

Accepted papers

- Xueying Wang, Guangli Li, Xiao Dong, Jiansong Li, Lei Liu and Xiaobing Feng, *Accelerating Deep Learning Inference with Cross-Layer Data Reuse on GPUs* (paper id:16)
- Xiao Meng and Lukasz Golab, *Parallel Scheduling of Data-Intensive Tasks* (paper id:18)
- Jonas Ellert, Johannes Fischer and Nodari Sitchinava, *LCP-Aware Parallel String Sorting* (paper id:31)
- Alexandro Baldassin, Rafael Murari, João Paulo Carvalho, Guido Araujo, Daniel Castro, João Barreto and Paolo Romano, *NVPhTM: An Efficient Phase-Based Transactional System for Non-Volatile Memory* (paper id:33)
- Giuseppe Antonio Di Luna, Paola Flocchini, Nicola Santoro, Giovanni Viglietta and Yukiko Yamauchi, *Mobile RAM and Shape Formation by Programmable Particles* (paper id:35)
- Elvis Rojas, Esteban Meneses Rojas, Terry Jones and Don Maxwell, *Towards a Model to Estimate the Reliability of Large-scale Hybrid Supercomputers* (paper id:37)
- Maha Alsayasneh and Noel De Palma, *A Learning-Based Approach for Evaluating the Capacity of Data Processing Pipelines* (paper id:43)
- Ming-Chang Lee, Jia-Chun Lin and Ernst Gunnar Gran, *Distributed Fine-Grained Traffic Speed Prediction for Large-Scale Transportation Networks based on Automatic LSTM Customization and Sharing* (paper id:46)
- Tobias Castanet, Olivier Beaumont, Nicolas Hanusse and Corentin Travers, *Approximation Algorithm for Estimating Distances in Distributed Virtual Environments* (paper id:57)
- Najeeb Ahmad, Buse Yilmaz and Didem Unat, *A Prediction Framework for Fast Sparse Triangular Solves* (paper id:62)
- Antoni Navarro, Arthur F. Lorenzon, Eduard Ayguadé and Vicenç Beltran, *Optimizing Resource Management through Prediction-based Policies* (paper id:65)
- Jorge Ejarque, Marta Bertran, Javier Álvarez Cid-Fuentes, Javier Conejero and Rosa M. Badiá, *Managing Failures in task-based parallel workflows in distributed computing environments* (paper id:72)
- Olivier Beaumont, Julien Langou, Willy Quach and Alena Shilova, *A Makespan Lower Bound for the Scheduling of the Tiled Cholesky Factorization based on ALAP scheduling* (paper id:75)
- Ismail Elhelw, Rutger Hofman and Henri Bal, *Accelerating Overlapping Community Detection: Performance Tuning a Stochastic Gradient Markov Chain Monte Carlo Algorithm* (paper id:76)
- Irvan Jahja, Haifeng Yu and Ruomu Hou, *On the Power of Randomization in Distributed Algorithms in Dynamic Networks with Adaptive Adversaries* (paper id:86)
- Dhruv Garg, Prathik Shiralkar, Anshu Shukla and Yogesh Simmhan, *TorqueDB: Distributed Querying of Time-series Data from Edge-local Storage* (paper id:90)
- Haewon Jeong, Yaoqing Yang, Christian Engelmann, Vipul Gupta, Tze Meng Low, Pulkit Grover, Viveck Cadambe and Kannan Ramchandran, *3D Coded SUMMA: Communication-Efficient and Robust Parallel Matrix Multiplication* (paper id:92)
- Fritz Goebel, Hartwig Anzt, Terry Cojean, Goran Flegar and Enrique S. Quintana-Orti, *Multiprecision block-Jacobi for Iterative Triangular Solves* (paper id:95)
- Fabian Schrammel, Florian Renk, Arya Mazaheri and Felix Wolf, *Efficient Ephemeris Models for Spacecraft Trajectory Simulations on GPUs* (paper id:101)
- Cristóbal Camarero, Carmen Martinez, Ramon Beivide and Javier Corral, *Modelling Standard and Randomized Slimmed Folded Clos Networks* (paper id:103)
- Nicolas Morew, Mohammad Norouzi, Ali Jannesari and Felix Wolf, *Skipping Non-essential Instructions Makes Data-dependence Profiling Faster* (paper id:105)
- Bertil Schmidt and Christian Hundt, *cuDTW++: Ultra-Fast Dynamic Time Warping on CUDA-enabled GPUs* (paper id:108)
- Bo Wang, Christian Terboven, Matthias Mueller and Julian Miller, *Operation-Aware Power Capping* (paper id:110)
- Olivier Beaumont, Lionel Eyraud-Dubois and Alena Shilova, *Optimal GPU-CPU Offloading Strategies for Deep Neural Network Training* (paper id:112)
- Tim Jammer, Christian Iwainsky and Christian Bischof, *A Comparison of the Scalability of OpenMP Implementations* (paper id:121)
- Lars van den Haak, Trevor McDonnel, Gabriele Keller and Ivo Gabe de Wolff, *Accelerating Nested Data Parallelism: Preserving Regularity* (paper id:124)
- Prithayan Barua, Jisheng Zhao and Vivek Sarkar, *OmpOptiMem: Optimized Memory Movement for Heterogeneous Computing* (paper id:135)
- Simone Economo, Sara Royuela Alcázar, Eduard Ayguadé Parra and Vicenç Beltran Querol, *A tool-chain to verify the parallelization of OmpSs-2 applications* (paper id:145)
- Qinglin Wang, Dongsheng Li, Xiandong Huang, Siqi Shen, Songzhu Mei and Jie Liu, *Optimizing FFT-based convolution on ARMv8 multi-core CPUs* (paper id:153)
- Changjiang Gou, Ali Al Zoobi, Anne Benoit, Mathieu Faverge, Loris Marchal, Grégoire Pichon and Pierre Ramet, *Improving mapping for sparse direct solvers: A trade-off between data locality and load balancing* (paper id:163)
- Rafael Campos, Diogo Marques, Sergio Santander-Jiménez, Leonel Sousa and Aleksandar Ilic, *Heterogeneous CPU+GPU Processing for Efficient Epistasis Detection* (paper id:164)
- Andrei Poenaru and Simon McIntosh-Smith, *Evaluating the Effectiveness of a Vector-Length-Agnostic Instruction Set* (paper id:165)
- Alexandre Denis, Emmanuel Jeannot, Philippe Swartvagher and Samuel Thibault, *Using Dynamic Broadcasts to improve Task-Based Runtime Performances* (paper id:175)
- Pedro Sanches, João A. Silva and Hervé Paulino, *Data-Centric Distributed Computing on Networks of Mobile Devices* (paper id:177)
- Sohan Lal, Aksel Alpay, Philip Salzmann, Biagio Cosenza, Alexander Hirsch, Nicolai Stawinoga, Peter Thoman, Thomas Fahringer and Vincent Heuveline, *SYCL-Bench: A Versatile Single-Source Benchmark Suite for Heterogeneous Computing* (paper id:181)
- S. Saberi, A. Vogel and G. Meschke, *Parallel Finite Cell Method with Adaptive Geometric Multigrid* (paper id:192)
- Sergi Vila Almenara, Josep Lluís Lerida, Fernando Cores, Fernando Guirado and Fabio Verdi, *WPSP: a multi-correlated weighted policy for VM selection and migration for Cloud computing* (paper id:194)
- Pedro Moreno, Miguel Areias and Ricardo Rocha, *A Compression-Based Design for Higher Throughput in a Lock-Free Hash Map* (paper id:210)
- Tariq Alturkestani, Hatem Ltaief and David Keyes, *Maximizing I/O Bandwidth for Reverse Time Migration on Heterogeneous Large-Scale Systems* (paper id:220)



Submission deadline for the Euro-Pas PhD Symposium has been extended to 27 May, 2023. Click here for more information - <https://t.co/wwxislCJSC>

17.05.2023 - 11:51

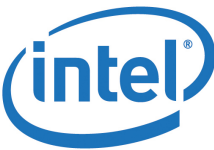
The Euro-Par PhD Symposium is a welcoming and supportive forum for PhD students to present their work. Click here for more information: <https://t.co/wwxislCJSC>

04.04.2023 - 09:25

Submit your paper for EURO-PAR 2023 Workshops and Minisymposia!
Click here for more information.
<https://t.co/UEseXWb3Dz>

07.03.2023 - 08:18

Abstract submission deadline is 24 Feb, 2023. Click here for more information: <https://t.co/UEseXWb3Dz>



SPONSOR DETAILS



SHARE ON:

