

Time	Wednesday 30 Aug, 2023	
08:30	Registration	
09:00 – 10:00	Plenary Session Keynote 1 <a href="#">Schahram Dustdar</a>	
10:00 – 10:30	Coffee Break	
10:30 – 12:30	Parallel Sessions	
	<p><b>T01: Programming, Compilers, And Performance</b></p> <p><b>176 Extending OpenSHMEM with Aggregation Support for Improved Message Rate Performance</b> <a href="#">Aaron Welch</a>, Oscar Hernandez and Stephen Poole</p> <p><b>204 Fault-Aware Group-Collective Communication Creation and Repair in MPI</b> <a href="#">Roberto Rocco</a> and Gianluca Palerm</p> <p><b>87 perun: Benchmarking Energy Consumption of High-performance Computing Applications</b> <a href="#">Juan Pedro Gutiérrez Hermosillo Muriedas</a>, Katharina Flügel, Charlotte Debus, Holger Obermaier, Achim Streit and Markus Götz</p> <p><b>82 DIPPM: a Deep Learning Inference Performance Predictive Model using Graph Neural Networks</b> <a href="#">Karthick Panner Selvam</a> and Mats Brorsson</p>	<p><b>T02: Track 2: Run-time and OS scheduling</b></p> <p><b>202 How Do OS and Application Schedulers Interact? An Investigation with Multithreaded Applications</b> <a href="#">Jonas H. Müller Korndörfer</a>, Ahmed Eleliemy, Osman Seckin Simsek, Thomas Ilsche, Robert Schöne and Florina M. Ciorba</p> <p><b>125 Hierarchical management of extreme-scale task-based applications</b> <a href="#">Francesc Lordan</a>, Gabriel Puigdemunt, Pere Verges, Javier Conejero, Jorge Ejarque and Rosa M. Badia</p> <p><b>147 ODIN: Overcoming Dynamic Interference in iNference pipelines</b> <a href="#">Pirah Noor Soomro</a>, Nikela Papadopoulou and Miquel Pericas</p> <p><b>170 INSTANT: a runtime framework to orchestrate in-situ workflows</b> Feng Li and <a href="#">Fengguang Song</a></p>
12:30 – 14:00	Lunch	
14:00 – 15:30	Parallel Sessions	
	<b>T06: Track 6: GPU based applications (1)</b>	<b>T04: Track 4: Computational Science</b>

	<p><b>75 Im2win: An Efficient Convolution Paradigm on GPU</b></p> <p><u>Shuai Lu</u>, Jun Chun, Luanzheng Guo and Xu T. Liu</p> <p><b>137 A Look at Performance and Scalability of the GPU Accelerated Sparse Linear System Solver Spliss</b></p> <p><u>Jasmin Mohnke</u> and Michael Wagner</p> <p><b>151 Faster Segmented Sort on GPUs</b></p> <p>Robin Kobus, Johannes Nelgen, Valentin Henkys and <u>Bertil Schmidt</u></p>	<p><b>91 Scalable Random Forest with Data-Parallel Computin</b></p> <p><u>Fernando Vázquez-Novoa</u>, Javier Conejero, Cristian Tatu and Rosa M. Badia</p> <p><b>140 SymED: Adaptive and Online Symbolic Representation of Data on the Edge</b></p> <p><u>Daniel Hofstätter</u>, Shashikant Ilager, Ivan Lujic and Ivona Brandic</p> <p><b>66 DeTAR: A Decision Tree-based Adaptive Routing in Networks-on-Chip</b></p> <p><u>Xiaoyun Zhang</u>, Yaohua Wang, Dezun Dong, Cunlu Li, Shaocong Wang and Liquan Xiao</p>
15:30 – 16:00	<b>Coffee Break</b>	
16:00 – 16:30	<b>Parallel Sessions</b>	
16:30 – 17:00	<p><b>ERC funding opportunities</b></p> <p>Speakers:</p>	<b>Posters</b>
17:00 – 17:30	<p><b>Mikaela Poulymenopoulou</b>, European Research Council Executive Agency (ERCEA)</p> <p><b>Thomas Fahringer</b>, 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><b>22 OpenCUBE Open-source Cloud Services on EPI systems: Overview, Early Results, and Roadmap</b></p> <p>Gabin Schieffer, Jacob Wahlgren, Daniel Araújo de Medeiros, Ivy Peng, Martin Schultz, Utz-Uwe Haus, Craig Prunty, Pedro Marcuello and Emanuele Danovaro</p> <p><b>24 HEAppE Middleware: From desktop to HPC</b></p> <p>Jakub Konvička, Václav Svatoň and Jan Křenek</p> <p><b>26 Exploring Factors Impacting Data Offloading Performance in Edge and Cloud Environments</b></p> <p>Gap-Joo Na, Youngwoo Jang, Harin Seo, Byungchul Tak and Young-Kyoon Suh</p> <p><b>27 Dynamic Memory Allocation in Hierarchical Parallelism for Performance Portability</b></p> <p>Aymeric Millan, Thomas Padioleau and Julien Bigot</p> <p><b>29 Blockchain-based Decentralized Authority for Complex Organizational Structures Management</b></p> <p>Kamil Jarosz, Patryk Wojtyczek and Renata G. Słota.</p> <p><b>35 BDDC preconditioning on GPUs for Cardiac Simulations</b></p> <p>Fritz Goebel, Terry Cojean and Hartwig Anzt</p>

**38 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

**42 Towards Energy-Aware Machine Learning in Geo-Distributed IoT Settings**

Demetris Trihinas and Lauritz Thamsen

**46 Designing A Sustainable Serverless Graph Processing Tool on the Computing Continuum**

Reza Farahani, Sashko Ristov and Radu Prodan

**55 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

**56 Experiences and Lessons Learned from PHYSICS: A Framework for Cloud Development with FaaS**

Marta Patino

**57 MPI-based Remote OpenMP Offloading**

Ilias K. Kasmeridis, Spyros Mantelos, Apostolos Piperis and Vassilios V. Dimakopoulos

**58 Improved IoT Application Placement in Fog through Postponement**

Aisha Aljohani and Rizos Sakellariou

**Demos:**

**15 Diorthotis: A Parallel Batch Evaluator for Programming Assignments**

Alexandros Karakasidis

**25 HEAppE Middleware: From desktop to HPC**

Pipelines Jakub Konvička, Václav Svatoň and Jan Křenek

**52 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