

EuroPar 2021 Program

Time				Conference Program		
New York (EDT)	Lisbon (WEST)	Brussels (CEST)	Beijing (CST)	SEPTEMBER 1	SEPTEMBER 2	SEPTEMBER 3
4:30 AM	9:30 AM	10:30 AM	4:30 PM		Regular Papers 4: Scheduling & Load Balancing I	EU Projects: European Initiative Projects Towards Exascale Computing
5:00 AM	10:00 AM	11:00 AM	5:00 PM			
5:30 AM	10:30 AM	11:30 AM	5:30 PM	Opening Session		
6:00 AM	11:00 AM	12 Noon	6:00 PM	Regular Papers 1: Parallel Methods & Applications	Regular Papers 5: Scheduling & Load Balancing II	Regular Papers 8: Tools & Environments
6:30 AM	11:30 AM	12:30 PM	6:30 PM			
7:00 AM	12 Noon	1:00 PM	7:00 PM			
7:30 AM	12:30 PM	1:30 PM	7:30 PM	Lunch Break	Lunch Break	Lunch Break
8:00 AM	1:00 PM	2:00 PM	8:00 PM			
8:30 AM	1:30 PM	2:30 PM	8:30 PM			
9:00 AM	2:00 PM	3:00 PM	9:00 PM	Regular Papers 2: Architectures & Accelerators	Regular Papers 6: Power & Performance Modeling	Regular Papers 9: Cloud & Edge Computing
9:30 AM	2:30 PM	3:30 PM	9:30 PM			
10:00 AM	3:00 PM	4:00 PM	10:00 PM			
10:30 AM	3:30 PM	4:30 PM	10:30 PM	Keynote I Manish Parashar	Keynote II Alba Melo	Keynote III Keshav Pingali
11:00 AM	4:00 PM	5:00 PM	11:00 PM			
11:30 AM	4:30 PM	5:30 PM	11:30 PM	Regular Papers 3: Machine Learning & Applications	Regular Papers 7: Theory & Algorithms	Regular Papers 10: Programming & Languages
12 Noon	5:00 PM	6:00 PM	12 Midnight			
12:30 PM	5:30 PM	6:30 PM	12:30 AM			
1:00 PM	6:00 PM	7:00 PM	1:00 AM			Closing Session

GO TO THE CONFERENCE

The conference is being help using Easychair Virtual Conference Solution (VCS). After [registering](#), access the conference by login in to [Easychair](#).

CONFERENCE PROCEEDINGS

Free access to the EuroPar 2021 conference proceedings until September 25. (available only by following the above link and only from the conference website)

KEYNOTE SPEAKERS

Keynote I: Manish Parashar

» University of Utah, USA

Chair: Leonel Sousa, University of Lisbon, Portugal

Big Data and Extreme-Scales: Computational Science in the 21st Century

September 1 - 10:30 AM New York | 3:30 PM Lisbon | 4:30 PM Brussels | 10:30 PM Beijing

Abstract: Extreme scales and big data are essential to computational and data-enabled science and engineering in the 21st century, promising dramatic new insights into natural and engineered systems. However, data-related challenges are limiting the potential impact of application workflows enabled by current and emerging extreme scale, high-performance, distributed computing environments. These data-intensive application workflows involve dynamic coordination, interactions and data coupling between multiple application processes that run at scale on different resources, and with services for monitoring, analysis and visualization and archiving, and present challenges due to increasing data volumes and complex data-coupling patterns, system energy constraints, increasing failure rates, etc. In this talk I will explore some of these challenges and investigate how solutions based on data sharing abstractions, managed data pipelines, data-staging service, and in-situ / in-transit data placement and processing can be used to help address them. This research is part of the DataSpaces project at the Scientific Computing and Imaging (SCI) Institute, University of Utah.

Keynote II: Alba Cristina Melo

» University of Brasilia, Brazil

Chair: Domenico Talia, University of Calabria, Italy

HPC for Bioinformatics: The Genetic Sequence Comparison Quest for Performance

September 2 - 10:30 AM New York | 3:30 PM Lisbon | 4:30 PM Brussels | 10:30 PM Beijing

Abstract: Genetic Sequence Comparison is an important operation in Bioinformatics, executed routinely worldwide. Two relevant algorithms that compare genetic sequences are the Smith-Waterman (SW) algorithm and Sankoff's algorithm. The Smith-Waterman algorithm is widely used for pairwise comparisons and it obtains the optimal result in quadratic time - $O(n^2)$, where n is the length of the sequences. The Sankoff algorithm is used to structurally align two sequences and it computes the optimal result in $O(n^4)$ time. In order to accelerate these algorithms, many parallel strategies were proposed in the literature. However, the alignment of whole chromosomes with hundreds of millions of characters with the SW algorithm is still a very challenging task, which requires extraordinary computing power. Likewise, obtaining the structural alignment of two sequences with the Sankoff algorithm requires parallel approaches.

In this talk, we first present our MASA-CUDAlign tool, which was used to pairwise align real DNA sequences with up to 249 millions of characters in a cluster with 512 GPUs, achieving the best performance in the literature in 2021. We will present and discuss the innovative features of the most recent version of MASA-CUDAlign: parallelogram execution, incremental speculation, block pruning and score-share balancing strategies. We will also show performance and energy results in homogeneous and heterogeneous GPU clusters. Then, we will discuss the design of our CUDA-Sankoff tool and its innovative strategy to exploit multi-level wavefront parallelism. At the end, we will show a covid-19 case study, where we use the tools discussed in this talk to compare the SARS-CoV-2 genetic sequences, considering the reference sequence and its variants.

Keynote III: Keshav Pingali

» Katana Graph & The University of Texas at Austin, USA

Chair: Fernando Silva, University of Porto, Portugal

Knowledge Graphs, Graph AI, and the Need for High-performance Graph Computing

September 3 - 10:30 AM New York | 3:30 PM Lisbon | 4:30 PM Brussels | 10:30 PM Beijing

Abstract: Knowledge Graphs now power many applications across diverse industries such as FinTech, Pharma and Manufacturing. Data volumes are growing at a staggering rate, and graphs with hundreds of billions edges are not uncommon. Computations on such data sets include querying, analytics, and pattern mining, and there is growing interest in using machine learning to perform inference on large graphs. In many applications, it is necessary to combine these operations seamlessly to extract actionable intelligence as quickly as possible. Katana Graph is a start-up based in Austin and the Bay Area that is building a scale-out platform for seamless, high-performance computing on such graph data sets. I will describe the key features of the Katana Graph Engine that enable high performance, some important use cases for this technology from Katana's customers, and the main lessons I have learned from doing a startup after a career in academia.

SESSIONS

Regular Papers 1: Parallel Methods & Applications

September 1 - 6:00 AM New York | 11:00 AM Lisbon | 12 Noon Brussels | 6:00 PM Beijing

Chair: Radu Prodan, University of Klagenfurt, Austria

A GPU Architecture Aware Fine-Grain Pruning Technique for Deep Neural Networks

» Kyusik Choi and Hoesook Yang

Mixed Precision Incomplete and Factorized Sparse Approximate Inverse Preconditioning on GPUs

» Fritz Goebel, Thomas Grützmacher, Tobias Ribizel and Hartwig Anzt

Outsmarting the Atmospheric Turbulence for Ground-Based Telescopes Using the Stochastic Levenberg-Marquardt Method

» Yuxi Hong, El Houcine Bergou, Nicolas Doucet, Hao Zhang, Jesse Cranney, Hatem Ltaief, Damien Gratadour, Francois Rigaut and David Keyes

GPU Accelerated Mahalanobis-average Hierarchical Clustering Analysis

» Adam Smeiko, Miroslav Kratochvil, Martin Kruliš and Tomáš Sieger

Regular Papers 2: Architectures & Accelerators

September 1 - 9:00 AM New York | 2:00 PM Lisbon | 3:00 PM Brussels | 9:00 PM Beijing

Chair: Roman Wyrzykowski, Czestochowa University of Technology, Poland

PriorAT: Criticality-Driven Prioritization Inside the On-Chip Memory Hierarchy

» Vladimir Dimić, Miquel Moretó, Marc Casas and Mateo Valero

Optimized Implementation of the HPCGBenchmark on Reconfigurable Hardware

» Alberto Zeni, Kenneth O'Brien, Michaela Blott and Marco Domenico Santambrogio

Exploiting co-execution with oneAPI: heterogeneity from a modern perspective

» Raúl Nozal and Jose Luis Bosque

Regular Papers 3: Machine Learning & Applications

September 1 - 11:30 AM New York | 4:30 PM Lisbon | 5:30 PM Brussels | 11:30 PM Beijing

Chair: João Sobral, University of Minho, Portugal

Designing a 3D Parallel Memory-Aware Lattice Boltzmann Algorithm on Manycore Systems

» Yuankun Fu and Fengguang Song

Fault-tolerant LU factorisation is low cost

» Camille Coti, Laure Petrucci and Daniel Alberto Torres Gonzalez

Efficient and Systematic Partitioning of Large and Deep Neural Networks for Parallelization

» Haoran Wang, Chang Li, Thibaut Tachon, Hongxing Wang, Sheng Yang, Sébastien Limet and Sophie Robert

Towards Flexible and Compiler-friendly Layer Fusion for CNNs on Multi-core CPUs

» Zhongyi Lin, Evangelos Georganas and John D. Owens

Regular Papers 4: Scheduling & Load Balancing I

September 2 - 4:30 AM New York | 9:30 AM Lisbon | 10:30 AM Brussels | 4:30 PM Beijing

Chair: Oliver Sinnen, University of Auckland, New Zealand

Collaborative GPU Preemption via Spatial Multitasking for Efficient GPU Sharing

» Zhuoran Ji and Cho-Li Wang

A Fixed-Parameter Algorithm for Scheduling Unit dependent Tasks with Unit Communication Delays

» Ning Tang and Alix Munier-Kordon

Plan-based Job Scheduling for Supercomputers with Shared Burst Buffers

» Jan Kopański and Krzysztof Rządca

A log-linear (2 5/6)-approximation algorithm for parallel machine scheduling with a single orthogonal resource

» Adrian Naruszko, Bartłomiej Przybylski and Krzysztof Rządca

Regular Papers 5: Scheduling & Load Balancing II

September 2 - 6:00 AM New York | 11:00 AM Lisbon | 12 Noon Brussels | 6:00 PM Beijing

Chair: Jorge Barbosa, University of Porto, Portugal

An MPI-Parallel Algorithm for Mapping Complex Networks onto Hierarchical Architectures

» Maria Predari, Charilaos Tzovas, Christian Schulz and Henning Meyerhenke

Taming Tail Latency in Key-Value Stores: a Scheduling Perspective

» Sonia Ben Mokhtar, Louis-Claude Canon, Anthony Dugois, Loris Marchal and Etienne Rivière

Pipelined Model Parallelism: Complexity Results and Memory Considerations

» Olivier Beaumont, Lionel Eyraud-Dubois and Alena Shilova

Enhancing Load-Balancing of MPI Applications with Workshare

» Thomas Dionisi, Stéphane Bouhrour, Julien Jaeger, Patrick Carribault and Marc Perache

Regular Papers 6: Power & Performance Modelling

September 2 - 9:00 AM New York | 2:00 PM Lisbon | 3:00 PM Brussels | 9:00 PM Beijing

Chair: Didem Unat, Koç University, Turkey

Trace-driven Workload Generation and Execution

» Yannis Sfakianakis, Eleni Kanelou, Manolis Morazakis and Angelos Bilas

Update on the Asymptotic Optimality of LPT

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

E2EWatch: An End-to-end Anomaly Diagnosis Framework for Production HPC Systems

» Burak Aksor, Benjamin Schwallier, Omar Aaziz, Vitus Leung, Jim Brandt, Manuel Egele and Ayse Coskun

Sustaining Performance While Reducing Energy Consumption: A Control Theory Approach

» Sophie Cerf, Raphaël Bleuse, Valentin Reis, Swann Perarnau and Eric Rutten

Regular Papers 7: Theory & Algorithms

September 2 - 11:30 AM New York | 4:30 PM Lisbon | 5:30 PM Brussels | 11:30 PM Beijing

Chair: Andrea Pietracaprina, Università di Padova, Italy

Algorithm design for Tensor Units

» Rezaul Chowdhury, Francesco Silvestri and Flavio Vella

TSLQueue: An Efficient Lock-free Design for Priority Queues

» Adones Rukundo and Philippos Tsigas

A Scalable Approximation Algorithm for Weighted Longest Common Subsequence

» Jeremy Buhler, Thomas Lavastida, Kefu Lu and Benjamin Moseley

G-Morph: Induced Subgraph Isomorphism Search of Labeled Graphs on a GPU

» Bryan Rowe and Rajiv Gupta

Regular Papers 8: Tools & Environments

September 3 - 6:00 AM New York | 11:00 AM Lisbon | 12 Noon Brussels | 6:00 PM Beijing

Chair: Gabriel Falcao, University of Coimbra, Portugal

ALONA: Automatic Loop Nest Approximation with Reconstruction and Space Pruning

» Daniel Maier, Biagio Cosenza and Ben Juurlink

Automatic low-overhead load-imbalance detection in MPI applications

» Peter Arzt, Yannic Fischler, Jan-Patrick Lehr and Christian Bischof

Smart Distributed DataSets for Streaming

» Tiago Lopes, Luís Veiga and Miguel Coimbra

Regular Papers 9: Cloud & Edge Computing

September 3 - 9:00 AM New York | 2:00 PM Lisbon | 3:00 PM Brussels | 9:00 PM Beijing

Chair: João Lourenço, Universidade Nova de Lisboa, Portugal

Colony: Parallel Functions as a Service on the Cloud-Edge Continuum

» Francesc-Josep Lordan Gomis, Rosa M. Badia and Daniele Lezzi

Horizontal Scaling in Cloud using Contextual Bandits

» David Delande, Patricia Stolf, Raphaël Feraud, Jean-Marc Pierson and André Bottaro

Geo-Distribute Cloud Application at the Edge

» Ronan-Alexandre Cherrueau, Marie Delavergne and Adrien Lèbre

A Fault Tolerant and Deadline Constrained Sequence Alignment Application on Cloud-based Spot GPU Instances

» Rafaela Brum, Wallisson Sousa, Alba Melo, Cristiana Bentes, Maria Clícia Castro and Lúcia Drummond

Regular Papers 10: Programming & Languages

September 3 - 11:30 AM New York | 4:30 PM Lisbon | 5:30 PM Brussels | 11:30 PM Beijing

Chair: Alfredo Goldman, University of São Paulo, Brazil

Particle-In-Cell simulation using Asynchronous Mapping

» Nicolas Gidotti, Pedro Ceyrat, João Barreto, José Monteiro, Rodrigo Rodrigues, Ricardo Fonseca, Xavier Martorell and Antonio J. Peña

Efficient GPU Computation using Task Graph Parallelism

» Dian-Lun Lin and Tsung-Wei Huang

Accelerating Graph Applications Using Phased Transactional Memory

» Catalina Munoz Morales, Rafael Murai, João P. L. de Carvalho, Bruno Chinelato Honório, Alexandro Baldassin and Guido Araujo

Towards High Performance Resilience using Performance Portable Abstractions

» Nicolas Morales, Keita Teranishi, Bogdan Nicolae, Christian Trott and Franck Cappello

