



# High Performance Computing [ 30th International Conference, ISC High Performance 2015, Frankfurt, Germany, July 12-16, 2015, Proceedings /

Kunkel, Julian M

Ludwig, Thomas

Computer science   Operating systems (Computers)   Computer network  
architectures   Computer system performance   Information theory  
Computer Science   System Performance and Evaluation   Performance and  
Reliability   Computer Systems Organization and Communication Networks  
Theory of Computation

Monografía

This book constitutes the refereed proceedings of the 30th International Conference, ISC High Performance 2015, [formerly known as the International Supercomputing Conference] held in Frankfurt, Germany, in July 2015. The 27 revised full papers presented together with 10 short papers were carefully reviewed and selected from 67 submissions. The papers cover the following topics: cost-efficient data centers, scalable applications, advances in algorithms, scientific libraries, programming models, architectures, performance models and analysis, automatic performance optimization, parallel I/O and energy efficiency

<https://rebiunoda.pro.baratznet.cloud:38443/OpacDiscovery/public/catalog/detail/b2FpOmNlbgVicmF0aW9uOmVzLmJhemF0ei5yZW4vMTc5Nzg5MjA>

**Título:** High Performance Computing [Recurso electrónico] 30th International Conference, ISC High Performance 2015, Frankfurt, Germany, July 12-16, 2015, Proceedings edited by Julian M. Kunkel, Thomas Ludwig

**Mención de serie:** Lecture Notes in Computer Science 9137

**Contenido:** Asynchronous Iterative Algorithm for Computing Incomplete Factorizations on GPUs -- Matrix Multiplication on High-Density Multi-GPU Architectures: Theoretical and Experimental Investigations -- A Framework for Batched and GPU-Resident Factorization Algorithms Applied to Block Householder Transformations -- Parallel Efficient Sparse Matrix-Matrix Multiplication on Multicore Platforms -- On the Design, Development, and Analysis of Optimized Matrix-Vector Multiplication Routines for Coprocessors -- Large-Scale Neo-Heterogeneous Programming and Optimization of SNP Detection on Tianhe-2 -- ACCOLADES: A Scalable

Workflow Framework for Large-Scale Simulation and Analyses of Automotive Engines -- Accelerating LBM and LQCD Application Kernels by In-Memory Processing -- On Quantum Chemistry Code Adaptation for RSC PetaStream Architecture -- Dtree: Dynamic Task Scheduling at Petascale -- Feasibility Study of Porting a Particle Transport Code to FPGA -- A Scalable, Linear-Time Dynamic Cutoff Algorithm for Molecular Dynamics -- BWTCP: A Parallel Method for Constructing BWT in Large Collection of Genomic Reads -- Lattice-CSC: Optimizing and Building an Efficient Supercomputer for Lattice-QCD and to Achieve First Place in Green500 -- An Efficient Clique-Based Algorithm of Compute Nodes Allocation for In-memory Checkpoint System -- A Scalable Algorithm for Radiative Heat Transfer Using Reverse Monte Carlo Ray Tracing -- Optimizing Processes Mapping for Tasks with Non-uniform Data Exchange Run on Cluster with Different Interconnects -- Dynamically Adaptable I/O Semantics for High Performance Computing -- Predicting Performance of Non-contiguous I/O with Machine Learning -- A Best Practice Analysis of HDF5 and NetCDF-4 Using Lustre -- Striping Layout Aware Data Aggregation for High Performance I/O on a Lustre File System -- Hop: Elastic Consistency for Exascale Data Stores -- Energy-Efficient Data Processing Through Data Sparsing with Artifacts -- Updating the Energy Model for Future Exascale Systems -- High-Order ADER-DG Minimizes Energy- and Time-to-Solution of SeisSol -- Modeling the Productivity of HPC Systems on a Computing Center Scale -- Taking Advantage of Node Power Variation in Homogenous HPC Systems to Save Energy -- A Run-Time System for Power-Constrained HPC Applications -- A Machine Learning Approach for a Scalable, Energy-Efficient Utility-Based Cache Partitioning -- A Case Study - Cost of Preemption for Urgent Computing on SuperMUC -- Designing Non-blocking Personalized Collectives with Near Perfect Overlap for RDMA-Enabled Clusters -- Design Methodology for Optimizing Optical Interconnection Networks in High Performance Systems -- Quantifying Communication in Graph Analytics -- Formal Metrics for Large-Scale Parallel Performance -- Hunting Down Load Imbalance: A Moving Target -- Orchestrating Docker Containers in the HPC Environment -- Performance and Scaling of WRF on Three Different Parallel Supercomputers

**Restricciones de acceso:** Acceso restringido a miembros del Consorcio de Bibliotecas Universitarias de Andalucía

**Detalles del sistema:** Modo de acceso: world wide web

**Fuente de adquisición directa:** Springer (e-Books)

**ISBN:** 9783319201191 978-3-319-20119-1 9783319201184

**Autores:** Kunkel, Julian M Ludwig, Thomas

---

## Baratz Innovación Documental

- Gran Vía, 59 28013 Madrid
- (+34) 91 456 03 60
- [informa@baratz.es](mailto:informa@baratz.es)