

HORIZON
2020

Pilot using Independent Local & Open Technologies

Ergebnisse

Projektinformationen

The European PILOT

ID Finanzhilfevereinbarung: 101034126

[Projektwebsite](#) 

DOI

[10.3030/101034126](https://doi.org/10.3030/101034126) 

EK-Unterschriftsdatum

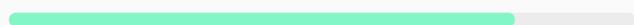
17 November 2021

Startdatum

1 Dezember 2021

Enddatum

28 Februar 2026



Finanziert unter

INDUSTRIAL LEADERSHIP - Leadership in enabling and industrial technologies - Information and Communication Technologies (ICT)

Gesamtkosten

€ 29 999 925,00

EU-Beitrag

€ 14 999 962,51

Koordiniert durch

BARCELONA
SUPERCOMPUTING CENTER
CENTRO NACIONAL DE
SUPERCOMPUTACION



Spain

CORDIS bietet Links zu öffentlichen Ergebnissen und Veröffentlichungen von HORIZON-Projekten.

Links zu Ergebnissen und Veröffentlichungen von RP7-Projekten sowie Links zu einigen Typen spezifischer Ergebnisse wie Datensätzen und Software werden dynamisch von [OpenAIRE](#)  abgerufen.

Leistungen

[Other \(2\)](#)



[Refined parallel Programming Interfaces](#)

D7.3. Refined parallel Programming Interfaces (FORTH, O, PU) [M30]. This document will elaborate on D7.1 based on the experience and results in D7.2. It will also provide results for the updated version of the implementations in D7.2.

[Compilation and Emulation infrastructure](#)

D9.1. Compilation and Emulation infrastructure (BSC, O, PU) [M9]. This deliverable will provide an updated version of the EPI compilation and Emulation infrastructure (Vehave) extended to support v1.0 of the RISC-V ISA. It will support C/C++ and will include automatic vectorization capabilities.

Documents, reports (10)

[Software and system integration specification and requirements](#)

D10.5 Software and system integration specification and requirements (BSC,R, PU) (M6). Working with a software/hardware co-design approach, this deliverable will provide architectural requirements extracted from the following applications: EC-Earth and GROMACS HPC applications, BLAS, FFTW and stencil numerical HPC kernels, image processing and deep learning applications as well as the molecular dynamics application utilizing both HPC and deep learning kernels.

[First Dissemination and Communication Report](#)

D2.2: First Dissemination and Communication Report (BSC, R, PU) [M12] This deliverable will report on the dissemination and communication activities of the project done in the first year.

[Collaboration roadmap and collaboration agreement with EUPEX](#)

D35 Collaboration roadmap and collaboration agreement with EUPEX EUPEX aims at delivering a largescale modular demonstrator based on the ARMbased general purpose processor design under development in EPI In contrast the European PILOT will deliver a demonstrator based on the RISC-V accelerators in EPI The European PILOT output could be integrated as an additional module into the EUPEX modular supercomputer For this reason we will define a collaboration roadmap between the two pilots to ensure the integration of the two projects into a global framework A joint Collaboration Agreement will be signed to that effect

[Parallel Programming Runtimes specifications](#)

D71 Parallel Programming Runtimes specifications BSC R PU M6 This deliverable will define the functionalities and interfaces that will have to be integrated in the Pilot Beyond the basic MPI and OpenMP support based in

MPICH and the LLVM OpenMP runtime it will include the TAMPI interface for improved interoperability between MPI and OpenMP resulting in more productive mechanisms to achieve communication/computation overlap Also the DLB interfaces to dynamically reassign cores between OpenMP threads in different processes The document will specify the fine grain resource management policies to be implemented by these runtimes within the processes and at the node level as well as the vertical interface to the coarser grain schedulers in WP5 It will also specify the optimizations to be implemented in the internals of the runtime like vectorization offloading to communication devices as well as mechanisms to be used to minimize the impact of noise OS communications in performance

[Design of AI frameworks for the Pilot platform](#)

D61 Design of AI frameworks for the Pilot platform ETH R PU M6 This deliverable will present the design of the AI frameworks ONNXDaCe TensorFlow Tarantella for accelerated ONNXDaCe TF and distributed Tarantella learning taking into account the requirements of the respective WP1 verticals

[Architecture specification and requirements for the MLS compute tile](#)

D10.2 Architecture specification and requirements for the MLS compute tile (ETH, R, PU) (M6). This deliverable will provide the architecture specification of the Machine Learning and Stencil (MLS) compute tile.

[Refined System Software components architecture](#)

D8.3. Refined System Software components architecture (CINI, R, PU) [M30]. This report will detail, based on the experience reported in D8.2. and the detailed knowledge of the actual design resulting from Stream 3 will refine the specification and architecture of the File system and Resource management high level system software components to be implemented on the actual Pilot.

[System Software components architecture](#)

D8.1. System Software components architecture (BSC, R, PU) [M6]. This deliverable will describe the functionality and interfaces of the File System and Global Resource Management components that will be provided to submit and execute workloads to the system. For general functionalities adopted from standards of existing developments, the document will refer to the appropriate external documentation. New features defined in this task (e.g. integration of the components) will be described. The document will also describe the global architecture of the implementation.

[Dissemination and Communication Plan](#)

D21 Dissemination and Communication Plan BSC R PU M3 This deliverable will set out the dissemination and communication strategy and the activities to be

undertaken to achieve it Results of the dissemination work will be reported in the periodic and final reports

[Project Management and Quality Guidelines](#) 

Open Research Data Pilot (1) 

[Data Management Plan](#) 

Veröffentlichungen

Conference proceedings (24) 

[Meet Monte Cimone: exploring RISC-V high performance compute clusters](#) 

Autoren: Federico Ficarelli, Andrea Bartolini, Emanuele Parisi, Francesco Beneventi, Francesco Barchi, Daniele Gregori, Fabrizio Magugliani, Marco Cicala, Cosimo Gianfreda, Daniele Cesarini, Andrea Acquaviva and Luca Benini

Veröffentlicht in: 2022

Herausgeber: ACM

DOI: 10.1145/3528416.3530869

[ControlPULP: A RISC-V Power Controller for HPC Processors with Parallel Control-Law Computation Acceleration](#) 

Autoren: Alessandro Ottaviano, Robert Balas, Giovanni Bambini, Antonio del Vecchio, Maicol Ciani, Davide Rossi, Luca Benini, Andrea Bartolini

Veröffentlicht in: Embedded Computer Systems: Architectures, Modeling, and Simulation - Lecture Notes in Computer Science (SAMOS22), 2022, Seite(n) 120-135, ISBN 978-3-031-15073-9

Herausgeber: Springer

DOI: 10.1007/978-3-031-15074-6_8

[STen: An Interface for Efficient Sparsity in PyTorch](#) 

Autoren: A. Ivanov, N. Dryden, T. Hoefler

Veröffentlicht in: Sparsity in Neural Networks workshop 2022, 2022

Herausgeber: ETH Zurich, Scalable Parallel Computing Laboratory

DOI: 10.48550/arxiv.2304.07613

[Experimenting with Emerging RISC-V Systems for Decentralised Machine Learning](#) 

Autoren: Gianluca Mittone, Nicolò Tonci, Robert Birke, Iacopo Colonnelli, Dariana Medić, Andrea Bartolini, Roberto Esposito, Emanuele Parisi, Francesco Beneventi, Mirko Polato, Massimo Torquati, Luca Benini, Marco Aldinucci

Veröffentlicht in: 2023

Herausgeber: ACM

DOI: 10.48550/arxiv.2302.07946

[Benchmarking Federated Learning Frameworks for Medical Imaging Tasks](#)

Autoren: Fonio, S.

Veröffentlicht in: Image Analysis and Processing - ICIAP 2023 Workshops. ICIAP 2023. Lecture Notes in Computer Science, 2024, ISBN 978-3-031-51026-7

Herausgeber: Springer Nature

DOI: 10.1007/978-3-031-51026-7_20

[The Italian research on HPC key technologies across EuroHPC](#)

Autoren: Marco Aldinucci, Giovanni Agosta, Antonio Andreini, Claudio A Ardagna, Andrea Bartolini, Alessandro Cilaro, Biagio Cosenza, Marco Danelutto, Roberto Esposito, William Fornaciari, Roberto Giorgi, Davide Lengani, Raffaele Montella, Mauro Olivieri, Sergio Saponara, Daniele Simoni, Massimo Torquati

Veröffentlicht in: 2021

Herausgeber: ACM

DOI: 10.1145/3457388.3458508

I/O-Optimal Cache-Oblivious Sparse Matrix-Sparse Matrix Multiplication

Autoren: Niels Gleinig, Maciej Besta, Torsten Hoefler

Veröffentlicht in: 36th IEEE International Parallel and Distributed Processing Symposium, 2022, ISBN 978-1-6654-8106-9

Herausgeber: 2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)

[The Red-Blue Pebble Game on Trees and DAGs with Large Input](#)

Autoren: Niels Gleinig, Torsten Hoefler

Veröffentlicht in: Structural Information and Communication Complexity. SIROCCO 2022, Lecture Notes in Computer Science, 2022, Seite(n) 135-153, ISBN 978-3-031-09992-2

Herausgeber: Springer, Cham

DOI: 10.1007/978-3-031-09993-9_8

[Specialization meets Flexibility: a Heterogeneous Architecture for High-Efficiency, High-flexibility AR/VR Processing](#)

Autoren: Arpan Suravi Prasad; Luca Benini; Francesco Conti

Veröffentlicht in: 2023 60th ACM/IEEE Design Automation Conference (DAC),

2023, ISBN 979-8-3503-2348-1

Herausgeber: IEEE

DOI: 10.1109/dac56929.2023.10247945

[Lifting C Semantics for Dataflow Optimization](#)

Autoren: Alexandru Calotoiu, Tal Ben-Nun, Grzegorz Kwasniewski, Johannes de Fine Licht, Timo Schneider, Philipp Schaad, Torsten Hoefler

Veröffentlicht in: ICS '22: Proceedings of the 36th ACM International Conference on Supercomputing, 2022

Herausgeber: ICS '22: Proceedings of the 36th ACM International Conference on Supercomputing

DOI: 10.1145/3524059.3532389

[A Federated Learning Benchmark for Drug-Target Interaction](#)

Autoren: Gianluca Mittone; Filip Svoboda; Marco Aldinucci; Nicholas Lane; Pietro Lió

Veröffentlicht in: In Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion), 2023, ISBN 978-1-4503-9419-2

Herausgeber: ACM DL

DOI: 10.1145/3543873.3587687

[MiniFloat-NN and ExSdotp: An ISA Extension and a Modular Open Hardware Unit for Low-Precision Training on RISC-V Cores](#)

Autoren: Bertaccini, Luca; Paulin, Gianna; Fischer, Tim; Mach, Stefan; Benini, Luca

Veröffentlicht in: 2022 IEEE 29th Symposium on Computer Arithmetic (ARITH), Ausgabe 5, 2022

Herausgeber: IEEE

DOI: 10.1109/arith54963.2022.00010

[Dallmi: Domain adaption for llm-based multi-label classifier](#)

Autoren: M. Betianu, A. Malan, M. Aldinucci, R. Birke, and L. Y. Chen

Veröffentlicht in: Advances in Knowledge Discovery and Data Mining - 28th Pacific-Asia Conference on Knowledge Discovery and Data Mining, PAKDD 2024, Lecture Notes in Computer Science, 2024, ISBN 978-981-97-2259-4

Herausgeber: Springer

DOI: 10.1007/978-981-97-2259-4_21

[Model-Agnostic Federated Learning](#)

Autoren: Gianluca Mittone; Walter Riviera; Iacopo Colonnelli; Robert Birke; Marco Aldinucci

Veröffentlicht in: Euro-Par 2023: Parallel Processing, Lecture Notes in Computer Science, Ausgabe 3, 2023, Seite(n) 383-396, ISBN 978-3-031-39698-4

Herausgeber: Springer Nature
DOI: 10.1007/978-3-031-39698-4_26

[Federated Learning meets HPC and cloud](#)

Autoren: Iacopo Colonnelli, Bruno Casella, Gianluca Mittone, Yasir Arfat, Barbara Cantalupo, Roberto Esposito, Alberto Riccardo Martinelli, Doriana Medic, Marco Aldinucci

Veröffentlicht in: 2022, Seite(n) 193-199, ISBN 978-3-031-34167-0

Herausgeber: Springer

DOI: 10.1007/978-3-031-34167-0_39

[AXI-Pack: Near-Memory Bus Packing for Bandwidth-Efficient Irregular Workloads](#)

Autoren: Chi Zhang; Paul Scheffler; Thomas Benz; Matteo Perotti; Luca Benini

Veröffentlicht in: 2023 Design, Automation & Test in Europe Conference & Exhibition (DATE), 2023, ISBN 979-8-3503-9624-9

Herausgeber: IEEE

DOI: 10.23919/date56975.2023.10137243

[Benchmarking FedAvg and FedCurv for Image Classification Tasks](#)

Autoren: Bruno Casella, Roberto Esposito, Carlo Cavazzoni, Marco Aldinucci

Veröffentlicht in: The 1st Italian Conference on Big Data and Data Science, 2022, Seite(n) 99-100

Herausgeber: CEUR-WS

DOI: 10.48550/arxiv.2303.17942

[Efficient Quantized Sparse Matrix Operations on Tensor Cores](#)

Autoren: Shigang Li; Kazuki Osawa; Torsten Hoeftler

Veröffentlicht in: SC '22: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, 2022, Seite(n) 1-15, ISBN 978-1-6654-5444-5

Herausgeber: IEEE

DOI: 10.1109/sc41404.2022.00042

[Exploring and Exploiting Data-Free Model Stealing](#)

Autoren: Chi Hong, Jiyue Huang, Robert Birke, Lydia Y. Chen

Veröffentlicht in: Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases 2023, 2023, Seite(n) 20-35, ISBN 978-3-031-43424-2

Herausgeber: Springer Nature

DOI: 10.1007/978-3-031-43424-2_2

[Boosting the Federation: Cross-Silo Federated Learning without Gradient Descent](#)

Autoren: Mirko Polato; Roberto Esposito; Marco Aldinucci

Veröffentlicht in: Proceedings of the International Joint Conference on Neural

Networks (IJCNN 2022), 2022, Seite(n) 1-10, ISBN 978-1-7281-8671-9

Herausgeber: IEEE

DOI: 10.1109/ijcnn55064.2022.9892284

[Efficient Direct Convolution Using Long SIMD Instructions](#) 

Autoren: Alexandre de Limas Santana; Adrià Armejach; Marc Casas

Veröffentlicht in: PPOPP '23: Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming, 2023, Seite(n) 342-353, ISBN 979-84-00-70015-6

Herausgeber: ACM DL

DOI: 10.1145/3572848.3577435

[Fast Arbitrary Precision Floating Point on FPGA](#) 

Autoren: Johannes de Fine Licht, Christopher A. Pattison, Alexandros Nikolaos Ziogas, David Simmons-Duffin, Torsten Hoefler

Veröffentlicht in: 2022 IEEE 30th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM), 2022

Herausgeber: 2022 IEEE 30th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)

DOI: 10.1109/fccm53951.2022.9786219

[Arax: A Runtime Framework for Decoupling Applications from Heterogeneous Accelerators](#) 

Autoren: Manos Pavlidakis, Stelios Mavridis, Antony Chazapis, Giorgos Vasiliadis, and Angelos Bilas.

Veröffentlicht in: SoCC '22: Proceedings of the 13th Symposium on Cloud Computing, 2022, ISBN 978-1-4503-9414-7

Herausgeber: Association for Computing Machinery

DOI: 10.1145/3542929.3563467

[A Data-Centric Optimization Framework for Machine Learning](#) 

Autoren: Oliver Rausch, Tal Ben-Nun, Nikoli Dryden, Andrei Ivanov, Shigang Li, Torsten Hoefler

Veröffentlicht in: ICS '22: Proceedings of the 36th ACM International Conference on Supercomputing, 2022, ISBN 978-1-4503-9281-5

Herausgeber: Association for Computing Machinery

DOI: 10.1145/3524059.3532364

Peer reviewed articles (9) 

[A Heterogeneous In-Memory Computing Cluster for Flexible End-to-End Inference of Real-World Deep Neural Networks](#) 

Autoren: Angelo Garofalo; Geethan Karunaratne; Francesco Conti; DAVIDE ROSSI; Irem Boybat; GIANMARCO OTTAVI; LUCA BENINI
Veröffentlicht in: IEEE Journal on Emerging and Selected Topics in Circuits and Systems, Ausgabe 1, 2022, ISSN 2156-3357
Herausgeber: IEEE Circuits and Systems Society
DOI: 10.1109/jetcas.2022.3170152

[Scalable Hierarchical Instruction Cache for Ultralow-Power Processors Clusters](#) 

Autoren: Jie Chen; Igor Loi; Eric Flamand; Giuseppe Tagliavini; Luca Benini; Davide Rossi
Veröffentlicht in: IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 31 (4), Ausgabe 8, 2023, ISSN 1557-9999
Herausgeber: IEEE
DOI: 10.1109/tvlsi.2022.3228336

[ControlPULP: A RISC-V On-Chip Parallel Power Controller for Many-Core HPC Processors with FPGA-Based Hardware-In-The-Loop Power and Thermal Emulation](#) 

Autoren: Ottaviano, A., Balas, R., Bambini, G. et al
Veröffentlicht in: International Journal of Parallel Programming, 2024, ISSN 1573-7640
Herausgeber: Springer
DOI: 10.1007/s10766-024-00761-4

[Dustin: A 16-Cores Parallel Ultra-Low-Power Cluster With 2b-to-32b Fully Flexible Bit-Precision and Vector Lockstep Execution Mode](#) 

Autoren: Gianmarco Ottavi; Angelo Garofalo; Giuseppe Tagliavini; Francesco Conti; Alfio Di Mauro; Luca Benini; Davide Rossi
Veröffentlicht in: IEEE Transactions on Circuits and Systems I: Regular Papers, 70 (6), Ausgabe 8, 2023, ISSN 1558-0806
Herausgeber: IEEE
DOI: 10.1109/tcsi.2023.3254810

[Darkside: A Heterogeneous RISC-V Compute Cluster for Extreme-Edge On-Chip DNN Inference and Training](#) 

Autoren: Angelo Garofalo; Yvan Tortorella; Matteo Perotti; Luca Valente; Alessandro Nadalini; Luca Benini; Davide Rossi; Francesco Conti
Veröffentlicht in: IEEE Open Journal of the Solid-State Circuits Society, Ausgabe 1, 2022, ISSN 2644-1349
Herausgeber: IEEE
DOI: 10.1109/ojsscs.2022.3210082

[Sparse Stream Semantic Registers: A Lightweight ISA Extension Accelerating General Sparse Linear Algebra](#) 

Autoren: Paul Scheffler; Florian Zaruba; Fabian Schuiki; Torsten Hoefler; Luca Benini

Veröffentlicht in: IEEE Transactions on Parallel and Distributed Systems, 2023, ISSN 1558-2183

Herausgeber: IEEE

DOI: 10.1109/tpds.2023.3322029

[MiniFloats on RISC-V Cores: ISA Extensions with Mixed-Precision Short Dot Products](#) 

Autoren: L. Bertaccini, G. Paulin, M. Cavalcante, T. Fischer, S. Mach and L. Benini

Veröffentlicht in: IEEE Transactions on Emerging Topics in Computing, 2024, ISSN 2168-6750

Herausgeber: IEEE Computer Society

DOI: 10.1109/tetc.2024.3365354

[FlooNoC: A Multi-Tb/s Wide NoC for Heterogeneous AXI4 Traffic](#) 

Autoren: Tim Fischer; Michael Rogenmoser; Matheus Cavalcante; Frank K. Gürkaynak; Luca Benini

Veröffentlicht in: IEEE Design&Test, 2023, ISSN 2168-2364

Herausgeber: IEEE

DOI: 10.3929/ethz-b-000638546

[A High-performance, Energy-efficient Modular DMA Engine Architecture](#) 

Autoren: Benz, Thomas, Rogenmoser, Michael, Scheffler, Paul, Riedel, Samuel, Ottaviano, Alessandro, Kurth, Andreas, Hoefler, Torsten, Benini, Luca

Veröffentlicht in: IEEE Transactions on Computers, 2024, ISSN 1557-9956

Herausgeber: IEEE

DOI: 10.1109/tc.2023.3329930

Other (1) 

[Programmatically Reaching the Roof: Automated BLIS Kernel Generator for SVE and RVV](#) 

Autoren: Stepan Nassyr, Kaveh Haghighi Mood, Andreas Herten

Veröffentlicht in: 2024

Herausgeber: Jülich Supercomputing Center (JSC)

DOI: 10.34734/fzj-2023-03437

Letzte Aktualisierung: 30 April 2025

Permalink: <https://cordis.europa.eu/project/id/101034126/results/de>

European Union, 2025

