17th Workshop on Parallel Programming and Run-Time Management Techniques for Many-Core Architectures

15th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms

PARMA-DITAM 2026, January 27, 2026, Krakow, Poland

Edited by
Davide Baroffio
Paola Busia
Lev Denisov
Nitin Shukla

Editors

Davide Baroffio
Politecnico di Milano, Italy
davide.baroffio@polimi.it

Paola Busia
University of Cagliari, Italy
paola.busia@unica.it

Lev Denisov
Politecnico di Milano, Italy
lev.denisov@polimi.it

Nitin Shukla
CINECA, Casalecchio di Reno, Italy
n.shukla@cineca.it

ACM Classification 2012
Computing methodologies → Distributed algorithms; Software and its engineering → Distributed systems organizing principles; Computing methodologies; Computer systems organization → Embedded software; Software and its engineering → Compilers; Computer systems organization → Quantum computing; Software and its engineering → Just-in-time compilers; Software and its engineering → Retargetable compilers; Computing methodologies → Parallel programming languages; Computing methodologies → Simulation tools; Computer systems organization → Parallel architectures; Computer systems organization → Cloud computing; Computer systems organization → Heterogeneous (hybrid) systems; Hardware → Emerging architectures; Computer systems organization → Reconfigurable computing; Hardware → Modeling and parameter extraction; Computing methodologies → Parallel computing methodologies; Software and its engineering → Parallel programming languages; Software and its engineering → Distributed programming languages; Software and its engineering → Runtime environments; Computer systems organization → Distributed architectures

ISBN 978-3-95977-416-1

Published online and open access by
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at https://www.dagstuhl.de/dagpub/978-3-95977-416-1.

Publication date
April, 2026

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists all publications of this volume in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at https://portal.dnb.de.

Digital Object Identifier: 10.4230/OASIcs.PARMA-DITAM.2026.0
ISBN 978-3-95977-416-1 **ISSN 1868-8969** **https://www.dagstuhl.de/oasics**

OASIcs – OpenAccess Series in Informatics

OASIcs is a series of high-quality conference proceedings across all fields in informatics. OASIcs volumes are published according to the principle of Open Access, i.e., they are available online and free of charge.

ISSN 1868-8969

https://www.dagstuhl.de/oasics

Contents

17th Workshop on Parallel Programming and Run-Time Management Techniques for Many-Core Architectures and 15th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms (PARMA-DITAM 2026).
Editors: Davide Baroffio, Paola Busia, Lev Denisov, and Nitin Shukla
OpenAccess Series in Informatics
OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Preface

This volume contains the papers selected for presentation at the 17th Workshop on Parallel Programming and Run-Time Management Techniques for Many-Core Architectures and 15th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms (PARMA-DITAM 2026), held in Krakow (Poland), on January 27, 2026, in conjunction with the 21st edition of the High Performance and Embedded Architecture and Compilation (HiPEAC) conference.

Each anonymously submitted paper was reviewed by at least three members of the program committee, who did not have any conflicts of interest with the authors of the paper, and was evaluated for significance, novelty, technical quality, and potential to spark discussion among attendees from the academic and industry research communities. The program committee's work was carried out electronically, and, where needed, additional reviews were solicited. The final acceptance decision of all 7 accepted works was made by consensus among the reviewers.

The workshop program consists of two sessions, entitled "Compilers and Cloud HPC Tools" and "Heterogeneous Deployment".

The first session consists of 4 papers, exploring the integration of advanced compiler techniques and orchestration tools to enhance performance and manageability in cloud and High-Performance Computing (HPC) environments. The first paper introduces an open-source tool that integrates the Asynchronous Successive Halving Algorithm (ASHA) with the SLURM Experiment Management Library (SEML) to enable scalable and fault-tolerant hyperparameter optimization on shared clusters. The second paper presents a framework that integrates the VisIVO visualization tool with cloud-based services at Cineca, enabling researchers to perform GPU-accelerated, real-time, interactive visualization of large astrophysical simulations via Jupiter notebooks. The third paper proposes an inter-procedural extension to the strength reduction compiler optimization, replacing expensive operations, like exponentiation, with incremental updates across function calls to improve the efficiency of code generated for embedded platforms. The fourth and last paper of the first session, instead, describes a library for MLIR to translate quantum dialects into machine-learning-oriented primitives, facilitating the integration of quantum accelerators into hybrid HPC workflows.

The second session consists of 3 papers, describing solutions for optimizing simulation and orchestration of heterogeneous hardware, focusing on GPGPUs and FPGAs to improve resource utilization. The first paper addresses the issue of long runtime cycle-accurate GPGPU simulators, exploiting SIMT-Core cluster execution using a thread pool to achieve significant speed-ups with minimal code modifications. The second paper introduces a framework leveraging Pareto-optimal designs shorten the gap between High-Level Synthesis (HLS) and runtime orchestration, increasing the efficiency of FPGA-as-a-service, balancing the trade-off between performance and resource management. The third and last paper of the second session provides an end-to-end framework for modeling and mapping LLM inference on Coarse-Grained Reconfigurable Architectures (CGRAs), achieving high performance per Watt through accurate micro-architectural design space exploration.

In addition to the themed sessions, the workshop included three invited talks from Paul Carpenter (Barcelona Supercomputing Center), Ivy Peng (KTH Royal Institute of Technology), and Daniele Gregori (E4). The proceedings include an extended abstract of the invited talk by Paul Carpenter, detailing recent advances in the OmpSs-2@Cluster framework, focusing on how task offloading and dynamic load balancing can mitigate

17th Workshop on Parallel Programming and Run-Time Management Techniques for Many-Core Architectures and 15th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms (PARMA-DITAM 2026).
Editors: Davide Baroffio, Paola Busia, Lev Denisov, and Nitin Shukla

OpenAccess Series in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

performance bottlenecks in distributed-memory systems. The talk from Ivy Peng examined the current landscape of schedulers and runtimes, shedding the light on the critical needs for orchestrating data movement and power dynamics, outlining the requirements for achieving system-level gains in future accelerator-rich systems. The talk from Daniele Gregori discussed the advancements of E4 in the adoption of RISC-V in HPC via the Monte Cimone cluster, the first RISC-V ISA cluster built for co-design activities to enable its use in scientific and engineering applications.

Putting together a successful PARMA-DITAM 2026 was the result of the time and effort devoted by many individuals. We are thankful to the program committee and external reviewers for their work in evaluating several papers in a short time-frame and their dedication to the quality of the program. Moreover, we are also grateful to Sasha Uhrig (HiPEAC 2026 Workshops & Tutorials Chair) and Tomasz Kryjak (HiPEAC 2026 General Chair) for their support during workshop organization and logistics. Finally, we thank the authors of all submitted papers, as well as the workshop's attendees, for contributing to the program and the discussions surrounding it. We hope you will find this program interesting, stimulating, and inspiring for your future research.

List of Authors

Giovanni Agosta (4)
Politecnico di Milano, Italy; Consorzio Interuniversitario Nazionale per l'Informatica, Roma, Italy

Georgios Alexandris (8)
National Technical University of Athens, Greece

Umer Arshad (3)
Department of Information Engineering, Computer Science and Mathematics, University of L'Aquila, Italy

Hendrik Borras (2)
Hardware and Artificial Intelligence (HAWAII) Lab, Heidelberg University, Germany

Paul Carpenter (1)
Barcelona Supercomputing Center, Spain

Panagiotis Chaidos (8)
National Technical University of Athens, Greece

Anusha Chattopadhyay (2)
Hardware and Artificial Intelligence (HAWAII) Lab, Heidelberg University, Germany

Henk Corporaal (8)
Eindhoven University of Technology, Netherlands

Juliette Fournis d'Albiat (1)
Barcelona Supercomputing Center, Spain

Marco De Pascale (5)
Leibniz Supercomputing Centre, München, Germany

Jorge Echavarria (5)
Leibniz Supercomputing Centre, München, Germany

Aggelos Ferikoglou (7)
Microprocessors and Digital Systems Laboratory, ECE, NTUA, Athens, Greece

Holger Fröning (2)
Hardware and Artificial Intelligence (HAWAII) Lab, Heidelberg University, Germany

Manil Dev Gomony (8)
Eindhoven University of Technology, Netherlands

Mario Hernández Vera (5)
Leibniz Supercomputing Centre, München, Germany

Dionysios Kefallinos (8)
National Technical University of Athens, Greece

Bernhard Klein (2)
Hardware and Artificial Intelligence (HAWAII) Lab, Heidelberg University, Germany

Sohan Lal (6)
Hamburg University of Technology, Germany

Tim Lühnen (6)
Hamburg University of Technology, Germany

Alexis Maras (8)
National Technical University of Athens, Greece

Dimosthenis Masouros (7)
Microprocessors and Digital Systems Laboratory, ECE, NTUA, Athens, Greece

Giuseppa Muscianisi (3)
CINECA, Bologna, Italy

Muhammad Nufail Farooqi (5)
Leibniz Supercomputing Centre, München, Germany

Isabel Piedrahita (1)
Barcelona Supercomputing Center, Spain

Fabio Pitari (3)
CINECA, Bologna, Italy

Jakob Sachs (6)
Hamburg University of Technology, Germany

Laura Schulz (5)
Argonne National Laboratory, Lemont, IL, USA

Martin Schulz (5)
Technical University of Munich, Germany

Eva Sciacca (3)
INAF Astrophysical Observatory of Catania, Italy

Omar Shaaban (1)
Barcelona Supercomputing Center, Spain

Dimitrios Soudris (7, 8)
Microprocessors and Digital Systems Laboratory, ECE, NTUA, Athens, Greece

Despoina Tomkou (7)
Microprocessors and Digital Systems Laboratory, ECE, NTUA, Athens, Greece

17th Workshop on Parallel Programming and Run-Time Management Techniques for Many-Core Architectures and 15th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms (PARMA-DITAM 2026).
Editors: Davide Baroffio, Paola Busia, Lev Denisov, and Nitin Shukla
OpenAccess Series in Informatics
OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Nicola Tuccari (3)
Department of Mathematics and Informatics,
University of Catania, Italy

Sotirios Xydis (7, 8)
Microprocessors and Digital Systems Laboratory,
ECE, NTUA, Athens, Greece