

2021 IEEE/ACM Sixth International Parallel Data Systems Workshop (PDSW 2021)

**St. Louis, Missouri, USA
15 November 2021**



**IEEE Catalog Number: CFP21W50-POD
ISBN: 978-1-6654-1416-6**

**Copyright © 2021 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:	CFP21W50-POD
ISBN (Print-On-Demand):	978-1-6654-1416-6
ISBN (Online):	978-1-6654-1837-9

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
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

2021 IEEE/ACM Sixth International Parallel Data Systems Workshop (PDSW) **PDSW 2021**

Table of Contents

Message from the Workshop Chairs	v
Workshop Organization	vi

Session 1

New Challenges of Benchmarking All-Flash Storage for HPC	1
<i>Glenn K. Lockwood (Lawrence Berkeley National Laboratory), Alberto Chiusole (Lawrence Berkeley National Laboratory), and Nicholas J. Wright (Lawrence Berkeley National Laboratory)</i>	
Understanding the I/O Impact on the Performance of High-Throughput Molecular Docking	9
<i>Stefano Markidis (KTH Royal Institute of Technology, Sweden; Politecnico di Milano, Italy), Davide Gadioli (Politecnico di Milano, Italy), Emanuele Vitali (Politecnico di Milano, Italy), and Gianluca Palermo (Politecnico di Milano, Italy)</i>	
I/O Bottleneck Detection and Tuning: Connecting the Dots using Interactive Log Analysis	15
<i>Jean Luca Bez (Lawrence Berkeley National Laboratory), Houjun Tang (Lawrence Berkeley National Laboratory), Bing Xie (Oak Ridge National Laboratory), David Williams-Young (Lawrence Berkeley National Laboratory), Rob Latham (Argonne National Laboratory), Rob Ross (Argonne National Laboratory), Sarp Oral (Oak Ridge National Laboratory), and Suren Byna (Lawrence Berkeley National Laboratory)</i>	
Data-Aware Storage Tiering for Deep Learning	23
<i>Cong Xu (Hewlett Packard Labs), Suparna Bhattacharya (Lawrence Berkeley National Lab), Martin Foltin (Hewlett Packard Labs), Suren Byna (Lawrence Berkeley National Lab), and Paolo Faraboschi (Hewlett Packard Labs)</i>	

Session 2

SCTuner: An Autotuner Addressing Dynamic I/O Needs on Supercomputer I/O Subsystems	29
<i>Houjun Tang (Lawrence Berkeley National Laboratory), Bing Xie (Oak Ridge National Laboratory), Suren Byna (Lawrence Berkeley National Laboratory), Philip Carns (Argonne National Laboratory), Quincey Koziol (Lawrence Berkeley National Laboratory), Sudarsun Kannan (Rutgers University), Jay Lofstead (Sandia National Laboratories), and Sarp Oral (Oak Ridge National Laboratory)</i>	

User-Centric System Fault Identification Using IO500 Benchmark	35
<i>Radita Liem (RWTH Aachen University), Dmytro Povaliaiev (RWTH Aachen University), Jay Loftstead (Sandia National Labs), Julian Kunkel (Göttingen University), and Christian Terboven (RWTH Aachen University)</i>	
Verifying IO Synchronization from MPI Traces	41
<i>Sushma Yellapragada (University of Illinois Urbana-Champaign, Illinois), Chen Wang (University of Illinois Urbana-Champaign, Illinois), and Marc Snir (University of Illinois Urbana-Champaign, Illinois)</i>	
Author Index	47