ACM Computing Frontiers 2026 Technical Program

May 19-21, 2026 · Catania, Sicily, Italy · www.computingfrontiers.org

SPONSORS

Tuesday, May 19

08:30-09:00 — GENERAL - ENTRANCE HALL
Registration

09:00-09:20 — GENERAL - SALA PINACOTECA
Opening

09:20-10:20 — KEYNOTE - SALA PINACOTECA
Bridging the gap between software and hardware.
Lana Josipović, ETH Zurich, Switzerland

10:20-10:40 — BREAK - FIRST FLOOR TERRACE
Break (20m)

10:40-12:55 — MAIN SESSION - SALA PINACOTECA
RISC-V
Chair: Carsten Trinitis, TUM Heilbronn, Germany

- 10:40 - An Open-Source Distributed Simulation Framework for RISC-V Systems incorporating Vector and Cryptographic Extensions
- 11:05 - Threads or Vectors? Evaluating SPMD and Vector Accelerators for Resource Constrained RISC-V Architectures
- 11:30 - Short: Robust Quantum Communication for Space Systems Through an Heterogeneous RISC-V based FPGA/SoC Platform
- 11:40 - Implications of Supporting Compressed Instructions in Area-Optimized Bit-Serial RISC-V Cores
- 12:05 - Accuracy-Performance-Resources Trade-Offs in RISC-V microarchitectures for Genetic Programming
- 12:30 - SMX: A RISC-V ISA Extension for Scale-Adaptive Quantized Softmax

10:40-13:00 — WORKSHOP - SALA POLIFUNZIONALE
Sustainable Computing (SCW)
Chair: Chiara Sandionigi, CEA, France

- 10:40 - All LCA models are wrong. Are some of them useful? Towards open computational LCA in ICT
- 11:10 - Analysis of the Relationship Between Carbon Footprint and Mineral Resource Depletion in the Life Cycle Assessment of Digital Systems
- 11:35 - Influence of Environmental Indicators Aggregation Methods for Eco-designing Integrated Circuits
- 12:05 - Exploring Instruction Set Extension Emulation for Long-Term Support of RISC-V Processors
- 12:30 - Poulpe: a software-hardware toolchain to clone legacy circuit boards

13:00-14:00 — BREAK - FIRST FLOOR TERRACE
Lunch

14:00-15:15 — MAIN SESSION - SALA PINACOTECA
Best Papers
Session Chairs: Rubén Salvador, CentraleSupélec, IRISA, Inria, France Estela Suarez, Jülich Supercomputing Centre & University of Bonn, Germany

- 14:00 - StAccato: Reducing the Performance Penalty of Randomness
- 14:25 - FIIP: Flow-Based Instruction Processing for Out-of-Order Scheduling in GPGPUs
- 14:50 - LightSerial: Accelerating In-Process Isolation via Implicit Dependency Exposure

14:20-15:20 — WORKSHOP - SALA POLIFUNZIONALE
Malicious Software and Hardware in Internet of Things (Mal-IoT)

15:20-15:40 — BREAK - FIRST FLOOR TERRACE
Break (20m)

15:40-16:50 — MAIN SESSION - SALA PINACOTECA
Energy Efficiency and Sustainability
Chair: Josef Weidendorfer, TU-Dresden, Germany

- 15:40 - Compact Thermal and Power Models for Manycore Multichiplet Architectures: A Case Study on Intel Sapphire Rapids Processor
- 16:05 - Short: SeT-Diff: Towards Semantic Foundation Models for HPC Telemetry and Time-Series
- 16:15 - How Much Energy Is Wasted in LLM operations? Evidence from Kernel-Level DVFS
- 16:40 - Short: M3SA: Exploring Datacenter Performance and Climate-Impact with Multi- and Meta-Model Simulation and Analysis

15:40-17:00 — WORKSHOP - SALA POLIFUNZIONALE
Malicious Software and Hardware in Internet of Things (Mal-IoT)

17:00-17:45 — POSTER SESSION - FIRST FLOOR ROOM
Poster Blitz
Chair: Josef Weidendorfer, TU-Dresden, Germany

- 17:00 - Introduction
- 17:05 - A Reliable Multi-FPGA RISC-V Based Cluster for Space AI Inference
- 17:08 - A Chip-level Monitoring Framework for Enhancing PCIe Observability
- 17:11 - AI Act Sandboxes: A Safety-by-Co-Design Framework
- 17:14 - Bypassing Blocking Instructions to enable Out-Of-Order Execution in GPGPUs
- 17:17 - TinyLLM: Evaluation and Optimization of Small Language Models for Agentic Tasks on Edge Devices
- 17:20 - Security and Cost Trade-offs of Side-Channel Countermeasures for AES Software on RISC-V SoCs
- 17:23 - Toward Energy-Efficient Approximate Computing for Mixture-of-Experts Based CNN Inference
- 17:26 - Measuring What Matters: Advancing Green Computing Through the Green Work Efficiency (GWE) Metric
- 17:29 - Hardware Acceleration for Graph Neural Networks
- 17:32 - Enabling Fine-Grain DVFS for Multi-Kernel GPU Workloads
- 17:35 - Reinforcement Learning-based QoS-aware Online Scheduling for Multi-Tenant DNN Inference on Heterogeneous Accelerators
- 17:38 - Towards a RISC-V-based SmartNIC Architecture on FPGA
- 17:41 - Closing

18:45-20:00 — SOCIAL - FIRST FLOOR TERRACE
Welcome Reception

Wednesday, May 20

08:30-09:00 — GENERAL - SALA PINACOTECA
Registration

09:00-10:00 — MAIN SESSION - SALA PINACOTECA
Performance Optimization
Chair: Estela Suarez, Jülich Supercomputing Centre & University of Bonn, Germany

- 09:00 - NoCWalk: In-Network Page Walks for Efficient Pointer-Chasing Workloads on Multicores
- 09:25 - A QSP Matrix Multiplication Method for Phase Factor Recovery
- 09:50 - Short: Simulating MPI Collectives on Tofino Smart Switches in SimGrid

08:40-10:00 — WORKSHOP - SALA POLIFUNZIONALE
High Performance and Quantum Computing Integration (HPQCI)
Chair: Davide Ferrari, University of Parma, Italy

- 08:40 - Workshop start
- 08:45 - Keynote: Operating quantum hardware in an HPC environment: potential, challenges, lessons learnt
- 09:15 - Compiling Linear Algebra Workloads from C to Quantum Circuits via Multi Level Intermediate Representation
- 09:35 - Short: From FFT Factorization to QFT Gate Decomposition: A Structural and Computational Bridge
- 09:55 - Ending remarks

10:00-10:20 — BREAK - FIRST FLOOR TERRACE
Break (20m)

10:20-11:20 — KEYNOTE - SALA PINACOTECA
Making or Breaking AI: Grand Challenges, Opportunities and What We Might be Missing.
Eren Kurshan, Princeton University, USA

11:20-13:00 — MAIN SESSION - SALA PINACOTECA
FPGA
Chair: Rubén Salvador, CentraleSupélec, IRISA, Inria, France

- 11:20 - Efficient and Accurate Graph Classification with Hyperdimensional Computing on FPGA
- 11:45 - LSFAF: A Layer-Sharing and FPGA-Accelerated Framework for Fast Collaborative Inference in Edge Scenarios
- 12:10 - Network Folding for Resource-Efficient Implementation of Stream-Dataflow Deep Neural Network Inference on FPGAs
- 12:35 - D3OF: DRL-Driven Dual-Layer Dynamic Obfuscation Framework for FPGAs

11:20-13:00 — SPECIAL SESSION - SALA POLIFUNZIONALE
Collaborative Initiatives (1)
Chair: TBD

- 11:20 - Towards Massively Parallel HyperDimensional Computing Architectures for Intelligent Satellite Links
- 11:36 - RILKOSAN: Holistic Resilient Communication System for Industrial EnvironmentsIndustrial Environments
- 11:52 - EPAC: The Last Dance
- 12:08 - Verification and Validation (V&V)-in-Loop for RISC-V Chip Design: The Holistic Vision of BZL
- 12:24 - Matrix Extensions for the RISC-V ISA
- 12:29 - TURANDOT - TUning Risc-V to Automotive Needs & Different Other Theme
- 12:34 - SEANERGYS: Integrating Monitoring, Analytics, and Resource Management for Energy-efficient HPC System Operations
- 12:39 - SATUQ: Qantum-Ready Cybersecurity for Integrated Space–Aerial–Terrestrial Networks
- 12:44 - End-to-end codesign in AI-enabled computational science (ENCODE)
- 12:49 - Computing Quantum, Not Quantum Computing (DOE SciDAC)

ACM Computing Frontiers 2026 Technical Program

May 19-21, 2026 · Catania, Sicily, Italy · www.computingfrontiers.org

SPONSORS

Wednesday, May 20

13:00-14:00 — BREAK - FIRST FLOOR TERRACE
Lunch + Posters

14:00-15:35 — MAIN SESSION - SALA PINACOTECA
Accelerators
Chair: Antonino Tumeo, PNNL, USA

- 14:00 - SABRE: A Compression-Aware BF16 Accelerator for Neuromorphic Attention
- 14:25 - Short: A 8.4 TFLOPS@16b/4.3W General-Purpose Programmable Accelerated Cluster for AI-Native RAN
- 14:35 - AExec: Asynchronous Multi-accelerator Execution and Management Mechanism
- 15:00 - Short: A System-Level Performance Analysis of On-Device Learning on Ultra-Low-Power Edge Systems
- 15:10 - HMix : An Efficient Hardware Accelerator for Quantized MLP-Mixer Inference

14:00-15:40 — SPECIAL SESSION - SALA POLIFUNZIONALE
Collaborative Initiatives (2)
Chair: TBD

- 14:00 - LibreRT: Portable Heterogeneous Real-Time Programming for the Embedded Computing Continuum
- 14:16 - The Energy-Oriented Centre of Excellence
- 14:32 - Cognitive Orchestration of Distributed Resources in the ENACT Cognitive Computing Continuum (CCC)
- 14:48 - CAPE's Composable Server Infrastructure for the Edge-Cloud Continuum
- 15:04 - Advancements in Cryptographic Algorithms, Design Tools and Evaluation Mechanisms through the PROACT project
- 15:20 - From the NSF Center for High-performance Reconfigurable Computing (CHREC) to the NSF Center for Space, High-performance, and Resilient Computing (SHREC)

15:40-16:00 — BREAK - FIRST FLOOR TERRACE
Break (20m)

16:00-17:25 — MAIN SESSION - SALA PINACOTECA
Storage and Memory
Chair: Josef Weidendorfer, TU Dresden, Germany

- 16:00 - AME-PIM: Can Memory be Your Next Tensor Accelerator?
- 16:25 - Hardware-Constrained Online Coordination for Hybrid CXL-RDMA Disaggregated Memory
- 16:50 - Short: Strata: Proactive Page Placement in Hybrid Memory Systems
- 17:00 - Ada-Store: An Adaptive Load-aware Hybrid Storage Architecture for Bursty Workloads

17:40-18:40 — SPECIAL SESSION - SALA PINACOTECA
Analog
Kristian Rietveld, Leiden University, Netherlands

- 17:40 - PARADIGM: Programmable, Analog, and Reconfigurable Active Dendrites Implementing Gain Modulation
- 18:00 - Accelerating Sparse Linear Solvers with an Optical Laser Processing Unit
- 18:20 - CRN2DB: Chemical Reaction Networks to Database Matching via Dynamic Programming in ChemComp

16:00-18:40 — WORKSHOP - SALA POLIFUNZIONALE
Scalable Performance, Portability, and Productivity for Scientific Computing (SP4Sci)
Chair: Davide Gadioli

- 16:00 - Welcome
- 16:10 - Keynote: Beyond Molecular Docking: Toward Extreme-Scale Virtual Screening on European HPC Infrastructures
- 16:50 - On the Limits of Performance Portability in Directive-Based GPU Programming
- 17:10 - Evaluating the benefits of Argument Fusion: Optimizing host-device communication
- 17:30 - A Distributed Virtual Screening Miniapp for Performance Portability and Cross-Architecture Benchmarking
- 17:50 - Design and Implementation of a Prediction-Serving System for Runtime and Parallel Performance in Quantum ESPRESSO
- 18:10 - Enabling Portable Collective Communication on Heterogeneous GPU Systems
- 18:30 - Conclusions and remarks

19:30-20:00 — SOCIAL - ENTRANCE HALL
Walking tour

20:30-23:00 — SOCIAL - ENTRANCE HALL
Social event

Thursday, May 21

08:30-09:00 — GENERAL - ENTRANCE HALL
Registration

09:00-10:20 — MAIN SESSION - SALA PINACOTECA
Frontier Computing: Quantum and Edge
Chair: Biagio Cosenza, University of Salerno, Italy

- 09:00 - Topology and Reliability Aware Qubit Mapping for Quantum Core Systems
- 09:25 - Quantum Walks for Collision-Based Information Set Decoding
- 09:50 - Short: Similarity-Aware Function Pre-Loading for Serverless Inference
- 10:00 - Short: PrudentCaster: A Tunable Broadcast Gossip Framework for Mobile Edge Synchronization

09:00-10:15 — WORKSHOP - SALA POLIFUNZIONALE
Open-Source Hardware (OSHW) (1)
Chair: Alfonso Rodriguez, Universidad Politécnica de Madrid, Spain

- 09:00 - Keynote #1: From Open Source to Silicon: Reptiles and the Journey with Open Hardware Tape-Outs
- 09:30 - Hardware-Enforced Throughput Quotas for Mitigating Accelerator Interference in Mixed-Criticality SoCs
- 09:45 - An Ethernet-Integrated Accelerator for On-the-Fly Spherical-to-Cartesian LiDAR Coordinates Conversion
- 10:00 - Deterministic Co-Simulation of Open RISC-V-Based Cyber-Physical Systems: A Regenerative Suspension Case Study

10:20-10:40 — BREAK - FIRST FLOOR TERRACE
Break (20m)

10:40-13:05 — MAIN SESSION - SALA PINACOTECA
LLMs: Training, Performance Portability & Applications
Chair: Hubertus Franke, IBM Research, USA

- 10:40 - EnsembleHealer: Autonomous Recovery from Model Poisoning in Decentralized Federated Learning
- 11:05 - Row-wise Inter-Phase Pipelining for Hardware-Efficient GCN Acceleration
- 11:30 - MKA: Memory-Keyed Attention for Efficient Long-Context Reasoning
- 11:55 - A Portable GPU Kernel Performance Modeling Method Based on LLVM IR Dynamic Feature Prediction
- 12:20 - On the Efficacy of PyTorch for High-Performance Computing: A Case Study in Computational Physics
- 12:45 - Short: LLM-Driven Optimization for High-Level Synthesis
- 12:55 - Short: wdCP: Windowed Incremental Checkpointing for Efficient and Bounded LLM Recovery

10:45-12:00 — WORKSHOP - SALA POLIFUNZIONALE
Open-Source Hardware (OSHW) (2)
Chair: Angelo Garofalo, ETH Zurich, Switzerland

- 10:45 - Keynote #2: Can an Open ISA Reach the Datacenter? Lessons from Monte Cimone on RISC-V for HPC
- 11:15 - Evaluating the Impact of a Vector Co-Processor on a Memory System through Hybrid Simulation
- 11:30 - Not All Faults Are Equal: Transient-Fault Sensitivity Characterization of an Open-Source RISC-V Vector Cluster
- 11:45 - An Embedded RISC-V Vector Extension for Edge-Oriented Acceleration

12:00-13:00 — SPECIAL SESSION - SALA POLIFUNZIONALE
CompSpace
Chair: Carsten Trinitis, TUM Heilbronn, Germany

- 12:00 - Reliability Analysis of TMR Configurations in SRAM-Based FPGAs Using Static Evaluation
- 12:20 - Distributed Deep Learning Inference on DPU-Based FPGA Systems for On-Board Earth Observation
- 12:40 - On the Analysis of Radiation-Induced Faults in FPGA-Based Optical Communication for Space Applications

13:00-14:00 — BREAK - FIRST FLOOR TERRACE
Lunch

14:00-15:20 — SPECIAL SESSION - SALA PINACOTECA
Big RISC-V Made Little
Session Chairs: Marcel Baunach, Graz University of Technology, Austria Sebastian Prebeck, Infineon Technologies AG, Germany

- 14:00 - About Big and Little: Motivation and Overview of Session, Projects and Talks
- 14:20 - Advances in Instruction Compression and Memory Footprint Reduction
- 14:40 - Lessons Learned from Sub-32-Bit RISC-V Core Implementations
- 15:00 - Application of Tiny RISC-V

15:20-15:40 — BREAK - FIRST FLOOR TERRACE
Break (20m)

15:40-16:00 — GENERAL - SALA PINACOTECA
Award and Closing Session