



PROGRAMME OVERVIEW

PROGRAMME OVERVIEW

	MONDAY 28 August 2023	TUESDAY 29 August 2023
09:00 - 09:30	Workshops and Minisymposia AMTE, HETEROPAR, PECS , QuickPar, RAW	Workshops and Minisymposia PhD Symposium, ABUMPIMP, WSCC, Tutorial
09:30 - 10:00		
10:00 - 10:30		
10:30 - 11:00	COFFEE BREAK	
11:00 - 11:30	Workshops and Minisymposia AMTE, HETEROPAR, PECS , QuickPar, RAW	Workshops and Minisymposia PhD Symposium, ABUMPIMP, WSCC
11:30 - 12:00		
12:00 - 12:30		
12:30 - 14:00	LUNCH BREAK	
14:00 - 14:30	Workshops and Minisymposia AMTE, HETEROPAR, ADAPIO, RAW	Workshops and Minisymposia ABUMPIMP, TDLPP, WSCC
14:30 - 15:00		
15:00 - 15:30		
15:30 - 16:00	COFFEE BREAK	
16:00 - 16:30	Workshops and Minisymposia AMTE, HETEROPAR, ADAPIO, RAW	Workshops and Minisymposia ABUMPIMP, TDLPP
16:30 - 17:00		
17:00 - 17:30		
17:30 - 18:00		
18:00 - 18:30		
18:30 - 19:00		
19:00 - 20:30		Welcome Reception

	WEDNESDAY 30 August 2023	THURSDAY 31 August 2023	FRIDAY 01 September 2023
09:00 - 09:30	Plenary Session Keynote 1 Schahram Dustdar	Plenary Session Keynote 2 Enrique S. Quintana-Orti	Plenary Session Keynote 3 Jahna Otterbacher
09:30 - 10:00			
10:00 - 10:30	COFFEE BREAK		
10:30 - 11:00	Parallel Sessions TO1 / TO2 (A)	Plenary Session Candidates for Best Paper Award	Parallel Sessions TO6 (C) / TO2 (C)
11:00 - 11:30			
11:30 - 12:00			
12:00 - 12:30	LUNCH		
12:30 - 14:00	LUNCH		
14:00 - 14:30	Parallel Sessions TO6 (A) / TO4 (A)	Parallel Sessions TO6 (B) / TO5 (A)	Parallel Sessions TO3 (A) / TO6 (D)
14:30 - 15:00			
15:00 - 15:30			
15:30 - 16:00	COFFEE BREAK		
16:00 - 16:30	Parallel Sessions ERC / Posters and Demos Sessions	Parallel Sessions TO2 (B) / TO5 (B)	Parallel Sessions TO3 (B) / TO6 (B)
16:30 - 17:00			
17:00 - 17:30			
17:30 - 18:00		Conference Dinner	
18:00 - 18:30			
18:30 - 19:00			
19:00 - 23:00			

Monday August 28, 2023

08:55 - 09:00	Chair: Patrick Diehl
09:00 - 09:40	Opening remarks
09:40 - 10:05	Keynote Talk (Thomas Sterling, Simultac LLC)
10:05 - 10:30	Making Uintah Performance Portable for Department of Energy Exascale Testbeds (John Holmen, ORNL)
10:30 - 11:00	Malleable APGAS Programs and their Support in Batch Job Schedulers (Patrick Finnerty, Kobe University)
11:00 - 11:40	COFFEE BREAK
11:40 - 12:05	Chair: Steven R. Brandt
12:05 - 12:30	Invited talk (Brad Richardson, Berkeley Lab)
	Task-Level Checkpointing for Nested Fork-Join Programs using Work Stealing (Lukas Reitz, University of Kassel)
	Benchmarking the Parallel 1D Heat Equation Solver in Chapel, Charm++, C++, HPX, Go, Julia, Python, Rust, Swift, and Java (Patrick Diehl, LSU)

Monday August 28, 2023

09:00 - 09:05	Opening note: Tery Cojean , workshop co-chair
09:05 - 09:50	Session A Keynote: "Building software for the next generation of heterogeneous systems" Mehdi Goli VP Research & Development, Codeplay Software
09:50 - 10:15	"ExaNBody : a HPC framework for N-Body applications" Thierry Carrard, Raphael Prat, Guillaume Latu, Laurent Soulard, Paul Lafourcade, Killian Babilotte and Lhassan Amarsid
10:15 - 10:40	"Power Estimation Models for Edge Computing Devices" Michalis Kasioulis, Moysis Symeonides, George Pallis and Marios Dikaiakos
10:40 - 11:00	COFFEE BREAK
11:00 - 11:25	Session B "A performance analysis of leading many-core technologies for Cellular Automata execution" Alessio De Rango, Donato D'Ambrosio, Alfonso Senatore, Giuseppe Mendicino, Kumudha Narasimhan, Mehdi Goli and Rod Burns
11:25 - 11:50	"MassiveClicks: A Massively-parallel Framework for Efficient Click Models Training" Skip Thijssen, Pooya Khandel, Andrew Yates and Ana Lucia Varbanescu
11:50 - 12:15	"Tall-and-Skinny QR Factorization for Clusters of GPUs using High-Performance Building Blocks" Andres Tomas and Enrique S. Quintana-Orti
12:15 - 12:40	"Leveraging MLIR for Loop Vectorization and GPU Porting of FFT Libraries" Yifei He, Artur Podobas and Stefano Markidis
12:40 - 14:00	LUNCH BREAK
14:00 - 14:25	Session C "Boosting the Performance of Object Tracking with a Half-Precision Particle Filter on GPU" Gabin Schieffer, Nattawat Pornthisan, Jacob Wahlgren, Daniel Medeiros, Stefano Markidis and Ivy Peng
14:25 - 14:50	"jPackFaaS: Profiling Java Serverless Functions Package Size in Federated FaaS" Thomas Larcher and Sashko Ristov
14:50 - 15:15	"Enabling Dynamic Selection of Implementation Variants in Component-Based Parallel Programming for Heterogeneous Systems" Suejb Memeti
15:15 - 15:40	"Sparse matrix-vector product for the bmSparse matrix format in GPUs" Gonzalo Berger, Ernesto Dufrechou and Pablo Ezzatti
15:40 - 16:00	COFFEE BREAK
16:00 - 16:25	Session D "An approach to performance portability through generic programming" Andreas Hadjigeorgiou, Christodoulos Stylianou, Michele Weiland, Dirk Jacob Verschuur and Jacob Finkenrath
16:25 - 16:50	"Scheduling Fork-Joins to Heterogeneous Processors" Huijun Wang and Oliver Sinnen
16:50 - 17:00	WORKSHOP CLOSING: Tery Cojean , workshop co-chair

Monday August 28, 2023

09:15 - 09:20	Chairs' Welcome to PECS 2023
09:20 - 10:30	<p>Keynote Speech Boosting the performance and efficiency of software and hardware designs with Cache-aware Roofline Model Aleksandar Ilic, Universidade de Lisboa, INESC-ID, Lisbon</p> <p>Abstract: In the first part of this talk, we will introduce the Cache-aware Roofline Model (CARM) and expose its basic principles when modelling the performance, power consumption and energy-efficiency upper-bounds of a processor. In the second part of this talk, we will showcase the ability of the CARM to drive the optimization of epistasis detection, an important application in bioinformatics. For both CPU and GPU devices, we will demonstrate how the CARM can be used to detect execution bottlenecks and provide useful hints on which type of optimizations to apply in order to fully exploit device capabilities. We will also discuss our recent research contributions in CARM-driven RISC-V hardware design, as well as several model extensions to improve its insightfulness for specialised hardware and applications, such as NPUs and sparse computations.</p>
10:30 - 11:00	COFFEE BREAK
11:00 - 12:30	<p>Paper Presentation Session Session Chair: Romolo Marotta</p>
11:00 - 11:25	<p>Energy Efficiency Impact of Processing in Memory: A Comprehensive Review of Workloads on the UPMEM Architecture Yann Falevoz and Julien Legriel</p>
11:25 - 11:50	<p>Online job failure prediction in an HPC system Francesco Antici, Andrea Borghesi and Zeynep Kiziltan</p>
11:50 - 12:15	<p>Enhancing Supercomputer Performance with Malleable Job Scheduling Strategies Jonas Posner, Fabian Hupfeld and Patrick Finnerty</p>
12:15 - 12:20	Greeting Message from PECS 2023 Chairs


Monday August 28, 2023

09:00 - 09:10

OPENING

09:10 - 10:00

Keynote "Urgent Computing with AI@Edge"

Speaker: Pete Beckman, Northwestern University / Argonne National Labs

Abstract: From the sensor to the supercomputer, scientific discovery is part of a connected digital continuum that is dynamic and fast. In this new digital continuum, Artificial intelligence (AI) is providing tremendous breakthroughs, making data analysis and automated real-time responses possible across the digital continuum. Sage is a National Science Foundation funded project to build a national cyberinfrastructure for programmable edge computing. This new edge computing programming framework gives scientists a new tool for exploring the impacts of global urbanization, natural disasters such as flooding and wildfires, and climate change on natural ecosystems and city infrastructure. The Sage infrastructure allows scientists to analyze data in situ, at the edge, at the highest resolution of data. Data from the edge computation are then transmitted to a cloud computing infrastructure where they can be archived and provided to the community as data products or used in real-time to trigger computational models or dynamically modify subsequent edge computation. This new edge computing programming framework gives scientists a new tool for exploring the impacts of global urbanization, natural disasters such as flooding and wildfires, and climate change on natural ecosystems and city infrastructure. Sage is deploying cyberinfrastructure in environmental test-beds in California, Montana, Colorado, and Kansas, in the National Ecological Observatory Network, and in urban environments in Illinois and Texas.

10:00:10:30

"Streaming IoT Data and the Quantum Edge: A Classic/Quantum Machine Learning Use Case"

Authors: Sabrina Herbst, Vincenzo De Maio and Ivona Brandic

10:30 - 11:00

COFFEE BREAK

11:00 - 11:30

"Performance Optimization for Cross-Domain Intelligent IoT Applications"

Authors: Osama Almurshed, Souham Meshoul, Asmail Muftah, Ashish Kumar Kaushal, Osama Almoghamis, Ioan Petri, Nitin Auluck and Omer Rana

11:30 - 12:00

"Dynamic Adaptation of Urgent Applications in the Cloud-to-Edge Continuum"

Authors: Daniel Balouek-Thomert and Helene Coullon

12:00 - 12:10

CLOSING

Monday August 28, 2023

09:00 - 09:15	<p>OPENING</p> <p>Maja Hanne Kirkeby and Gordana Rakić</p>
09:15 - 10:00	<p>Invited talk: Unveiling the potential of the edge-cloud continuum: a deep dive into video streaming and graph processing use cases)</p> <p>Reza Farahani</p>
10:00 - 10:30	<p>Cormas: the software for participatory modelling and its application for managing natural resources in senegal</p> <p>Oleksandr Zaitsev, Francois Vendel And Etienne Delay</p>
10:30 - 11:00	COFFEE BREAK
11:00 - 11:30	<p>Towards resource-efficient dnn deployment for traffic object recognition: from edge to fog</p> <p>Dragan Stojanovic, Stefan Sentic And Natalija Stojanovic</p>
11:30 - 12:00	<p>Towards a simulation as a service platform For the cloud-to-things continuum</p> <p>Wilson Valdez, Hamza Baniata, Andras Markus And Attila Kertesz</p>
12:00 - 12:30	<p>The implementation of battery charging strategy for iot nodes</p> <p>Petar Rajkovic, Dejan Aleksić and Dragan Jankovic</p>
12:30 - 14:15	LUNCH BREAK
14:15 - 15:00	<p>Invited talk: Meta-heuristic search for the optimization of complex problems and its applications in the context of resource-aware development of cyber-physical systems</p> <p>Aitor Arrieta Marcos</p>
15:00 - 15:30	<p>Gpprmon: gpu runtime memory performance and power monitoring tool</p> <p>Burak Topçu And Işıl Öz</p>
15:30 - 16:00	COFFEE BREAK
16:00 - 16:30	<p>Analysis of approaches to the hardware implementation of a convolutional neural network</p> <p>Dmitry Pertsau And Mikhail Tatur</p>
16:30 - 17:00	<p>Submfl: compatible submodel generation for federated learning in device heterogeneous environment</p> <p>Zeyneddin Oz, Ceylan Soygul Oz, Abdollah Malekjafarian, Nima Afraz And Fatemeh Golpayegani</p>
17:00 - 17:30	<p>Performance and energy-aware training of a deep neural network in a multi-gpu environment with power capping</p> <p>Grzegorz Koszczał, Jan Dobrosolski, Mariusz Matuszek And Pawel Czarnul</p>
17:30	PANNEL DISCUSSION AND CLOSING


Monday August 28, 2023

14:00

OPENING [Jesus Carretero](#)

14:00 – 14:30

How dynamic workload may change HPC I/O for good
[Jean Baptiste Besnard](#). Paratools. France

14:30 – 15:00

Extending I/O POSIX semantics for data-streaming in scientific workflows with CAPIO
[Massimo Torquati](#). University of Pisa (CINI). Italy

15:00 – 15:30

ElastiSim: A Batch-System Simulator for Malleable Workloads
[Taylan Özden](#) & [Felix Wolf](#). Technical University of Darmstadt. Germany

15:30 – 16:00

COFFEE BREAK

16:00 – 16:30

Benefits of malleable ad hoc file systems for HPC applications
[Marc-André Vef](#). [Johannes Gutenberg](#) University Mainz, Germany

16:30 – 17:00

Computational malleability and environmental modeling: practice and experience
[Raffaele Montela](#). University of Naples “Parthenope” – Consorzio Interuniversitario Nazionale per Informatica. Italy

17:00 – 17:30

Exploring in-situ techniques and resource allocation on Quantum ESPRESSO
[Yi Ju](#). Max Planck Computing and Data Facility

17:30

CLOSING

Tuesday August 29, 2023

14:00 – 14:30

Leveraging HPC Profiling & Tracing Tools to Understand the Performance of Particle-in-Cell Monte Carlo Simulations

[Jeremy Williams](#), [David Tskhakaya](#), [Stefan Costea](#), [Ivy Peng](#), [Marta Garcia-Gasulla](#) and [Stefano Markidis](#)

14:30 – 15:00

Analyzing One-Sided Communication Using Memory Access Diagrams

[Olaf Krzikalla](#), [Arne Rempke](#) and [Ralph Mueller-Pfefferkorn](#)

15:00 – 15:30

Enhancing Performance Monitoring in C/C++ Programs with EDPM: A Domain-Specific Language for Performance Monitoring

[David Weisskopf Holmqvist](#) and [Suejb Memeti](#)

15:30 – 16:00

COFFEE BREAK

16:00 – 16:30

Performance Prediction for Sparse Matrix Vector Multiplication using Structure-dependent Features

[Konstantin Pogorelov](#), [James Trotter](#) and [Johannes Langguth](#)

16:30 – 17:00

Sparse-aware CARM: Rooflining locality of sparse computations

[Afonso Coutinho](#), [Diogo Marques](#), [Leonel Sousa](#) and [Aleksandar Ilic](#)

17:00 - 18:00

TOOL DEMO SESSION


Tuesday August 29, 2023

09:00 - 09:05	Session welcome and aims UPMEM
09:05 - 10:00	Keynote: UPMEM PIM platform for Data-Intensive Applications Yann FALEVOZ (UPMEM) /Julien LEGRIEL (UPMEM)
10:00 - 10:30	Invited talk: Understanding the potential of real processing-in-memory for modern workloads Juan GOMEZ LUNA (ETHZ)
10:30 - 11:00	COFFEE BREAK
11:00 - 11:30	Research paper: pimDB: From Main-Memory DBMS to Processing-In-Memory DBMS-Engines on Intelligent Memories Arthur BERNHARDT (Reutlingen University)
11:30 - 12:00	Invited talk: A Fast Processing-in-DIMM Join Algorithm Exploiting UPMEM DIMMs Chaemin LIM (Yonsei University)
12:00 - 12:30	Invited talk: PIM Performance and Economics for In-Memory Databases Hanna KRUPPE (SAP)
12:30 - 14:00	LUNCH BREAK
14:00 - 14:30	Research paper: Banded Dynamic Programming Algorithms on UPMEM PIM Architecture Meven MOGNOL (Univ. Rennes, CNRS-IRISA, Inria & UPMEM)
14:30 - 15:00	Research paper: Protein Alignment on UPMEM PIM Architecture Dominique LAVENIER (Univ. Rennes, CNRS-IRISA & Inria)
15:00 - 15:30	Research paper: An Experimental Evaluation of Machine Learning Training on a Real Processing-in-Memory System Juan GOMEZ LUNA (ETHZ) / Sylvain BROCARD (UPMEM)
15:30 - 16:00	COFFEE BREAK
16:00 - 16:30	Research paper: Implementation and Evaluation of Deep Neural Networks in Commercially Available Processing in Memory Hardware Purab Sutradhar (RIT)
16:30 - 17:00	Research paper: Privacy-Preserving Computing on UPMEM Elaheh Sadredini (UCR)
17:00 - 17:15	CLOSING UPMEM

Tuesday August 29, 2023

09:00 - 09:10	OPENING
09:10 - 10:00	Keynote: Compute continuum: the missing abstraction in data-intensive systems Alessandro Margara
10:00 - 10:30	Compute Continuum: What Lies Ahead? Matteo Nardelli , Gabriele Russo Russo , Valeria Cardellini
10:30 - 11:00	COFFEE BREAK
11:00 - 11:30	An Algorithm for Tunable Memory Compression of Time-Based Windows for Stream Aggregates, Vincenzo Gulisano
11:30 - 12:00	Latency-aware Placement of Stream Processing Operators Raphael Ecker , Vasileios Karagiannis , Michael Sober , Elmira Ebrahimi , Stefan Schulte
12:00 - 12:30	"Towards a Scalable Compute Continuum platform applied to electrical energy forecasting" Mohamad Moussa , Nabil Abdennadher , Raphaël Couturier , Giovanna Di Marzo Serugendo
12:30 - 14:00	LUNCH BREAK
14:00 - 14:30	"MTCL: a Multi-Transport Communication Library", Federico Finocchio , Nicolò Tonci , Massimo Torquati
14:30 - 15:00	"Scalable and Efficient Architecture for Random Forest on FPGA-based Edge Computing" Cuong Pham-Quoc
15:00 - 15:30	"Evaluation of Adaptive Micro-batching Techniques for GPU-accelerated Stream Processing" Ricardo Leonarczyk , Dalvan Griebler , Gabriele Mencagli , Marco Danelutto
15:30 - 15:40	CLOSING


Tuesday August 29, 2023

09:00 - 09:15

Welcome by the PhD Symposium Chairs
 Dr. Herodotos Herodotou and Dr. Demetris Trihinas

09:15 - 10:35

Session 1**Chair: Dr. Herodotos Herodotou**

Each presentation will be allocated 10min + 5min for Q/A

Exploring Mapping Strategies for Co-allocated HPC Applications

Ioannis Vardas, Sascha Hunold, Philippe Swartvagher and Jesper Larsson Träff

Online Job Failure Prediction in an HPC System

Francesco Antici, Andrea Borghesi and Zeynep Kiziltan

Parallel Auto-scheduling of Counting Queries in Machine Learning Applications on HPC Systems

Paweł Bratek, Lukasz Szustak and Jaroslaw Zola

Algorithm Selection of MPI Collectives Considering System Utilization

Majid Salimi Beni, Sascha Hunold and Biagio Cosenza

A Polynomial-time Algorithm for Detecting Potentially Unbounded Places in a Petri net-based Concurrent System

Marcin Wojnakowski, Remigiusz Wiśniewski and Mateusz Popławski

10.35 – 11.00

COFFEE BREAK

11.00 – 12.30

Session 2**Chair: Dr. Demetris Trihinas**

Massively Parallel EEG Algorithms for Pre-exascale Architectures

Zeyu Wang and Zoltan Juhasz

Service Management in Dynamic Edge Environments

Claudia Torres-Pérez, Estefanía Coronado, Cristina Cervelló-Pastor and Muhammad Shuaib Siddiqui

Efficiently Distributed Federated Learning

Gianluca Mittone, Robert Birke and Marco Aldinucci

Path Plan Optimization for UAV Assisted Data Collection in Large Areas

Ruben Enrique Padilla Robles and Rizos Sakellariou

Transitioning to Smart Sustainable Cities Based on Cutting-Edge Technological Improvements

Andreas Andreou and Constandinos Mavromoustakis

Data Assimilation with Ocean Models: A Case Study of Reduced Precision and Machine Learning in the Gulf of Mexico

Daniel Voss, Gary Tyson, Olmo Zavala-Romero, Alexandra Bozec and Ashwanth Srinivasan

12:30

CLOSING REMARKS by the PhD Symposium Chairs
 Dr. Herodotos Herodotou and Dr. Demetris Trihinas

Tuesday August 29, 2023



An Introduction To Sycl And The Celerity And Synergy Extensions

Biagio Cosenza

Abstract:

SYCL is a single-source, high-level, standard C++ programming model that can target a wide range of heterogeneous platforms equipped with CPUs, GPUs, FPGAs and more. The main advantage of using SYCL over other heterogeneous programming models is the single programming language approach, which allows you to target multiple devices with the same model, resulting in a cleaner, more portable and more readable code base. In addition, with the growing interest in heterogeneous hardware, moving to an open standard, platform-independent model such as SYCL is essential for today's software developers.

This tutorial provides an introduction to the SYCL programming model and the key language semantics, using a by-example approach. All major semantics of SYCL 2020 will be presented, including subgroups, kernel reductions, unified shared memory, task scheduling and data dependency. The second part of the tutorial presents two examples of SYCL-based frameworks that extend the SYCL semantics to provide additional functionality to the program. The first one is Celerity, which extends SYCL semantics such as the accessor to execute SYCL code on a distributed cluster of accelerators. The second one is SYnergy, which implements an energy-aware queue to enable energy-efficient techniques such as frequency scaling.

Attendees will be able to test the code presented during the tutorial on online compilers, so that they will leave the tutorial with an understanding of how to write SYCL programs, use advanced features, and understand the guidelines that enable performance portability across different platforms and architectures. Additional web-based resources are provided for follow-up studies.

Bio:

Biagio Cosenza is a tenure-track Assistant Professor in the Department of Computer Science at the University of Salerno, Italy, and a member of the Khronos SYCL Working Group. He joined the University of Salerno in August 2019 through a national brain gain program (Attraction and International Mobility), and received the Abilitazione for Italian Associate Professorship. From 2015 to 2019, he was Senior Research at the TU Berlin, Germany, where he was Principal Investigator for the DFG project Celerity and received the Habilitation from the Faculty IV. From 2011 to 2015, he was Postdoctoral researcher at the University of Innsbruck, Austria, where he contributed to the Insieme Compiler project and the DK-Plus multidisciplinary platform for Scientific Computing. His research is currently funded by the European HPC Joint Undertaking (LIGATE project), the Italian Ministry of Research (LibreRT project, PRIN 2022), and several industrial projects. Cosenza's main research interests are in the field of high performance computing, in particular with respect to programming models, compiler technology, optimization and tuning.



CONFERENCE PROGRAMME

Wednesday August 30, 2023

08:30 - 09:00	Registration	
09:00 - 09:30	ROOM: PANORAMA	
09:00 - 09:30	Plenary Session Keynote 1 Distributed Intelligence in the Computing Continuum Schahram Dustdar Session Chair: Marios Dikaiakos	
10:00 – 10:30	COFFEE BREAK	
	Parallel Sessions	
	ROOM: MEGARON AB	ROOM: PANORAMA
10:30 - 12:30	T01: Programming, Compilers, And Performance Session Co-Chairs: Biagio Cosenza Thomas Fahringer Extending OpenSHMEM with Aggregation Support for Improved Message Rate Performance Aaron Welch, Oscar Hernandez and Stephen Poole Fault-Aware Group-Collective Communication Creation and Repair in MPI Roberto Rocco and Gianluca Palerm perun: Benchmarking Energy Consumption of High-performance Computing Applications Juan Pedro Gutiérrez Hermsillo Muriedas, Katharina Flügel, Charlotte Debus, Holger Obermaier, Achim Streit and Markus Götz DIPPM: a Deep Learning Inference Performance Predictive Model using Graph Neural Networks Karthick Panner Selvam and Mats Brorsson	T02: Track 2: Run-time and OS scheduling Session Chair: Javier Garcia Blas How Do OS and Application Schedulers Interact? An Investigation with Multithreaded Applications Jonas H. Müller Korndörfer, Ahmed Eleliemy, Osman Seckin Simsek, Thomas Ilsche, Robert Schöne and Florina M. Ciorba Hierarchical management of extreme-scale task-based applications Francesc Lordan, Gabriel Puigdemunt, Pere Verges, Javier Conejero, Jorge Ejarque and Rosa M. Badia ODIN: Overcoming Dynamic Interference in Inference Pipelines Pirah Noor Soomro, Nikela Papadopoulou and Miquel Pericas INSTANT: a runtime framework to orchestrate in-situ workflows Feng Li and Fengguang Song
12:30 – 14:00	LUNCH	

Wednesday August 30, 2023

14:00 – 15:30	ROOM: MEGARON AB	ROOM: PANORAMA
	<p>T06: Track 6: GPU based applications (1) Session Chair: Francisco F. Rivera</p> <p>Im2win: An Efficient Convolution Paradigm on GPU Shuai Lu, Jun Chun, Luanzheng Guo and Xu T. Liu</p> <p>A Look at Performance and Scalability of the GPU Accelerated Sparse Linear System Solver Spliss Jasmin Mohnke and Michael Wagner</p> <p>Faster Segmented Sort on GPUs Robin Kobus, Johannes Nelgen, Valentin Henkys and Bertil Schmidt</p>	<p>T04: Track 4: Computational Science Session Chair: Pawel Czarnul</p> <p>Scalable Random Forest with Data-Parallel Computing Fernando Vázquez-Novoa, Javier Conejero, Cristian Tatu and Rosa M. Badia</p> <p>SymED: Adaptive and Online Symbolic Representation of Data on the Edge Daniel Hofstätter, Shashikant Ilager, Ivan Lujic and Ivona Brandic</p> <p>DeTAR: A Decision Tree-based Adaptive Routing in Networks-on-Chip Xiaoyun Zhang, Yaohua Wang, Dezun Dong, Cunlu Li, Shaocong Wang and Liqun Xiao</p>
15:30 – 16:00	COFFEE BREAK	
16:00 – 17:00	Parallel Sessions	
	ROOM: MEGARON AB	ROOM: PANORAMA
	<p>ERC funding opportunities Session Chair: Rizos Sakellariou</p> <p>Speakers: Mikaela Poulymenopoulou, European Research Council Executive Agency (ERCEA) Thomas Fahringer, University of Innsbruck, Austria</p> <p>In this session, Mikaela Poulymenopoulou, scientific officer in the Computer Science and Informatics panel, will present the ERC and give an overview of the available ERC funding schemes in Horizon Europe (Starting, Consolidator, Advanced Grants) and will describe the evaluation process of proposals, by peer review panels. This presentation will be complemented by the talk of Prof. Thomas Fahringer, who will share his ERC experience as Panel Member.</p>	<p>Posters and Demos Session Chair: George Pallis</p> <ol style="list-style-type: none"> OpenCUBE Open-source Cloud Services on EPI systems: Overview, Early Results, and Roadmap Gabin Schieffer, Jacob Wahlgren, Daniel Araújo de Medeiros, Ivy Peng, Martin Schultz, Utz-Uwe Haus, Craig Prunty, Pedro Marcuello and Emanuele Danovaro HEAppE Middleware: From desktop to HPC Jakub Konvička, Václav Svatoň and Jan Křenek Exploring Factors Impacting Data Offloading Performance in Edge and Cloud Environments Gap-Joo Na, Youngwoo Jang, Harin Seo, Byungchul Tak and Young-Kyoon Suh Dynamic Memory Allocation in Hierarchical Parallelism for Performance Portability Aymeric Millan, Thomas Padioleau and Julien Bigot Blockchain-based Decentralized Authority for Complex Organizational Structures Management Kamil Jarosz, Patrik Wojtyczek and Renata G. Stota

Wednesday August 30, 2023

17:00 – 17:30

ROOM: PANORAMA

6. BDDC preconditioning on GPUs for Cardiac Simulations
[Fritz Goebel, Terry Cojean and Hartwig Anzt](#)

7. A Layered Approach for Direct Convolution Using Matrix Multiplication Acceleration Extensions
[Victor Ferrari, Rafael Sousa, Marcio Pereira, João Carvalho, José Amaral, José Moreira and Guido Araujo](#)

8. Towards Energy-Aware Machine Learning in Geo-Distributed IoT Settings
[Demetris Trihinas and Lauritz Thamsen](#)

9. Designing A Sustainable Serverless Graph Processing Tool on the Computing Continuum
[Reza Farahani, Sashko Ristov and Radu Prodan](#)

10. High-Performance Distributed Computing with Smartphones
[Nadeem Ishikawa, Hayato Nomura, Yuya Yoda, Osamu Uetsuki, Keisuke Fukunaga, Seiji Nagoya, Junya Sawara, Hiroaki Ishihata and Junsuke Senoguchi](#)

11. Experiences and Lessons Learned from PHYSICS: A Framework for Cloud Development with FaaS
[Marta Patino](#)

12. MPI-based Remote OpenMP Offloading
[Illias K. Kasmeridis, Spyros Mantelos, Apostolos Piperis and Vassilios V. Dimakopoulos](#)

13. Improved IoT Application Placement in Fog through Postponement
[Aisha Aljohani and Rizos Sakellariou](#)

Demos:

1. Diorthotis: A Parallel Batch Evaluator for Programming Assignments
[Alexandros Karakasidis](#)

2. HEAppE Middleware: From desktop to HPC Pipelines
[Jakub Konvička, Václav Svatoň and Jan Křenek](#)

3. DAPHNE Runtime: Harnessing Parallelism for Integrated Data Analysis
[Aristotelis Vontzalidis, Stratos Psomadakis, Constantinos Bitsakos, Mark Dokter, Kevin Innerebner, Patrick Damme, Matthias Boehm, Florina M. Ciorba, Ahmed Eleliemy, Vasileios Karakostas, Aleš Zamuda and Dimitrios Tsoumakos](#)

Thursday August 31, 2023

08:30	Registration	
	ROOM: PANORAMA	
09:00 – 10:00	Plenary Session Keynote 2 A Continuum of Matrix Multiplications: From Scientific Computing to Deep Learning Enrique S. Quintana-Orti Session Chair: Chris Lengauer	
10:00 – 10:30	COFFEE BREAK	
	ROOM: PANORAMA	
10:30 – 12:30	Plenary Session Candidates for Best Paper Award Session Chair: Paul Kelly MMExit: Enabling Fast and Efficient Multi-modal DNN Inference with Adaptive Network Exits Xiaofeng Hou, Jiacheng Liu, Xuehan Tang, Chao Li, Kwang-Ting Cheng, Li Li and Minyi Guo Optimizing Data Movement for GPU-Based In-Situ Workflow Using GPUDirect RDMA Bo Zhang, Philip E Davis, Nicolas Morales, Zhao Zhang, Keita Teranishi and Manish Parashar Distributed k-Means with Outliers in General Metrics Enrico Dandolo, Andrea Pietracaprina and Geppino Pucci An efficient parallel adaptive GMG solver for large-scale Stokes problems S. Saberi, G. Meschke and A. Vogel	
12:30 – 14:00	LUNCH	
14:00 – 15:30	Parallel Sessions	
	ROOM: MEGARON AB	ROOM: PANORAMA
	T06: GPU based applications (2) Session Chair: Francisco F. Rivera GPU Code Generation of Cardiac Electrophysiology Simulation with MLIR Tiago Trevisan Jost, Arun Thangamani, Raphaël Colin, Vincent Loechner, Stéphane Genaud and Bérenger Bramas Parareal with a physics-informed neural network as coarse propagator Abdul Qadir Ibrahim, Sebastian Götschel and Daniel Ruprecht Accelerating Reduction Operation in AutoDock-GPU with Tensor Cores Gabin Schieffer and Ivy Bo Peng	T05: Graphs and Networks Session Chair: Chryssis Georgiou Distributed Deep Multilevel Graph Partitioning Peter Sanders and Daniel Seemaier On size hiding protocols in Beeping Model Dominik Bojko, Mateusz Marciniak, Marek Klonowski and Piotr Syga Efficient Protective Jamming in 2D SINR Networks Dominik Bojko, Mateusz Marciniak, Marek Klonowski and Dariusz Kowalski
15:30 – 16:00	COFFEE BREAK	

Thursday August 31, 2023

16:00 – 17:30

Parallel Sessions

ROOM: MEGARON AB

T02: Scheduling Algorithms

Session Chair:

A Poisson-Based Approximation
Algorithm for Stochastic Bin Packing of
Bernoulli Items

Tomasz Kanas and **Krzysztof Rządca**

Parametrized Analysis of a Dynamic
Programming Algorithm for a Parallel
Machine Scheduling Problem

İstenç Tarhan, **Jacques Carlier**, **Claire
Hanan**, **Antoine Jouglet** and **Alix
Munier Kordon**

Asymptotic Performance and Energy
Consumption of SLACK

Anne Benoit, **Louis-Claude Canon**,
Redouane Elghazi and **Pierre-Cyrille
Héam**

ROOM: PANORAMA

T05 : Parallel Processing and Scheduling

Session Chair: Christos Kaklamanis

A Parallel Scan Algorithm in the Tensor
Core Unit Model

Anastasios Zouzias and **William McColl**

Improved Algorithms for Monotone
Moldable Job Scheduling using
Compression and Convolution

Kilian Grage, **Klaus Jansen** and **Felix
Ohnesorge**

TrainBF: High-Performance DNN Training
Engine using BFloat16 on AI Accelerators

Zhen Xie, **Siddhisanket Raskar**, **Murali
Emani** and **Venkatram Vishwanath**

17:45 – 23:00

Tour/Conference Dinner
and Euro-Par 2023 Awards

Friday September 1, 2023

08:30	Registration	
	ROOM: PANORAMA	
09:00 – 10:00	Plenary Session Keynote 3 Bias in Data and Algorithms: Problems, Solutions and Stakeholders Jahna Otterbacher Session Chair: George A. Papadopoulos	
10:00 – 10:30		
10:30 – 12:30	Parallel Sessions	
	ROOM: MEGARON AB	ROOM: PANORAMA
	T06: HPC applications Session Chair: Francisco F. Rivera SWSPH: A Massively Parallel SPH Implementation for Hundred-Billion-Particle Simulation on New Sunway Supercomputer Ziyu Zhang , Junshi Chen, Zhanming Wang, Yifan Luo, Jineng Yao, Hong An and Shenghong Huang Hercules: scalable and network portable in-memory ad-hoc file system for data-centric and high-performance applications Javier Garcia Blas , Genaro Sanchez-Gallegos, Cosmin Petre, Alberto Riccardo Martinelli, Marco Aldinucci and Jesus Carretero Optimizing Distributed Tensor Contractions using Node-Aware Processor Grids Andreas Irmeler , Raghavendra Kanakagiri, Sebastian Ohlmann, Edgar Solomonik and Andreas Grüneis FedCML: Federated Clustering Mutual Learning Framework with non-IID Data Zekai Chen, Fuyi Wang, Shengxing Yu, Ximeng Liu and Zhiwei Zheng	T02: Resource management in the Cloud-Edge continuum Session Chair: Gianluca Mittone Geofence-Based Service Discovery in the Computing Continuum Kurt Horvath , Dragi Kimovski, Christoph Uran, Helmut Wöllik and Radu Prodan Assessing power needs to run a workload with quality of service on green datacenters Louis-Claude Canon, Damien Landré , Laurent Philippe, Jean-Marc Pierson and Paul Renaud-Goud SparkEdgeEmu: An Emulation Framework for Edge-enabled Apache Spark Deployments Moysis Symeonides , Demetris Trihinas, George Pallis and Marios D. Dikaiakos MetaLive: Meta-Reinforcement Learning Based Collective Bitrate Adaptation for Multi-Party Live Streaming Yi Yang , Xiang Li, Yeting Xu, Wenzhong Li, Jiangyi Hu, Taishan Xu, Xiancheng Ren and Sanglu Lu MetaLive:
12:30 – 14:00	LUNCH	

Friday September 1, 2023

14:00 – 15:30

Parallel Sessions

ROOM: MEGARON AB	ROOM: PANORAMA
<p>T03: Architectures and Accelerators (1) Session Chair: Jesus Carretero Perez</p> <p>Lock-free Bucketized Cuckoo Hashing Wenhai Li, Zhiling Cheng, Yuan Chen, Ao Li and Lingfeng Deng</p> <p>A Multi-level Parallel Integer/Floating-Point Arithmetic Architecture for Deep Learning Instructions Hongbing Tan, Jing Zhang, Libo Huang, Xiaowei He, Dezun Dong, Yongwen Wang and Liquan Xiao</p> <p>Computational Storage for an Energy-Efficient Deep Neural Network Training System Shiju Li, Kevin Tang, Jin Lim, Chul-Ho Lee and Jongryool Kim</p>	<p>T06: Various Session Chair: Francisco F. Rivera</p> <p>Transactional-Turn Causal Consistency Benoît Martin, Laurent Prosperi and Marc Shapiro</p> <p>Parallel Cholesky Factorization for Banded Matrices using OpenMP Tasks Felix Liu, Albin Fredriksson and Stefano Markidis</p> <p>Efficient Concurrent Transaction Execution Framework for Blockchains Manaswini Piduguralla, Saheli Chakraborty, Anjana Parwat Singh and Sathya Peri</p>

15:30 – 16:00

COFFEE BREAK

16:00 – 17:30

Parallel Sessions

ROOM: MEGARON AB	ROOM: PANORAMA
<p>T03: Architectures and Accelerators (2) Session Chair: Leonel Sousa</p> <p>Improving Utilization of Dataflow Architectures through Software and Hardware Co-Design Fan Zhihua, Wenming Li, Shengzhong Tang, Xuejun An, Xiaochun Ye and Dongrui Fan</p> <p>BitHist: A Precision-Scalable Sparse-Awareness DNN Accelerator Based on Bit Slices Products Histogram Zhaoteng Meng, Long Xiao, Xiaoyao Gao, Lin Shu and Jie Hao</p>	<p>T04: Data Analytics and AI Session Chair: Krzysztof Rzdca</p> <p>Model-Agnostic Federated Learning Gianluca Mittone, Walter Riviera, Iacopo Colonnelli, Robert Birke and Marco Aldinucci</p> <p>Auto-Divide GNN: Accelerating GNN Training with Subgraph Division Hongyu Chen, Zhejiang Ran, Keshi Ge, Zhiquan Lai, Jingfei Jiang and Dongsheng Li</p> <p>FedGM: Heterogeneous Federated Learning via Generative Learning and Mutual Distillation Chao Peng, Yiming Guo, Yao Chen, Qilin Rui, Zhengfeng Yang and Chenyang Xu</p>

SOCIAL EVENTS

Welcome Reception

Date: **29 AUG, 2023**

Time: **19:00 – 20:30**

Where: **Venue Hotel gardens**

Welcome Cocktail is the first social gathering between all conference delegates and it will take place at the Venue Hotel. It will be a relaxing evening during which delegates will have the opportunity to talk to colleagues and peers, while enjoying local drinks and ample canapés.

The Welcome Cocktail is included in All Registration fees except 'Workshops only'.

Ticket per accompanying person or "Workshop Only" participant: **€30.00**



Tour & Conference Dinner

Date: **31 AUG 2023**

Departure Time: **17:45**

Departure From: **The Venue Hotel**

We will get together at the lobby of the Venue Hotel, from where we will promptly depart in air-conditioned coaches for a city tour. A professional guide will tell us about the history of Cyprus and Limassol Town in particular. Dinner will take place at a restaurant serving scrumptious dishes of Cypriot cuisine complimented with local drinks and desserts. There will also be folkloric entertainamnt.

The Conference Dinner is included in All Registration fees except 'Workshops only'.

Ticket per accompanying person or "Workshop Only" participant: **€60.00**

