordeaux), Alberto Miranda (Barcelona Supercomputing Center), Ramon Nou (Barcelona Supercomputing Center), Toni Cortes (Polytechnic University of Catalonia), and Philippe Olivier Alexandre Navaux (Federal University of Rio Grande do Sul)*

### Session 2

Gauge: An Interactive Data-Driven Visualization Tool for HPC Application I/O Performance Analysis .15.....  
*Eliakin del Rosario (Texas A&M University), Mikaela Currier (Texas A&M University), Mihailo Isakov (Texas A&M University), Sandeep Madireddy (Argonne National Laboratory (ANL)), Prasanna Balaprakash (Argonne National Laboratory (ANL)), Philip Carns (Argonne National Laboratory (ANL)), Robert B. Ross (Argonne National Laboratory (ANL)), Kevin Harms (Argonne National Laboratory (ANL)), Shane Snyder (Argonne National Laboratory (ANL)), and Michel A. Kinsy (Texas A&M University)*

Fractional-Overlap Declustered Parity: Evaluating Reliability for Storage Systems .22.....  
*Huan Ke (University of Chicago), Haryadi S. Gunawi (The University of Chicago), Dominic Manno (Los Alamos National Laboratory), David Bonnie (Los Alamos National Laboratory), and Bradley W. Settlemyer (Los Alamos National Laboratory)*

GPU Direct I/O with HDF5 .28.....  
*John Ravi (North Carolina State University, USA), Suren Byna (Lawrence Berkeley National Laboratory, USA), and Quincey Koziol (Lawrence Berkely National Laboratory, USA)*

## Session 3

Emulating I/O Behavior in Scientific Workflows on High Performance Computing Systems .34.....	
<i>Fahim Chowdhury (Florida State University), Yue Zhu (Florida State University), Francesco Di Natale (Lawrence Livermore National Laboratory), Adam Moody (Lawrence Livermore National Laboratory), Elsa Gonsiorowski (Lawrence Livermore National Laboratory), Kathryn Mohror (Lawrence Livermore National Laboratory), and Weikuan Yu (Florida State University)</i>	
Pangeo Benchmarking Analysis: Object Storage vs. POSIX File System .40.....	
<i>Haiying Xu (National Center for Atmospheric Research, USA), Kevin Paul (National Center for Atmospheric Research, USA), and Anderson Banihirwe (National Center for Atmospheric Research, USA)</i>	
Fingerprinting the Checker Policies of Parallel File Systems .46.....	
<i>Runzhou Han (Iowa State University, USA), Duo Zhang (Iowa State University, USA), and Mai Zheng (Iowa State University, USA)</i>	
<b>Author Index 53</b> .....	