

Important Dates

- Abstract Submission: 7 February 2020 extended to 10 February 2020, 23:59 Anywhere on Earth (AoE), final extension
- Paper Submission Deadline: 14 February 2020 extended to 21 February 2020, 23:59 Anywhere on Earth (AoE), final extension
- Author Notification: 30 April 2020
- Camera-Ready Papers: 5 June 2020

Conference

Euro-Par is the prime European conference covering all flavors of parallel and distributed processing.

- from theory to practice,
- from multi-core processors to accelerators, supercomputers and clouds,
- from fundamental algorithmic problems to systems and tools,
- from applications in computational sciences to machine learning and artificial intelligence.

Euro-Par proceedings are published by Springer in the LNCS series.

The main audience of Euro-Par are the researchers in academic institutions, government laboratories and industrial organisations. Euro-Par's objective is to be the primary choice of such professionals for the presentation of new results in their specific areas. Euro-Par's unique organization into topics provides an excellent forum for focused technical discussion, as well as interaction with a large, broad and diverse audience. In addition, Euro-Par conferences provide a platform for a number of accompanying Technical Workshops for smaller and emerging communities.

This call for conference papers will be followed by a separate call for workshop papers (in February 2020) and for posters for the doctoral symposium (in March 2020).

Submission Guidelines

Submit your papers through EasyChair: <https://easychair.org/conferences/?conf=europar2020>

- Papers must be in PDF format and should not exceed 12 14 pages (including references) and 7500 words.
- Papers must be formatted in the Springer LNCS style: <https://www.springer.com/la/complex-science/lns/conference-tocguidelines>
- Papers that don't meet these requirements might be rejected without a review.
- Only contributors that are not submitted elsewhere or currently under review will be considered.
- All submitted papers will be checked for originality by Springer iThenticate. Papers which show an insufficient originality might be rejected without a review.

Topics

We invite submissions of high-quality, novel and original research results in areas of parallel and distributed computing covered by the List of Topics.

Topic 1: Support Tools and Environments

Global Chair: Michael Gernat, Technical University of Munich, Germany

Local Chair: Marusz Sterzel, Academic Computer Centre Cyfronet AGH, Krakow, Poland

Topic 2: Performance and Power Modeling, Prediction and Evaluation

Global Chair: Arnaud Legrand, Univ. Grenoble Alpes, Inria, CNRS, France

Local Chair: Ariel Oleksak, Poznan Supercomputing and Networking Center, Poland

Topic 3: Scheduling and Load Balancing

Global Chair: Sascha Hunold, Technical University Wien, Austria

Local Chair: Joanna Berlinska, Adam Mickiewicz University, Poznan, Poland

Topic 4: High Performance Architectures and Compilers

Global Chair: Leonel Sousa, INESC-ID, IST, University of Lisbon, Portugal

Local Chair: Pawel Czumal, Gdansk University of Technology, Poland

Topic 5: Data Management, Analytics and Machine Learning

Global Chair: Morris Riedel, Forschungszentrum Juelich, Germany

Local Chair: Jacek Sroka, University of Warsaw, Poland

Topic 6: Cluster, Cloud and Edge Computing

Global Chair: Maria S. Pérez, Universidad Politécnica de Madrid, Spain

Local Chair: Lukasz Dutka, Academic Computer Centre Cyfronet AGH, Krakow, Poland

Topic 7: Theory and Algorithms for Parallel and Distributed Processing

Global Chair: Ben Moseley, Carnegie-Mellon University, UA

Local Chair: Marek Klonowski, University of Wroclaw, Poland

Topic 8: Parallel and Distributed Programming, Interfaces, and Languages

Global Chair: Phil Trinder, University of Glasgow, United Kingdom

Local Chair: Wojciech Turek, AGH University of Science and Technology, Krakow, Poland

Topic 9: Multicore and Manycore Parallelism

Global Chair: Arturo Gonzalez, Universidad de Valladolid, Spain

Local Chair: Witold Rudnicki, University of Białystok, Poland

Topic 10: Parallel Numerical Methods and Applications

Global Chair: Hatem Ltaief, KAUST, Saudi Arabia

Local Chair: Erin Carson, Charles University, Prague, Czechia

Topic 11: Accelerator Computing

Global Chair: Alba Cristina Melo, University of Brasilia, Brasil

Local Chair: Lukasz Szustak, Czestochowa University of Technology, Poland

Per-Topic CFP

TOPIC 1: SUPPORT TOOLS AND ENVIRONMENTS

Global Chair: Michael Gernat, Technical University of Munich, Germany

Local Chair: Marusz Sterzel, Academic Computer Centre Cyfronet AGH, Krakow, Poland

Despite an impressive body of research, parallel and distributed programming remains a complex task, prone to subtle software issues that can affect both the correctness and the performance of the application. This topic focuses on tools and techniques to help tackling that complexity. We solicit contributions on tools and environments that address any of the many challenges of parallel and distributed programming related to programmability, portability, correctness, reliability, scalability, efficiency, performance and energy consumption.

This topic aims to bring together tool designers, developers, and users to share their concerns, ideas, solutions, and products for a wide range of parallel platforms. We particularly value contributions with solid theoretical foundations and with strong experimental validations on production-level parallel and distributed systems. We encourage submissions that detail novel program development tools and environments that address the expected complexity of exascale systems.

Focus

- Debugging and correctness tools
- Hybrid shared memory and message passing tools
- Instrumentation and monitoring tools and techniques
- Program development tools
- Programming environments, interoperable tool environments
- Integration of tools, compilers and operating systems
- Performance and reliability analysis (manual and automatic)
- Energy efficiency and savings tools
- Performance and code structure visualization
- Testing and analysis tools
- Computational steering
- Tool infrastructure and scalability
- Tool evaluations and comparisons in production environments
- Tools for extreme-scale systems
- Tools for code modernization
- Tools for heterogeneous and heterogeneous multi/many-core processors
- Tools and environments for clusters, clouds, and grids
- Autotuning techniques and tools
- Tool success stories

TOPIC 2: PERFORMANCE AND POWER MODELING, PREDICTION AND EVALUATION

Global Chair: Arnaud Legrand, Univ. Grenoble Alpes, Inria, CNRS, France

Local Chair: Ariel Oleksak, Poznan Supercomputing and Networking Center, Poland

In recent years, a range of novel methods and tools have been developed for the evaluation, design, and modeling of parallel and distributed systems and applications. At the same time, the term 'performance' has broadened to also include scalability and energy efficiency, and touching reliability and robustness in addition to the classic resource-oriented metrics. The aim of this topic is to gather researchers working on different aspects of performance modeling, evaluation, and prediction, be it for systems or for applications running on the whole range of parallel and distributed systems (multi-core and heterogeneous architectures, HPC systems, grid and cloud contexts etc.). Authors are invited to submit novel research in all areas of performance modeling, prediction and evaluation, and to help bring together current theory and practice.

Focus

- Design of experiments, reproducible experiments
- Novel techniques and tools for performance measurement, evaluation, and prediction
- Advanced simulation techniques and tools
- Measurements, benchmarking, and tracing
- Workload modeling
- Performance-driven code optimization
- Verification and validation of performance models
- Performance visualization
- Power consumption modeling and prediction
- Performance modeling and simulation of emerging exascale systems

TOPIC 3: SCHEDULING AND LOAD BALANCING

Global Chair: Sascha Hunold, Technical University Wien, Austria

Local Chair: Joanna Berlinska, Adam Mickiewicz University, Poznan, Poland

New computing systems offer the opportunity to reduce the response times and the energy consumption of the applications by exploiting the levels of parallelism. Heterogeneity and complexity are the main characteristics of modern architectures. Thereby, the optimal exploitation of modern platforms is a challenging. Scheduling and load balancing techniques are key instruments to achieve higher performance, lower energy consumption, reduced resource usage, and real-time properties of applications.

This topic invites papers on all aspects related to scheduling and load balancing on parallel and distributed machines, from theoretical foundations for modeling and designing efficient and robust scheduling policies to experimental studies, applications and practical tools and solutions. It applies to multi-/many-core processors, embedded systems, servers, heterogeneous and accelerated systems, HPC clusters as well as distributed systems such as clouds and global computing platforms.

Focus

All aspects related to scheduling and load balancing on parallel and distributed machines including but not limited to:

- Scheduling algorithms for homogeneous and heterogeneous platforms
- Theoretical foundations of scheduling algorithms
- Real-time scheduling on parallel and distributed machines
- Robustness of scheduling algorithms
- Feedback-based load balancing
- Multi-objective scheduling
- Resilient scheduling
- Scheduling, coordination and overhead at extreme scales
- Online scheduling
- Energy and temperature awareness in scheduling and load balancing
- Workload characterization and modeling
- Workflow scheduling
- Performance models for scheduling and load balancing
- Reproducibility of scheduling

TOPIC 4: HIGH PERFORMANCE ARCHITECTURES AND COMPIERS

Global Chair: Leonel Sousa, INESC-ID, IST, University of Lisbon, Portugal

Local Chair: Pawel Czumal, Gdansk University of Technology, Poland

This topic deals with architecture design, languages, and compilation for parallel high-performance systems. The areas of interest range from microprocessors to large-scale parallel machines (including multi-/many-core, possibly heterogeneous, architectures); from general-purpose to specialized hardware platforms (e.g., graphic coprocessors, low-power embedded systems); and from architecture design to compiler technology and language design.

On the compilation side, topics of interest include programmer productivity issues, concurrent and/or sequential language aspects, vectorization, program analysis, program transformation, automatic discovery and/or management of parallelism at all levels, autotuning and feedback-directed compilation, and the interaction between the compiler and the system at large. On the architecture side, the scope spans system architectures, processor/micro-architecture, memory hierarchy, and multi-threading, architectural support for parallelism, and the impact of emerging hardware technologies.

Focus

- Compiling for multi-threaded/multi-core/many-core/vector and heterogeneous processors/architectures
- Compiling for emerging architectures (low-power embedded systems, reconfigurable hardware, processors in memory, coprocessors)
- Iterative, just-in-time, feedback-oriented, dynamic, and machine-learning-based compilation
- Static analysis and interaction between static and dynamic analysis
- Programmer productivity tools and analysis for high-performance architectures
- Program transformation systems
- High-level programming models and tools for multi-/many-core and heterogeneous architectures
- Interaction between compiler, runtime system, application, hardware, and operating system
- Parallel computer architecture design – ILP, DLP, multi-threaded, and multi-core processors
- Designs and compiler optimizations for power/performance efficiency
- Software and hardware fault-tolerance techniques
- Memory hierarchy, emerging memory technologies, and 3D stacked memories
- Application-specific, reconfigurable and embedded parallel systems
- Compiler, run-time, and architectural support for dynamic adaptation
- Optimizing compilers for domain-specific languages

TOPIC 5: DATA MANAGEMENT, ANALYTICS AND MACHINE LEARNING

Global Chair: Morris Riedel, Forschungszentrum Juelich, Germany

Local Chair: Jacek Sroka, University of Warsaw, Poland

Many areas of science, industry, and commerce are producing extreme-scale data that must be processed – stored, managed, analyzed – in order to extract useful knowledge. This topic seeks papers in all aspects of distributed and parallel data management and data analysis. For example, cloud and grid data-intensive processing, parallel and distributed machine learning, HPC in situ data analytics, parallel storage systems, scalable data processing workflows, and distributed stream processing are all in the scope of this topic.

Focus:

- Parallel, replicated, and highly-available distributed databases
- Cloud and HPC storage architectures and systems
- Scientific data analytics (Big Data or HPC based approaches)
- Middleware for processing large-scale data
- Programming models for parallel and distributed data analytics
- Workflow management for data analytics
- Coupling HPC simulations with in situ data analysis
- Parallel data visualization
- Distributed and parallel transaction query processing and information retrieval
- Internet-scale data-intensive applications
- Sensor network data management
- Data-intensive clouds and grids
- Parallel data streaming and data stream mining
- New storage hierarchies in distributed data systems
- Parallel and distributed machine learning, knowledge discovery and data mining
- Privacy and trust in parallel and distributed data management and analytics systems
- IoT data management and analytics

TOPIC 6: CLUSTER, CLOUD AND EDGE COMPUTING

Global Chair: Maria S. Pérez, Universidad Politécnica de Madrid, Spain

Local Chair: Lukasz Dutka, Academic Computer Centre Cyfronet AGH, Krakow, Poland

While the term Cluster Computing is hardware oriented and determines the organization of large computer systems at one location, the term Cloud Computing addresses the use of such large computer systems. Since Cluster and Cloud Computing complement each other, there are interdependencies between many research questions addressing these topics. In this Topic of EuroPar, we will particularly focus on these interdependencies in addition to the topics specifically addressing issues belonging only to one of both areas.

In Cluster computing, important research topics focus on performance, reliability, and energy efficiency as well as the impact of novel processor architectures. Since Cloud Computing tries to hide hardware and system software details from the end-user, research issues include various forms of virtualization and their impact on performance, resource management, and business models that address system owner and user interests.

Further, it is interesting to address Cloud Computing on top of several smaller clusters and its advantages with respect to reliability and load balancing on a high abstraction level as well as the consideration of networks.

Finally, the combination of local computer installation together with Cloud Computing, also referred to as 'fog/edge' computing has received growing interest in recent time. This concept has led to many research questions, like an appropriate distribution of subtasks to the available systems under the consideration of various constraints.

Since many research studies in this area use experimental evaluation, we expect the authors reporting such studies to provide sufficient study details, if necessary complemented by a possibly web-based supplement, to allow a technical evaluation during the review process and reproducibility and replicability of results if the submission is accepted.

Focus

- Cloud-enabled applications and platforms
- Interoperability and portability in Cloud Computing
- Aggregation and Federation of Clouds
- Hybrid Fog and edge computing
- Energy efficiency in Cluster and Cloud Computing
- Resource/Service/Information discovery in Clouds
- Resource management and scheduling in Clusters and Clouds
- Cloud programming models, tools, and algorithms
- Dependability, adaptability, and scalability of Cloud applications
- Security and privacy for Clouds
- Accounting, billing and business models for Cloud Computing
- Management of resources and applications in Clusters and Clouds
- Quality-of-Service and Service-Level-Agreement in Clouds
- Containers and serverless computing

TOPIC 7: THEORY AND ALGORITHMS FOR PARALLEL AND DISTRIBUTED PROCESSING

Global Chair: Ben Moseley, Carnegie-Mellon University, USA

Local Chair: Marek Klonowski, University of Wroclaw, Poland

Nowadays distributed and parallel data processing is ubiquitous. Parallel cores are available on smartphones, laptops, servers and supercomputing nodes. Many devices cooperate in a fully distributed and heterogeneous systems to provide even basic services. Despite astonishing progress in recent years, many challenges remain. We urgently need better or specific solutions for scalability, load balancing or more efficient communication in more and more complex systems. We also need more robust algorithms to cope with failures, malicious or selfish behaviour.

For a better design of algorithms, we need theoretical tools to understand parallel and distributed algorithms.

High quality, original papers are solicited on this general topic of the theory of parallel and distributed algorithms.

Focus

The focus is on, but not limited to, the theoretical aspects of the following:

- Foundations, complexity theory, models, and emerging paradigms for parallel and distributed computation
- Design and practice of distributed and parallel algorithms
- Algorithmic aspects of packing, scheduling, and resource management in distributed and parallel systems
- Scalability, concurrency and performance
- Fault tolerance, error resilient and self-stabilizing algorithms
- Distributed storage and distributed data processing
- Dependable, secure and privacy-preserving distributed systems
- Power/energy-efficient algorithms
- Distributed operating systems
- Algorithms on GPUs and accelerators
- Data structures for parallel and distributed parallel computing
- Algorithms and models for big data/data-intensive parallel computing
- Algorithms for routing and information dissemination, communication networks
- Algorithms for cloud computing
- Algorithmic game theory related to parallel and distributed systems
- Algorithms for computational and collaborative learning
- Algorithms for social networks
- Instruction level parallelism research
- Lower bounds

TOPIC 8: PARALLEL AND DISTRIBUTED PROGRAMMING, INTERFACES, AND LANGUAGES

Global Chair: Phil Trinder, University of Glasgow, United Kingdom

Local Chair: Wojciech Turek, AGH University of Science and Technology, Krakow, Poland

Parallel and distributed applications require appropriate programming abstractions and models, efficient design tools, parallelization techniques and practices. This topic is open for the presentation of new results and practical experience in this domain. Efficient and effective parallel languages, interfaces, libraries and frameworks, as well as solid practical and experimental validation. It emphasizes research on high-performance, correct, portable, and scalable parallel programs via appropriate parallel and distributed programming model, interface and language support. Contributions that assess programming abstractions, models and methods for usability, performance prediction, scalability, self-adaptation, rapid prototyping and fault-tolerance, as needed. For instance, in dynamic heterogeneous (i.e. parallel and distributed) infrastructures, are welcome. Authors are invited to include quantitative evaluations of their claims.

Focus

- Programming paradigms and techniques for novel infrastructures like accelerators, exascale systems, low-power architectures and clouds
- Design and implementation, performance analysis and performance portability of programming models across parallel and distributed platforms
- Innovative paradigms, programming models, languages and libraries for parallel and distributed applications
- Programming models, languages and applications in clusters and cloud environments
- Programming tools for application design, implementation and performance tuning
- Application case studies for benchmarking and comparative studies of parallel programming models
- Domain-specific libraries and languages
- Parallel and distributed programming productivity, usability, and component-based parallel programming

TOPIC 9: MULTICORE AND MANYCORE PARALLELISM

Global Chair: Arturo Gonzalez, Universidad de Valladolid, Spain

Local Chair: Witold Rudnicki, University of Białystok, Poland

Modern homogeneous and heterogeneous multicore and manycore architectures are now part of the high-end, embedded, and mainstream computing scene and can offer impressive performance for many applications. This architecture trend has been driven by the need to reduce power consumption, increase processor utilization, and deal with the memory-processor speed gap. However, the complexity of these new architectures has created several programming challenges, and achieving performance on these systems is often a difficult task. This topic seeks to explore productive programming of multi- and manycore systems as well as stand-alone systems with large numbers of cores like GPUs and various types of accelerators. This can also include hybrid and heterogeneous systems with different types of multicore processors. The focus is on novel research and solutions in the form of programming models, algorithms, languages, compilers, libraries, runtime and analysis tools to increase the programmability of multicore, many-core, and heterogeneous systems, in the context of general-purpose, high-performance, and embedded parallel computing.

Focus

- Programming techniques, models, frameworks and languages
- Compiler optimizations and techniques
- Lock-free algorithms, transactional memories
- Libraries and runtime systems
- Tools for discovering and understanding parallelism
- Advances in algorithms and data-structures
- Hardware support for programming models and runtime systems
- Models, methods and tools for innovative many-core architectures
- Performance and power trade-offs and scalability
- Innovative applications and case studies

TOPIC 10: PARALLEL NUMERICAL METHODS AND APPLICATIONS

Global Chair: Hatem Ltaief, KAUST, Saudi Arabia

Local Chair: Erin Carson, Charles University, Prague, Czechia

The need for high-performance computation is driven by the need for large-scale simulation and data analysis in science and engineering, finance, life sciences, etc. This requires the design of highly scalable numerical methods and algorithms that are able to efficiently exploit modern computer architectures. The scalability of these algorithms and methods and their ability to efficiently utilize high-performance heterogeneous resources is critical to improving the performance of computational and data science applications.

This conference topic aims to provide a forum for presenting and discussing recent developments in parallel numerical algorithms and their implementation on current parallel architectures, including many-core and hybrid architectures. We encourage submissions that address algorithmic design, implementation details, performance analysis, as well as integration of parallel numerical methods in large-scale applications.

Focus

The focus is on, but not limited to, the following topics:

- Numerical linear algebra for dense and sparse matrices
- Synchronization-reducing and communication-avoiding algorithms
- Optimization and non-linear problems
- Mixed precision algorithms exploiting low-precision hardware
- High-dimensional problems and reduction methods
- Numerical methods for large-scale data analysis
- Uncertainty quantification
- Applications of numerical algorithms in science, engineering, and data analysis

TOPIC 11: ACCELERATOR COMPUTING

Global Chair: Alba Cristina Melo, University of Brasilia, Brasil

Local Chair: Lukasz Szustak, Czestochowa University of Technology, Poland

Hardware accelerators of various kinds offer a potential for achieving massive performance in applications that can leverage their high degree of parallelism and customization. Examples include graphics processors (GPUs), manycore co-processors, as well as more customizable devices, such as FPGA-based systems, and streaming data flow architectures.

The research challenge for this topic is to explore new directions for actually realizing this potential. We encourage submissions in all areas related to accelerators: architectures, algorithms, languages, compilers, libraries, runtimes systems, coordination of accelerators and CPU, and debugging and profiling tools. Application-related submissions that contribute new insights into fundamental problems or novel solution approaches in this domain are welcome as well, including big data, data analytics, machine learning and computational science/engineering.

Focus

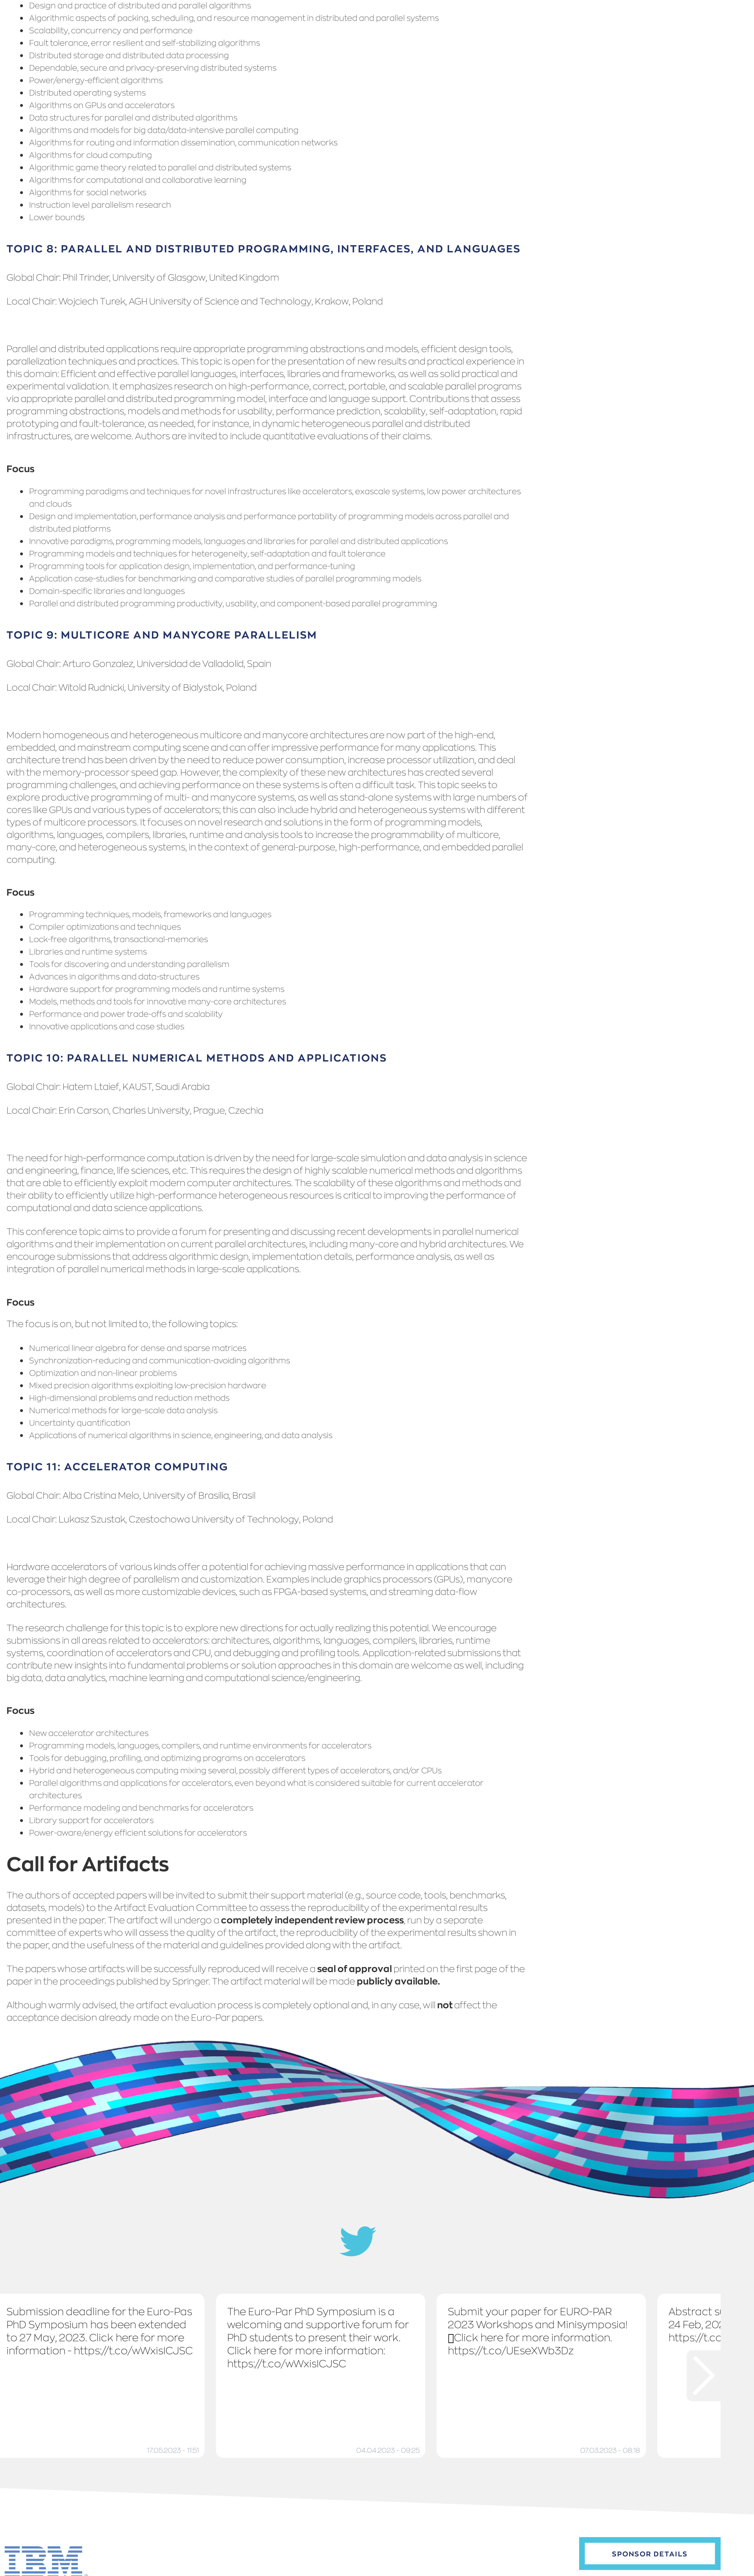
- New accelerator architectures
- Programming models, languages, compilers, and runtime environments for accelerators
- Tools for debugging, profiling, and optimizing programs on accelerators
- Hybrid and heterogeneous computing mixing several, possibly different types of accelerators, and/or CPUs
- Parallel algorithms and applications for accelerators, even beyond what is considered suitable for current accelerator architectures
- Performance modeling and benchmarks for accelerators
- Library support for accelerators
- Power-aware/energy efficient solutions for accelerators

Call for Artifacts

The authors of accepted papers will be invited to submit their support material (e.g., source code, tools, benchmarks, datasets, models) to the Artifact Evaluation Committee to assess the reproducibility of the experimental results presented in the paper. The artifact will undergo a **completely independent review process**. Only by a separate committee of experts who will assess the quality of the artifact, the reproducibility of the experimental results shown in the paper, and the usefulness of the material and guidelines provided along with the artifact.

The papers whose artifacts will be successfully reproduced will receive a **seal of approval** printed on the first page of the paper in the proceedings published by Springer. The artifact material will be made **publicly available**.

Although warmly advised, the artifact evaluation process is completely optional and, in any case, will **not** affect the acceptance decision already made on the Euro-Par papers.



Submission deadline for the Euro-Pas PhD Symposium has been extended to 27 May, 2023. Click here for more information: <https://t.co/wWxisJCISC>

The Euro-Par PhD Symposium is a welcoming and supportive forum for PhD students to present their work. Click here for more information: <https://t.co/wWxisJCISC>

Submit your paper for EURO-PAR 2023 Workshops and Minisymposia! Click here for more information: <https://t.co/UEseXWb5Dz>

Abstracts 24 Feb, 2023: <https://t.co/UEseXWb5Dz>



SHARE ON:

