



**March 11-13, 2020
Västerås, Sweden**

Conference Program

PDP 2020

**28th Euromicro International Conference on Parallel,
Distributed and Network Based Processing**



PDP 2020

28th Euromicro International Conference on
Parallel, Distributed and Network-based Processing

CONFERENCE PROGRAM

Overview

Wednesday 11th of March 2020

9:30 – 10:30 **Keynote: Schahram Dustdar (Edge Intelligence: The Co-evolution of Humans, IoT, and AI)**

10:30 – 10:45 **Break**

10:45 – 12:00

Session 1.1: PDP: GPU kernels

Session 1.2: GMICC

12:00 – 13:00 **Lunch**

13:00 – 14:30

Session 2.1: PDP: Accelerators

Session 2.2: GMICC

14:30 – 15:00 **Coffee break**

15:00 – 16:30

Session 3.1: PDP: Distributed Systems

Session 3.2: AHPBB + HPC4NS

Thursday 12th of March 2020

9:30 – 10:30 Keynote: Hadi Esmaeilzadeh (AI for Accelerated AI Execution)

10:30 – 10:45 Break

10:45 – 12:00

Session 4.1: PDP: Cloud

Session 4.2: OPNS

Session 4.3: HPCM

12:00 – 13:00 Lunch

13:00 – 14:30

Session 5.1: PDP: Machine learning, deep learning

Session 5.2: OPNS

Session 5.3: HPCM

14:30 – 15:00 Coffee break

15:00 – 16:45

Session 6.1: PDP: Performance vs. Energy Efficiency

Session 6.2: OPNS

Session 6.3: HPCM

Friday 13th of March 2020

9:30 – 10:30

Session 7.1: PDP: Protocols (classification and communication)

Session 7.2: CISA + EMPS

10:30 – 10:45 Break

10:45 – 11:30

Session 8.1: PDP: Embedded systems

Session 8.2: SPDNBC

Session 8.3: Work in progress

12:00 – 13:00 Lunch

13:00 – 14:30

Session 9.1: PDP: Parallelism

Session 9.2: SPDNBC

CONFERENCE PROGRAM

Extended

Wednesday 11th of March 2020

9:30 – 10:30 Keynote: Schahram Dustdar (Room: 1)
Edge Intelligence: The Co-evolution of Humans, IoT, and AI

10:30 – 10:45 Break

10:45 – 12:00

Session 1.1: PDP: GPU kernels

[Room: 1]

- XKBlas: a high performance implementation of BLAS-3 kernels on multi-GPU server
Thierry Gautier and João Vicente Ferreira Lima
- Efficient NAS Parallel Benchmark Kernels with CUDA
Gabriell Araujo, Dalvan Griebler, Marco Danelutto and Luiz Gustavo Fernandes
- Using Parallel Programming Models for Automotive Workloads on Heterogeneous Systems - a Case Study (short)
Lukas Sommer, Florian Stock, Leonardo Solis-Vasquez and Andreas Koch

Session 1.2: GMICC

[Room: 2]

- Lessons learned from Comparing C-CUDA and Python-Numba for GPU-Computing
Lena Oden
- CoopCL: Cooperative Execution of OpenCL Programs on Heterogeneous CPU-GPU Platforms
Konrad Moren and Diana Goehringer
- A GPU enhanced LIDAR Perception System for Autonomous Vehicles (short)
Abderrahim Haneche, Mohammed Yazid Lachachi, Hamza Ouarnoughi and Smail Niar

12:00 – 13:00 Lunch

13:00 – 14:30

Session 2.1: PDP: Accelerators

[Room: 1]

- Parallel Comparison of Huge DNA Sequences in Multiple GPUs with Pruning
Marco Figueiredo Jr, Alba Cristina M. A. Melo and George Teodoro
- Implementation and Optimization of a 1D2V PIC Method for Nonlinear Kinetic Models on GPUs
Matthias Korch, Philipp Raithel and Tim Werner
- The Non-Uniform Compute Device (NUCD) Architecture for Lightweight Accelerator Offload
Mochamad Asri, Curtis Dunham, Roxana Rusitoru, Andreas Gerstlauer and Jonathan Beard

Session 2.2: GMICC

[Room: 2]

- Performance and efficiency investigations of SIMD programs of Coulomb solvers on multi- and many-core systems with vector units
Ronny Kramer and Gudula Rünger,
- Parallelization of Variable Rate Decompression through Metadata
Lennart Noordsij, Steven van der Vlugt, Mohamed Bamakhrama, Peter Lindstrom and Zaid Al-Ars,
- Heuristic Algorithms with Near Optimal Broadcast Time in Cactus Graphs (short)
Neil Conlan, Hovhannes Harutyunyan and Edward Maraachlian

14:30 – 15:00 Coffee break

15:00 – 16:30

Session 3.1: PDP: Distributed Systems

[Room: 1]

- Inter-Server RSS: Extending Receive Side Scaling for Inter-Server Workload Distribution
Andreas Oeldemann, Franz Biersack, Thomas Wild and Andreas Herkersdorf
- DSPG: Decentralized Simultaneous Perturbations Gradient Descent Scheme
Arunselvan Ramaswamy
- Performance Meets Programmability: Enabling Native Python MPI Tasks In PyCOMPSs (short)
Hatem Elshazly, Francesc Lordan, Jorge Ejarque and Rosa M. Badia
- Accurate Contention Estimate Scheduling Method Using Multiple Clusters of Many-core Platform for Applications Requiring Parallel Computations (short)
Shingo Igarashi, Yuto Kitagawa, Takuro Fukunaga and Takuya Azumi

Session 3.2: AHPBB + HPC4NS

[Room: 2]

- Accelerating Human Genome Phenotypic Analysis with Bitwise Search and Batched Computation
Yuichiro Miyamoto, Masao Okita and Fumihiko Ino
- Implementation of Syncytial Models in NEURON Simulator for Improved Efficiency
Shailesh Appukkuttan, Darshan Mandge and Rohit Manchanda

Thursday 12th of March 2020

9:30 – 10:30 Keynote: Hadi Esmaeilzadeh (Room: 1)
AI for Accelerated AI Execution

10:30 – 10:45 Break

10:45 – 12:00

Session 4.1: PDP: Cloud

[Room: 1]

- M3AT: Monitoring Agents Assignment Model for Data-Intensive Applications
Vladislav Kashansky, Dragi Kimovski, Radu Prodan, Prateek Agrawal, Fabrizio Marozzo, Iuhasz Gabriel, Marek Justyna and Javier Garcia-Blas
- An Interference-Aware Application Classifier Based on Machine Learning to Improve Scheduling in Clouds
Vinicius Meyer, Dionatrã F. Kirchoff, Matheus L. da Silva and César A. F. De Rose
- Modeling and Simulation of QoS-Aware Power Budgeting in Cloud Data Centers (short)
Jakub Krzywda, Vinicius Meyer, Miguel Xavier, Ahmed Ali-Eldin, P-O Östberg, Cesar De Rose and Erik Elmroth

Session 4.2: OPNS

[Room: 2]

- Efficient On-Chip Multicast Routing based on Dynamic Partition Merging
Binayak Tiwari, Mei Yang, Yingtao Jiang and Xiaohang Wang
- Off-Chip Congestion Management in GPU-based Non-Uniform Processing-in-Memory Networks
Kishore Punniyamurthy and Andreas Gerstlauer
- Performance study of HPC applications on an Arm-based cluster using a generic efficiency model
Fabio Banchelli, Kilian Peiro, Andrea Querol, Guillem Ramirez-Gargallo, Guillem Ramirez-Miranda, Joan Vinyals, Pablo Vizcaino, Marta Garcia-Gasulla and Filippo Mantovani.

Session 4.3: HPCM

[Room: 3]

- Optimizing Cellular Automata Execution by Distributed Discrete Event Simulation Techniques
Andrea Giordano, Alessio De Rango, Donato D'Ambrosio, William Spataro and Rocco Rongo,
- Elastic and Real-time Capacity Planning for Web Search Engines
Veronica Gil Costa, Mauricio Marin and Alonso Inostroza Psijas
- Efficient Wavefront Parallel Processing for HEVC CABAC Decoding (short)
Philipp Habermann, Chi Ching Chi, Mauricio Alvarez-Mesa and Ben Juurlink

12:00 – 13:00 Lunch

13:00 – 14:30

Session 5.1: PDP: Machine learning, deep learning

[Room: 1]

- Adaptive Load Balancing based on Machine Learning for Iterative Parallel Applications
Anna Victoria C. R. Oikawa, Vinicius Freitas, Márcio Castro and Laércio L. Pilla
- Accelerating Deep Learning using Multiple GPUs and FPGA-Based 10GbE Switch
Tomoya Itsubo, Michihiro Koibuchi, Hideharu Amano and Hiroki Matsutani
- On the Resilience of Deep Learning for Low-voltage FPGAs
Kamyar Givaki, Behzad Salami, Reza Hojabr, Sm. Reza Tayaranian, Ahmad Khonsari, Dara Rahmati, Saeed Gorgin, Adrian Cristal and Osman Unsal

Session 5.2: OPNS

[Room: 2]

- Selective Caching: Avoiding Performance Valleys in Massively Parallel Architectures
Amin Jadidi, Chita Das and Mahmut Kandemir
- NoC Design Methodologies for Heterogeneous Architecture
Lulwah A M J Alhubail, Masoomeh Jasemi and Nader Bagherzadeh

Session 5.3: HPCM

[Room: 3]

- A p2p environment to validate ensemble-based approaches in the cybersecurity domain
Francesco Folino, Gianluigi Folino and Luigi Pontieri
- Recent Trends in Modelling and Simulation with Machine Learning
Giuseppe A. Trunfio
- Natural language processing approach for distributed health data management (short)
Agostino Forestiero and Giuseppe Papuzzo
- Health Data Information Retrieval For Improved Simulation (short)
Mario Ciampi, Giuseppe De Pietro, Elio Masciari and Stefano Silvestri

14:30 – 15:00 Coffee break

15:00 – 16:45

Session 6.1: PDP: Performance vs Energy Efficiency

[Room: 1]

- Robustness and Energy-elasticity of Crown Schedules for Sets of Parallelizable Tasks on Many-core Systems with DVFS
Christoph Kessler, Sebastian Litzinger and Jörg Keller
- Decreasing the Learning Cost of Offline Parallel Application Optimization Strategies
Gustavo Berned, Fábio Rossi, Marcelo Luizelli, Arthur Lorenzon and Antonio Carlos Schneider Beck.
- Maximizing Profit in Energy-Efficient Moldable Task Execution with Deadline (short)
Sebastian Litzinger, Jörg Keller and Christoph Kessler
- Voltage Island-Aware Energy-Efficient Scheduling of Parallel Streaming Tasks on Many-Core Processors (short)
Nicolas Melot, Christoph Kessler and Jörg Keller

Session 6.2: OPNS

[Room: 2]

- Temperature-Aware Core Mapping for Heterogeneous 3D NoC Design Through Constraint Programming
Ayhan Demiriz, Hamzeh Ahangari and Özcan Öztürk
- RSMCC: Enabling Ring-based Software Managed Cache-Coherent Embedded SoCs (short)
Georgios Kornaros
- Switching at flit level: A Congestion Efficient Flow Control Strategy for Network-on-Chip (short)
Avik Bose and Prasun Ghosal
- Parallelization-Aware Pruning for Dense Neural Networks (short)
Sina Shahhosseini, Ahmad Albaqsami, Masoomeh Jasemi and Nader Bagherzadeh. Partition Pruning

Session 6.3: HPCM

[Room: 3]

- An efficient and scalable SPARK methodology to mine Association Rule from Genome Wide Association Studies
Giuseppe Agapito, Pietro Hiram Guzzi and Mario Cannataro
- Practical Cost-Effective Parallelization of Scientific Applications
Valentina Cesare, Iacopo Colonnelli and Marco Aldinucci,
- Acceleration of Radiofrequency Ablation Process for Liver Cancer Using GPU (short)
Claudio Schepke and Marcelo Cogo Miletto
- Scalable Parallel Genetic Algorithm For Solving Large Integer Linear Programming Models Derived From Behavioral Synthesis (short)
Mohammad K Fallah, Mina Mirhosseini, Mahmood Fazlali and Masoud Daneshlab
- Automation of High-Fidelity CFD Analysis for Aircraft Design and Optimization Aided by HPC (short)
Mengmeng Zhang, Jing Gong, Lilit Axner and Michaela Barth

Friday 13th of March 2020

9:30 – 10:30

Session 7.1: PDP: Protocols (classification and communication)

[Room: 1]

- TLB-based Block-Grain Classification of Private Data
Bhargavi R. Upadhyay, Alberto Ros and N S Murty
- Enhancing Two Phase-Commit Protocol for Replicated State Machines (short)
Tolga Ovatman and Halit Uyanik

Session 7.2: CISA + EMPS

[Room: 2]

- Windsurfing with APPA: Automating the Computational Fluid Dynamics Simulations of Wind Flow using the Cloud Computing
Anshul Jindal, Benedikt Strahm, Vladimir Podolskiy and Michael Gerndt
- Container Anomaly Detection Using Neural Networks Analyzing System Calls (short)
Christoph Reich, Holger Gantikow and Tom Zöhner
- Evaluating the Energy Efficiency of OpenCL-accelerated AutoDock Molecular Docking (short)
Leonardo Solis-Vasquez, Diogo Santos-Martins

10:30 – 10:45 Break

10:45 – 11:30

Session 8.1: PDP: Embedded systems

[Room: 1]

- Multi-level Binarized LSTM in EEG Classification for Wearable Devices
Najmeh Nazari, Seyed Ahmad Mirsalari, Sima Sinaei, Mostafa Ersali Salehi and Masoud Daneshtalab
- Mapping Method of MATLAB/Simulink Model for Embedded Many-Core Platform (short)
Kentaro Honda, Sasuga Kojima, Hiroshi Fujimoto, Masato Edahiro and Takuya Azumi

Session 8.2: SPDBNC

[Room: 2]

- SEPAD – Security Evaluation Platform for Autonomous Driving
Daniel Zelle, Roland Rieke, Christian Plappert, Christoph Krauß, Dmitry Levshun and Andrey Chechulin
- Augmented reality for visualizing security data for cybernetic and cyberphysical systems
Maxim Kolomeets, Andrey Chechulin, Ksenia Zhernova, Igor Kotenko and Diana Gaifulina
- Optimizing Secure Information Interaction in Distributed Computing Systems by the Sequential Concessions Method (short)
Igor Kotenko, Yury Sineshchuk and Igor Saenko

Session 8.3: Work in progress

[Room: 3]

- Embracing Heterogeneity for Exascale Computing: The EXA2PRO High-Level Programming Model
Christoph Kessler, August Ernstsson, Suejb Memeti, Johan Ahlqvist
- Model-Based Design Automation for High-Performance Acquisition and Control Systems on Heterogeneous Hardware
Rens Baeyens, Jan Steckel, Walter Daems
- HPC-Enabled Aerodynamic Optimization Studies Using CFD and Design Suite SU2
Mengmeng Zhang, Jing Gong, Lilit Axner

12:00 – 13:00 Lunch

13:00 – 14:30

Session 9.1: PDP Parallelism

[Room: 1]

- A Hybrid Approach to Parallel Pattern Discovery in C++, Christopher Brown, Murray Cole (short)
Adam Barwell, Jose Daniel Garcia, Vladimir Janjic, David Del Rio Astorga, Roberto Castaneda Lozano, Bjoern Franke and John Thomson
- Automatic Placement of Tasks to NUMA Nodes in Iterative Applications (short)
Jiri Dokulil and Siegfried Benkner
- Recursive Task Generation for Scalable SDF Graph Execution on Multicore Processors (short)
Georgios Georgakarakos and Johan Lilius
- Evaluating the performance and improving the usability of parallel and distributed Word Embeddings tools(short)
Matheus L. da Silva, Vinicius Meyer, Dionatrã F. Kirchoff, Joaquim Santos, Renata Vieira and Cesar A. F. De Rose
- Modeling and Simulating Daily Power Budgets for Sustainable Data Centers (short)
Rumenigie Hohemberger, Arthur Lorenzon, Marcelo Luizelli and Fábio Rossi
- F2MH Cryptosystem: Preliminary analysis of an original attempt to revive Knapsack-based public-key encryption schemes (short)
Yuri Villas Boas, Jean Martina and Charles Barros

Session 9.2: SPDBNC

[Room: 2]

- A Trusted Agent Strategy in Decentralized Network Environments
Karishma Karishma and Shrisha Rao
- Decidability of Deterministic Process Equivalence for Finitary Deduction Systems (short)
Yannick Chevalier and Benito Fabian Romero Jimenez