# 2026 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW 2026)

## New Orleans, Louisiana, USA
## 25-29 May 2026

## Pages 1–455

IEEE Catalog Number: CFP2651J-POD
ISBN: 979-8-3195-3090-5

IEEE Advancing Technology for Humanity

Copyright © 2026, IEEE

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: | CFP2651J-POD |
| ISBN (Print-On-Demand): | 979-8-3195-3090-5 |
| ISBN (Online): | 979-8-3195-3089-9 |
| ISSN: | 2639-3867 |

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



# TABLE OF CONTENTS

Raiders of the Lost Log: Synchronous Parallel In-Place Models and Algorithms ..... 1
*Michael T. Goodrich, Vinesh Sridhar*

Accelerating HPC Applications on Heterogeneous Distributed Systems with Standard C++ ..... 11
*Beatrice Branchini, Ian Di Dio Lavore, Vito Giovanni Castellana, Antonino Tumeo, Marco Santambrogio*

FastShard: Fast Transaction Scheduling in Blockchain Sharding ..... 21
*Ramesh Adhikari, Costas Busch, Miroslav Popovic*

Radiation Hydrodynamics at Scale: Comparing MPI and Asynchronous Many-Task Runtimes with FleCSI ..... 31
*Alexander Strack, Hartmut Kaiser, Dirk Pflüger*

Multithreaded Fine-Grained Asynchronous BSP for Integer Sorting with LCI and OpenMP ..... 41
*Minyu Cheng, Jiakun Yan, Marc Snir*

A Two-Tier Secure Event Model for Privacy-Preserving IoT with Encrypted Reductions: A Parametric Feasibility Study ..... 51
*Yushan Li, Satoshi Fujita*

Heterogeneous Non-Linear Performance Modeling of Parameter Interactions ..... 55
*Steven D. Harris, Roger D. Chamberlain, Christopher D. Gill*

Efficient Degree-3 HUBO Formulation with Auxiliary-Variable-Minimum QUBO Reduction for the Angular Distance Traveling Salesman Problem ..... 65
*Yutaro Nakahara, Shunsuke Tsukiyama, Koji Nakano, Yasuaki Ito, Victor Parque, Takumi Kato, Kaiki Ii Kanto Teranishi*

Immediate Deterministic Leader Election (IDLE) Algorithm for Reducing Downtime in Distributed Systems ..... 73
*Masahiro Tanaka, Hideyuki Kawashima*

Improving Matricized Tensor Times Khatri-Rao Product Performance for Blocked Sparse Tensors ..... 83
*Kobe Bergmans, Karl Meerbergen, Raf Vandebril*

Fully Dynamic Rooted Spanning Tree on GPU ..... 92
*Abhijeet Sahu, Harmit Singh, Soham Nandy, G. Ramakrishna*

Time and Communication Balanced Transaction Scheduling on Trees under the Dual-flow Model ..... 102
*Alec Broadhurst, Michael Meyers, Pavan Poudel*

Accelerating Clarke and Wright’s savings algorithm by multi-threaded parallel processing with SIMD assistance for solving large-scale Euclidean CVRPs ..... 112
*Kenji Kanazawa*

Energy–Performance Trade-offs in Privacy-Preserving Federated Learning on SmartNIC-Enabled HPC Systems ..... 122
*Olivera Kotevska, Nguyen Duc, Oscar Hernandez*

Optimizing Allreduce Operations for Modern Heterogeneous Architectures with Multiple Processes per GPU ..... 131
*Michael Adams, Amanda Bienz*

Incidence Constraints in Hypergraph Partitioning on GPU ..... 141
*Marco Ronzani, Cristina Silvano*

A Benchmark Suite for Evaluating Scientific AI Workloads on GPUs ..... 145
*Zheming Jin, David Bernholdt, Feiyi Wang, Junqi Yin*

Quantitative Characterization of Host-Initiated CUDA Memory Allocators: Performance, Fragmentation, and Concurrency Trade-offs ..... 153
*Xun Jia, Shuangshuang Jin*

Adapting Workloads for CXL-Based Disaggregated-Memory Systems ..... 163
*Sayan Ghosh, Elise Song, Andrés Márquez*

Towards an Adaptive Runtime System for Cloud-Native HPC ..... 167
*Aditya Bhosale, Advait Tahilyani, Laxmikant Kale, Sara Kokkila-Schumacher*

A Statically Scheduled Vector Accelerator with Explicit Data Movement Control ..... 176
*Selim Sandal, Yusuf Can Akbas, Ismail Akturk*

Exploring C++ Standard Parallelism for GPU Programming in a Particle-In-Cell Application ..... 186
*Ester El Khoury, Mathieu Lobet, Laurent Colombet, Julien Bigot, Juan Jose Silva-Cuevas, Kévin Peyen*

ET: Bridging the Gap on Energy Telemetry for Multi-GPU Communication Collectives ..... 196
*Brandon Tran, Matthias Maiterth, Woong Shin, Matthew D. Sinclair, Shivaram Venkataraman*

Constraint-aware Optimization in Auto-Tuning ..... 200
*Floris-Jan Willemsen, Stijn Heldens, Rob V. van Nieuwpoort, Ben van Werkhoven*

ParCoder: Harnessing Transformer Models for Parallel Code Generation with OpenMP and MPI ..... 210
*Le Chen, Arijit Arijit, Nesreen Ahmed, Niranjan Hasabnis, Mihai Capota, Pablo Munoz, Gal Oren, Vy Vo Barbara Chapman, Ali Jannesari*

Modality Inflation: Energy Characterization and Optimization Opportunities for MLLM Inference ..... 220
*Mona Moghadampanah, Adib Rezaei Shahmirzadi, Farhana Amin, Dimitrios S. Nikolopoulos*

VibeCodeHPC: A Multi LLM Agent System for HPC Code Auto-Tuning ..... 229
*Shun-Ichiro Hayashi, Koki Morita, Daichi Mukunoki, Tetsuya Hoshino, Takahiro Katagiri*

Manual, Translated, or LLM-Generated? Comparing Paths to OpenMP Target Offload for LULESH ..... 242
*Yehonatan Fridman, Gal Oren*

OpenSHMEM and Rust: Towards PGAS in Memory-Safe Environments ..... 251
*Yash Kumar, Michael Beebe, Brody Williams, Yong Chen, Wendy Poole, Steve Poole*

Dynamic Contract Analysis for Parallel Programming Models ..... 261
*Yussur Mustafa Oraji, Alexander Huck, Christian Bischof*

Assessing the impact of MPI process placement on communication in many-core systems ..... 271
*Jackson Wesley, Whit Schonbein, David DeBonis, Matthew G. F. Dosanjh, Amanda Bienz*

Multi-GPU Memory Coherence for BLAS Matrices ..... 280
*Romain Pereira, Pierre-Etienne Polet, Thierry Gautier, Swann Perarnau*

Deriving Symbolic AET Using Locality Identity Equation ..... 291
*Fangzhou Liu, Yifan Zhu, Chen Ding, Yekai Pan, Yanghui Wu*

Performance-portable Acceleration of a DNS Turbulent Reacting Flows Application with OPS ..... 301
*Ashutosh S. Londhe, Hamid Kavari, Vishnu Mohan, Joe Kaushal, Istvan Z. Reguly, Nilanjan Chakraborty R. Stewart Cant, Gihan R. Mudalige*

Evaluating LLM Generated Task Codes ..... 312
*Sanjana Yasna, Simon Garcia De Gonzalo, Michael Robson*

Mitigating Power Fluctuation in Performance Counter Measurements on NVIDIA H100 GPU ..... 316
*Ryoma Ohara, Toshihiro Hanawa, Yohei Miki*

Pipelined FPGA Implementation of the QP-DYN Cipher with Scalable Dynamical System Parallelism ..... 326
*Paolo Palazzari, Luigi Accardi, Pasquale Tommasino, Alesandro Trifiletti*

Using FPGA-based Network-Attached Accelerators for Energy-Efficient AI Training in HPC Datacenters ..... 334
*Niklas Schelten, Steffen Christgau, Merit Hutzler, Philipp Kreowsky, Marco De Lucia, Bettina Schnor Hannes Signer, Johannes Spazier, Benno Stabernack, Serhii Yahdzhyiev*

100 Gbps Hash-Based Reconfigurable Pattern Matching ..... 341
*Magnus Östgren, Anna-Maria Unterberger, Ioannis Sourdis*

A Power-Centric Methodology to Characterize Edge AI SoC with Limited Telemetry Capabilities ..... 349
*Giulio Mantovi, Davide Paltrinieri, Giuseppe Sorrentino, Christian Pilato, Davide Conficconi*

Programming Domain-Specific FPGA Hardblocks from HLS: An RTL Blackbox Approach ..... 356
*Ruthwik Reddy Sunketa, Jeevesh Choudhury, Aman Arora*

Configurable Architecture Design Automation (CADA) - Enabling Efficient, Dynamically Reconfigurable Systems ..... 363
*Tim Ling, Jack Burd, Hongseok Lee, Michael Kuan-Tse Lee, Chia-Hsiang Yang, Dejan Markovic*

Design Space Exploration of Convolutional Neural Network Accelerators for Predictive Maintenance ..... 370
*Florian Rokohl, Sudatta Mondal, Edris Zaman Farsa, Marc Reichenbach*

From Algorithm to RTL: A Comparative Study of Compiler-Driven and LLM-Driven FPGA Design Flows ..... 375
*Md Mahfuzul Haque Gazi, Rickard Ewetz, Hao Zheng, Md Rubel Ahmed*

Comparing High-Level Synthesis and RTL for Simple FPGA Kernels – Revisited ..... 380
*Vinicius R. Fertrin, Christopher D. Gill, Roger D. Chamberlain*

Hokidachi: Low-Latency Random Forest Inference on FPGAs through full Horizontal Parallelism ..... 381
*Andrea Bellocci, Alessandro Verosimile, Marco D. Santambrogio*

A Minimal and Jitter-Free FPGA Telemetry Core for Cycle-Accurate Determinism in Real-Time Systems ..... 382
*Ceren Korkmaz, Büşra Bulut, Ayşe Yayla*

A Lightweight Slow Control Interlock System Targeting the European GateMate™ FPGA ..... 383
*Daniel Anders, Vitalii Burtsev, Daniele Passaretti, Christof Pfannenmüller, Fabian Lurz, Thilo Pionteck*

End-to-End Reconfigurable Rack-Scale Systems ..... 384
*C. Bash, R. Beausoleil, P. Bruel, C. Davidson, A. Dhakal, R. Enriquez, M. Fiorentino, J. Ignowski K. McAuliffe, D. Milojicic, A. Mishra, V. Mittal, C. Patel, S. Pasricha, P. Prakash, G. Ratihalli, D. Roweth S. Sury, B. Tossoun, A. Veprinsky, J. Zeiler*

High-Performance Hyperspectral Brain Tumor Analysis on Versal ACAP ..... 385
*Dismas Ezechukwu, Mark Sears, Chenfeng Zhao, Dinesh Bhatia*

A Top-Down/Bottom-Up Framework for Hardware Design Exploration of Mathematical Kernels ..... 386
*Angelos Ioannou, Mario Vega, Fabien Chaix, Dania Mosuli, Blair Reasoner, Tan Nguyen, Xiaokun Yang John Shalf, Doru Thom Popovici*

A Persistent-State Dataflow Accelerator for Memory-Bound Linear Attention Decode on FPGA ..... 387
*Neelesh Gupta, Peter Wang, Rajgopal Kannan, Viktor K. Prasanna*

GPU-Parallelizable Randomized Sketch-and-Precondition for Linear Regression using Sparse Sign Sketches ..... 388
*Tyler Chen, Pradeep Niroula, Archan Ray, Pragna Subrahmanya, Marco Pistoia, Niraj Kumar*

LLM-HPC++: Evaluating LLM-Generated Modern C++ and MPI+OpenMP Codes for Scalable Mandelbrot Set Computation ..... 398
*Patrick Diehl, Noujoud Nader, Deepti Gupta*

Where Did My NCCL Calls Go? A Profiler Comparison ..... 408
*Ruben Laso, Majid Salimi Beni, Ioannis Vardas, Siegfried Benkner, Sascha Hunold*

Efficient Communication-Aware Distributed Δ-Stepping for Single-Source Shortest Paths ..... 413
*Rakibul Hassan, Shaikh Arifuzzaman*

ScalSALE: A Scalable MPI+X Benchmark for Lagrangian Hydrodynamics on Heterogeneous Systems ..... 422
*Re'em Harel, Ron Wagner, Yehonatan Fridman, Matan Rusanovsky, Harel Levin, Gal Oren*

GPU synchronization-free multicolor DILU for Nvidia AmgX library ..... 430
*Franco Seveso, Ernesto Dufrechou, Pablo Ezzatti*

Communication-Aware Diffusion Load Balancing for Persistently Interacting Objects ..... 440
*Maya Taylor, Kavitha Chandrasekar, Laxmikant V. Kale*

Optimizing Multiobjective Genetic Algorithms For High-Order Epistasis Detection ..... 448
*Miguel Graça, Sergio Santander-Jiménez, Aleksandar Ilic*

Memory-Efficient Community Detection on Large Graphs Using Weighted Sketches ..... 456
*Subhajit Sahu, Abhijeet Sahu*

Edge-Native Data-Aware Runtime Systems for Distributed AI: Bridging HPC and Edge Workloads ..... 466
*Arne Hendricks*

Multi-Agent Based Time Critical FaaS Scheduling in Edge-Cloud Continuum ..... 473
*Yitao Lei, Mengbing Zhou, Paul Daniëlse, Zhiming Zhao*

Optimized Data Distribution Strategy in the Computing Continuum ..... 481
*Christian Schneeweiss, Kurt Horvath, Dragi Kimovski*

Observability and Self Healing of Quantum Secure V2X communication in Edge-Cloud Systems ..... 489
*Rohit Bantupalli, Kartick Sutradhar, Bheemappa Halavar*

Towards Heterogeneous Consumer Group Autoscaling for Distributed Event Queues ..... 493
*Garegin Tadevosyan, Miguel Areias, Jorge G. Barbosa*

Online Placement of Distributed Machine Learning Services over Heterogeneous IoT Platforms ..... 497
*Jérémie Rodez, Thierry Monteil, Patricia Stolf*

Zero-Shot Labeling with Foundation Models for Federated Learning on Edge Accelerators ..... 505
*Roopkatha Banerjee, Akash Sharma, Yogesh Simmhan*

ViSched: Reinforcement Learning-Based Scheduling for Vision-Language Model Inference on a Heterogeneous GPU Cluster ..... 509
*Ravi Gupta*

DYNAMIX: A Real-Time Adaptive Batch Scheduler for Heterogeneous Distributed ML ..... 513
*Yuanjun Dai, Keqiang He, An Wang*

Constrained Cuts, Flows, and Lattice-Linearity ..... 523
*Robert P. Streit, Vijay K. Garg*

Parallel Fractal Decomposition Search for Dynamic Continuous Optimization ..... 533
*Arcadi Llanza, Nadiya Shvai, Amir Nakib*

Cryptographically Computable Self-Learning Activation Functions for Privacy-Preserving Logistic Regression with Homomorphic Encryption ..... 544
*Bernardo Pulido-Gaytan, Horacio González-Vélez, Andrei Tchernykh*

A Constraint-Handling Capability Selection Method for Quantum-Inspired Annealers ..... 553
*Kotaro Bannai, Kazuhiko Komatsu, Takamasa Nakasone, Shintaro Momose, Hiroaki Kobayashi*

Compiler-Based Performance Results for Regular Application Kernels on the HPC CGRA HiPReP ..... 562
*Markus Weinhardt*

bitSMM: A bit-Serial Matrix Multiplication Accelerator ..... 569
*Pedro Antunes, Artur Podobas*

Predication in Elastic CGRAs ..... 578
*Omkar Bhilare, Omar Ragheb, Boma Adhi, Kentaro Sano, Jason Anderson, Tomohiro Ueno*

Control-Flow Execution on CGRAs: A Comprehensive Survey of Architectural and Compilation Techniques ..... 586
*Hisako Ito, Takuya Kojima, Hideki Takase, Hiroshi Nakamura*

FlowSpec: A Flexible and Scalable Simulation Framework for Coarse-Grained Spatial Architectures ..... 594
*Youngno Kim, Hyeonseo Kim, Eunseok Cho, Sanhtet Aung, Jongeun Lee*

Prefill/Decode-Aware Evaluation of LLM Inference on Emerging AI Accelerators ..... 602
*Shun Usami, Venkatram Vishwanath, E. Wes Bethel*

Performance of Vision-Language Models in Blast Severity Analysis ..... 610
*Angel M. Beltre, Douglas C. Crowder, Ritesh Kasamsetty, Jaimie Murdock, Lekha Patel Salvador G. Sanchez, Christopher M. Siefert, Ryan Roche, Nick Winovich*

Generalizing Scaling Laws for Dense and Sparse Large Language Models ..... 620
*Md Arafat Hossain, Xingfu Wu, Valerie Taylor, Ali Jannesari*

Evaluating the Cerebras Wafer Scale Engine 3 Against the NVIDIA H100 GPU for LLM Training and Inference ..... 628
*Aaron Shumer, Yang Ho, Chris Siefert*

PDE foundation model-accelerated inverse estimation of system parameters in inertial confinement fusion ..... 633
*Mahindra Rautela, Alexander Scheinker, Bradley Love, Diane Oyen, Nathan DeBardeleben, Earl Lawrence Ayan Biswas*

Evaluating LLM Coding Agents on SZ-Family Lossy Compression Across Architectures ..... 641
*Changqing Li, Shouwei Gao, Kai Zhao, Sheng Di, Wenqian Dong*

CelloAI Benchmarks: Toward Repeatable Evaluation of AI Assistants ..... 646
*Mohammad Atif, Kriti Chopra, Fang-Ying Tsai, Ozgur O. Kilic, Tianle Wang, Zhihua Dong Douglas Benjamin, Charles Leggett, Meifeng Lin, Paolo Calafiura, Salman Habib*

CoAttn: Concurrent Attention Engine for Optimized Mixed-Batch Inference in Large Language Models ..... 652
*Yuhao Wang, Zechun Zhou, Jingwei Sun, Guangzhong Sun*

EdgeCudaCoder: A Multi-Agent Framework to Generate Energy-Efficient CUDA Codes for Edge Computing Systems ..... 660
*Kshitij Bhardwaj, Harshitha Menon*

A Data-Driven, Forecast-Triggered EdgeOrchestration Workflow for Adaptive Air Quality Sensing Using NDP Catalog and SciDx Streaming ..... 671
*Sumaiya Azad, Jess Tate, Saleem Alharir, Daniel Mendoza, Simon Brewer, Manish Parashar*

IGEA: a Framework for Sleep Tracking via Edge-based Federated Learning on Multimodal Data ..... 679
*Susanna Bardini, Alessandro Verosimile, Giovanni Vaccarino, Marco D. Santambrogio*

Evaluating Deployment Trade-offs of Compact Vision-Language Models on Edge Systems ..... 683
*Alejandra Rios, Fatima Mora Garcia, Michael E. Papka*

A Preliminary Study of LLM Inferencing on Sandboxed Environments for Jetson Edge Platforms ..... 688
*Yash Vijay Kamble, Mayank Arya, Yogesh Simmhan*

Development and Validation of a Parallel and Distributed Computing Exam for Early Computing Students & Research* ..... 692
*Chris Bourke, David P. Bunde, Charlotte Gruner, Mary Smith, Xiaoyuan Suo*

An Educational Framework for Teaching Performance Engineering Using Deep Learning Workloads on CPU-only Clusters ..... 700
*Xiaoyuan Suo, Hayden Reinhold*

Leveraging Graph-Based Retrieval-Augmented Generation for Parallel and Distributed Computing Education ..... 708
*Rithvik Dhanpal, Vinay Regu, Ramaprasanna Devarapally, Jiayin Wang, Sushil Prasad, Srishti Srivastava Liudong Zuo, Michelle Zhu*

A Collaborative Classroom Activity for Teaching Cloud Architecting ..... 716
*Steven Santillan, Cristina L. Abad*

Online Self-Service Learning Modules for Running AI-Driven Scientific Workflows ..... 724
*Pari Hirenkumar Patel, Vani Seth, Bishwas Wagle, Roshan Lal Neupane, Ashish Pandey, Sanjit Subhash Sharan Srinivas, Prasad Calyam*

Evaluating LLM-Generated Instructional Content for Parallel Computing Education ..... 732
*Xinyao Yi, Yonghong Yan*

Knoxcraft Terps: Parallel Minecraft for Novices, a Peachy Assignment ..... 740
*Jurdana Masuma Iqrah, Jaime Spacco, Michael Gerten*

OMPeek: A Lightweight OpenMP Performance Monitor for Education ..... 744
*Esteban Stafford*

Learning-Augmented Algorithms for Graph Sparsification and Steiner Trees ..... 748
*Reyan Ahmed, Alvin Chiu, Mithun Ghosh, Michael Goodrich, Kwang-Sung Jun, Stephen Kobourov*

CHROMA: High-Quality Robust Graph Coloring on GPUs ..... 758
*Chou-Ying Hsieh, Chien-Hsien Wu, Sy-Yen Kuo*

Approximate Matching for Fuzzy Set Similarity ..... 768
*Michael Mandulak, S. M. Ferdous, Sayan Ghosh, Mahantesh Halappanavar, George M. Slota*

Alternative Decomposed Message Passing for Efficient Geometric GNNs ..... 772
*Kutay Tasci, Sanmukh Kuppannagari*

A Multi-Queue Design for High-Performance GPU Priority Queues ..... 783
*ZhongYang Li*

DistFusedMM: A Distributed-Memory Fused SDDMM-SpMM Algorithm for Graph Machine Learning ..... 792
*Isuru Ranawaka, Suhita Anubha, Dhaura Kariyawasam Pathiranage, Ariful Azad*

Sparsity-Aware Roofline Models for Sparse Matrix-Matrix Multiplication ..... 800
*Matthew Qian, Yahia Ramadan, Suhita Anubha, Ariful Azad*

Greedy Adaptive Task Splitting with Prediction and Placement for Heterogeneous Systems ..... 808
*Thomas Morin, Nathalie Furmento, Abdou Guermouche, Samuel Thibault, Pierre-André Wacrenier*

Heterogeneous Acceleration Framework for Unified Rendering of Hybrid NeRF-Mesh Representation ..... 818
*Yaqian Zhao, Xiaohui Zhang, Yuefeng Zhang, Lihua Lu, Ruyang Li, Hui Wei, Xin Lou*

SkipPar: A Hybrid CPU-GPU Framework for Accelerating LLM Training via Efficient Scheduling of Parameter Updates ..... 828
*Dheemanth Joshi, Karthik Namboori, Kevin Mahesh Kuriakose, Sathish Vadhiyar, Avanindra Singh Subhasis Banerjee*

Identifying and Addressing HMM Data Migration Overhead on NVIDIA GPUs ..... 838
*Nicholas Cassarino, Tyler Allen, Erik Saule*

AutoBAS: Automatic Topology-Aware Broadcast Algorithm Selection ..... 848
*Maren Wessel-Berg, Magnus Jahre, Jan Christian Meyer*

Hybrid Quantum-Classical Complex-Valued Classifier: A Statevector-Based Approach ..... 858
*Taoreed Akinola, Xiangfang Li, Lijun Qian*

Low-Latency Interconnects & HPC Resources: What Does Quantum Computing Really Need? ..... 864
*Elise Jennings, Philip Clarke, Owen Arnold*

Enhancing Quantum Reservoir Computing: Optimization Strategies and Real-World Applications ..... 869
*Sanket Shah, E. Wes Bethel*

EmuCSD: A Scalable Framework For Emulating Computational Storage Devices ..... 876
*Saleh AlSaleh, Wahid Uz Zaman, Mahmut Taylan Kandemir*

Chasing the Rabbit: A Systematic Exploration of Software-Defined HPC Storage ..... 885
*Hariharan Devarajan, Brian Behlendorf, Blake Devcich, Dean Roehrich, Matt Richerson, Tony Floeder*
*Brent Petit, Don Boyd, Trace Hamada, Tim McCree, Gian-Carlo Defazio, James Corbett*
*Elsa Gonsiorowski, Cameron Stanavige, Nikoli Dryden, Roger Pearce, Kathryn Mohror*
*Bronis R. de Supinski*

NZFS: A Null File System with Zero-Copy I/O Support for Application Benchmarking ..... 894
*Shingo Hattori, Osamu Tatebe*