1. Record Nr. UNINA9911047655603321

Autore	Neuwirth Sarah
Titolo	High Performance Computing : ISC High Performance 2025 International Workshops, Hamburg, Germany, June 10–13, 2025, Revised Selected Papers / / edited by Sarah Neuwirth, Arnab Kumar Paul, Tobias Weinzierl, Erin Claire Carson
Pubbl/distr/stampa	Cham : , : Springer Nature Switzerland : , : Imprint : Springer, , 2026
ISBN	3-032-07612-9
Edizione	[1st ed. 2026.]
Descrizione fisica	1 online resource (947 pages)
Collana	Lecture Notes in Computer Science, , 1611-3349 ; ; 16091
Disciplina	621.39 004.6
Soggetti	Computer engineering Computer networks Computer Engineering and Networks
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	-- 6th ISC HPC International Workshop on Monitoring \& Operational Data Analytics (MODA25). -- Duration-Informed Workload Scheduler. -- Monitoring Energy Consumption of Workloads on HPC Vega. -- A Unified I/O Monitoring Framework Using eBPF. -- Supporting HPC Users with LLview. -- What Time Taught Us: Monitoring a Computing Technology Testbed Across Multiple Years. -- 9th International Workshop on In Situ Visualization (WOIV'25). -- Enabling Modular In-situ Workflows through CatalystMaestro and CatalystComposer. -- Updating Inshimtu with Catalyst2 and Integrating an HPC MiniApp: Lessons Learned. -- Issues and challenges of deploying in-situ visualization for SPH codes. -- 5th International Workshop on Computational Aspects of Deep Learning (CADL). -- Direct Feedback Alignment for Recurrent Neural Networks. -- Assessing Tenstorrent's RISC-V MatMul Acceleration Capabilities. -- Automatically parallelizing batch inference on deep neural networks using Fiats and Fortran 2023 "do concurrent". -- Optimizing edge AI models on HPC systems with the edge in the loop. -- Evaluation of Distributed Asynchronous Checkpointing in High-Performance Computing. -- Energy Efficiency

with Sustainable Performance: Techniques, Tools, and Best Practices (EESP). -- Characterizing GPU Energy Usage in Exascale-Ready Portable Science Applications. -- What A Waste. -- Pinpointing Idle-Power Regressions in Linux. -- DARE-ML: Democratized Accessible Resource Environment for Machine Learning in the SUPERCOM platform. -- Power-Capping Metric Evaluation for Improving Energy Efficiency. -- Experience on clock rate adjustment for energy-efficient GPU-accelerated real-world codes. -- Running Energy-Efficient HPL on APUs: Strategies and Best Practices. -- Analysis of Application Power Characteristics Using Performance Counters on A64FX. -- Fifth workshop on Compiler-assisted Correctness Checking and Performance Optimization for HPC (C3PO'25). -- CGPatch: Streamlining Static Call Graph Validation Using Selective Instrumentation. -- Speculative Recursion Unrolling. -- From C to Rust: Evaluating LLM Capabilities in Transpilation through Compilation Errors. -- CompilerGPT: Leveraging Large Language Models for Analyzing and Acting on Compiler Optimization Reports. -- Improving compiler support for SIMD offload using Arm Streaming SVE. -- Fifth Workshop on Interactive and Urgent HPC (WIUHPC). -- Dynamic Resource Management Framework for Elastic Computing. -- Enabling Seamless Transitions from Experimental to Production HPC for Interactive Workflows. -- A novel approach to dynamic computing using Slurm. -- Fifth workshop on Communication, I/O, and Storage at Scale on Next-Generation Platforms – Scalable Infrastructures (ixpug-comm-io-storage). -- Combining Malleability and Distributed Control Mechanisms to Reduce I/O Contention. -- DOCA UROM: A Vehicle for Offloading HPC and AI to DPUs. -- Accelerating I/O in Scientific Workflows with the Impact of Apache Ignite’s In-Memory File System. -- HPC on Heterogeneous Hardware (H3). -- Generation of Mixed-precision Kernels for Quantized Transformer Encoders with Exo. -- Training an Image Classification Model on a Supercomputer with AMD Genoa Compute Nodes. -- Exploring QUBO on LPUs for Engineering. -- Investigating Matrix Repartitioning to Address the Over- and Undersubscription Challenge for a GPU-based CFD Solver. -- Stream-K++: Adaptive GPU GEMM Kernel Selection and Scheduling for AI using Bloom Filters. -- Accelerating Electrostatics Simulations with GPUs. -- International workshop on RISC-V for HPC at ISC. -- Streamlining Fedora Linux Distributions for RISC-V: A Scalable and Automated Approach. -- Evaluating RISC-V processor as an alternative for High Performance Computing. -- Evaluation of RVV-enabled COTS Platforms with Matrix Multiplication and EXO. -- Advancing the RISC-V Performance Simulation Ecosystem with Data Prefetching. -- RISC-V in HPC: a look into tools for performance monitoring. -- Monte Cimone v2: HPC RISC-V Cluster Evaluation and Optimization. -- Parallel FFTW on RISC-V: A Comparative Study including OpenMP, MPI, and HPX. -- Exploring Fast Fourier Transforms on the Tenstorrent Wormhole. -- The First International Workshop on Foundational Large Language Models Advances for HPC (LLM4HPC). -- Leveraging AI for Productive and Trustworthy HPC Software: Challenges and Research Directions. -- LLM & HPC:Benchmarking DeepSeek's Performance in High-Performance Computing Tasks. -- Analysis of MPI Parallel Code Generated by GPT-4o. -- Workshop Review Contributions. -- QRUCH : Quantum Resources for Unified Computing Hub. -- High Performance Container Workshop 2025. -- A2SD: Advancing Autonomous Scientific Discovery Workshop. -- The Future of Benchmarks in Supercomputing. -- Research Software Engineering in High-Performance Computing - Tools and Techniques for Continuous Integration and Benchmarking.

Sommario/riassunto

This volume constitutes the revised selected papers of 40th

International Conference on High Performance Computing, ISC High Performance 2025, Hamburg, Germany, during June 10–13, 2025.