31st Euromicro International Conference on Parallel, Distributed and Network-Based Processing

Naples, Italy
March 1-3, 2023

Book of Abstracts

Day 1. Wednesday 1 March 2023

	TEATRO	Sala Procida	AWS Academy
9.15 AM – 10.30 AM	Opening Ceremony		
10.30 AM – 11.15 AM	Plenary session: Brendan Bouffler pag. 2		
11.15 AM – 11.30 AM	Sala Wagner - Coffee Break		
11.30 AM – 1.00 PM	Main Track pag. 7	BDCSA2023 pag. 10	SALTCSMLNHPC2023 pag. 13
1.00 PM – 2.30 PM	Sala Wagner - Lunch		
2.30 PM – 3.15 PM	Plenary session: Angelo Ciaramella pag. 2		
3.15 PM – 3.30 PM	Sala Wagner - Coffee Break		
3.30 PM – 5.30 PM	Main Track pag. 16	BDCSA2023 pag. 19	SALTCSMLNHPC2023 pag. 22

Day 2. Thursday 2 March 2023

	TEATRO	Sala Procida	AWS Academy
10.00 AM – 10.45 AM	Plenary session: Jorge Ejarque Artigas pag. 3		
10.45 AM – 11.00 AM	Sala Wagner - Coffee Break		
11.00 AM – 12.30 PM	Main Track pag. 26	HPCMS2023 pag. 29	FlexMPI (tutorial) pag. 32
12.30 PM – 2.00 PM	Sala Wagner - Lunch		
2.00 PM – 2.45 PM	Plenary session: Alessandro Mei pag. 3		
2.45 PM – 3.00 PM	Sala Wagner - Coffee Break		
3.00 PM – 4.30 PM	Main Track pag 33	HPCMS2023 pag 36	FlexMPI (tutorial) pag. 32
4.30 PM – 6.00 PM	Sala Wagner – Industrial Session (E4 Computer Engineering)		

Day 3. Friday 3 March 2023

	TEATRO	Sala Procida	AWS Academy
10.00 AM – 10.45 AM	Plenary session: Giuseppe Coviello pag. 4		
10.45 AM – 11.00 AM	Sala Wagner - Coffee Break		
11.00 AM – 12.30 PM	Main Track pag. 39	CCIaaSA2023 pag. 42	CC2023 pag. 44
12.30 PM – 2.00 PM	Sala Wagner - Lunch		
2.00 PM – 2.45 PM	Plenary session: Pasqua D'Ambra pag. 4		
2.45 PM – 3.00 PM	Sala Wagner - Coffee Break		
3.00 PM – 4.30 PM	Main Track pag. 47	Parallel Track	
4.30 PM – 5.30 PM	Awards and conclusions		

TEATRO 10:30 AM - 11:15 AM

Distributed computing disrupts. Discuss.

Brendan Bouffler

HPC Engineering

Distributed computing is becoming more integrated into the lives of humans than most would recognize, leading to some great outcomes (and some less-than-great ones, we should also admit). To understand what's net, it's worth looking at the drivers of this success, the problems yet to be solved, and the optimistic view of how we can all help.

TEATRO 2:30 PM - 3:15 PM

Soft Computing in data integration and decision-making

Angelo Ciaramella

University of Naples "Parthenope"

Nowadays, data-driven decision-making methodologies are influencing many aspects of society. In the era of Big Data, data is generated, collected, and analyzed on a large scale.

For example, in biological and biomedical fields, high throughput methods acquire a large number of molecular parameters, collected in "omics" datasets, by a single experiment. Recently, a number of studies pointed out the best performance comes from the integration of multi-omics data. Moreover, record linkage aims to identify records from multiple data sources that refer to the same entity in the real world.

On the other hand, Soft computing is a set of algorithms, including neural networks, fuzzy logic, and evolutionary algorithms that are tolerant of imprecision, uncertainty, partial truth, and approximation, with a great rule in decision support systems.

This talk explores the progress made in decision-making, data integration, and record linkage in addressing novel challenges by using Soft Computing methodologies and identifies a range of open challenges for the community.

TEATRO 10:00 AM - 10:45 AM

Introducing the FaaS model in Complex HPC Workflows: The eFlows4HPC approach

Jorge Ejarque Artigas

Barcelona Supercomputing Center

The evolution of High-Performance Computing (HPC) platforms enables the design and execution of progressively larger and more complex workflow applications in these systems. The complexity comes not only from the number of elements that compose the workflows but also from the type of computations they perform. While traditional HPC workflows target simulations and modelling of physical phenomena, current needs require additional data analytics (DA) and artificial intelligence (AI) tasks. However, these workflows are not just complex to develop, they are also difficult to deploy and execute in shared and restricted HPC sites. Workflow developers implement tools to orchestrate the deployment, configuration and execution of different frameworks which are used in the different workflow steps (HPC simulations or DA and AI algorithms). In this talk, the speaker will present the problems and challenges to facilitate the development, deployment, and execution of these complex workflows in the HPC sites and their analogies with the Function-as-a-Service model used in the Cloud Computing environment. He will introduce the methodology developed by the EuroHPC eFlows4HPC project which focuses on how similar approaches can be followed to simplify the deployment and execution of Complex workflows in HPC systems.

TEATRO 2:00 PM – 2.45 PM

Frauds in the Cryptocurrency Ecosystem

Alessandro Mei

University of Rome “La Sapienza”

The cryptocurrency market is loosely regulated. Even if policymakers are making some progress, building a safer environment for cryptocurrency investors is a complex task that needs time. Meanwhile, blockchain-related technologies evolve fast, and, with the birth of the DeFi, investors begin to move from centralized exchanges (CEX) like Binance to decentralized exchanges (DEX). While regulating the standard cryptocurrencies market is not easy, ruling the on-chain trading platform is harder. In this talk, we will see what are the major security challenges for investors in the cryptocurrency environment, what are the main frauds, and what we can do to build a safer eco-system.

TEATRO 10:00 AM - 10:45 AM

A solution for real-time streaming applications

Giuseppe Coviello

NEC Labs America

Throughout history, humanity has collected data by observing natural phenomena. Collecting data from observations, analyzing it, and deriving higher-level insights are natural inclinations of human beings and essential processes for developing past, present, and future civilizations. The exponential growth in smart sensors and rapid progress in 5G networks create a world awash with data streams. However, high programming complexity is a critical barrier to building performant multi-sensor distributed stream processing applications. We propose DataX, a novel platform that improves programmer productivity by enabling easy exchange, transformations, and fusion of data streams. DataX abstraction simplifies the application's specification and exposes parallelism and dependencies among the application functions (microservices). DataX runtime automatically sets up appropriate data communication mechanisms and enables effortless reuse of microservices and data streams across applications. DataX makes it easy to write, deploy, and reliably operate distributed applications at scale. Synthesizing these capabilities into a single platform is more transformative than any available stream processing system.

TEATRO 2:00 PM – 2.45 PM

Node-level efficiency and scalability issues in iterative sparse linear solvers at scale

Pasqua D'Ambra

Institute for Applied Computing of the National Research Council

In this talk, I will present activities aimed at designing and developing algorithms and mathematical software that enable scientific and engineering applications to face the exascale challenge. Resolution needs of high-fidelity numerical simulations reach numbers of degrees of freedom which go beyond O(10^10) and are in an increasing trend. This generally imposes a rethinking of numerical algorithms and solvers to exploit high degrees of parallelism and complexity of the current (and near-future) high-end supercomputers. I will focus on iterative solvers for linear algebraic equations arising in physics-driven scientific simulation and present parallel efficiency and scalability results from engineering applications in the energy sector. The activities are developed at the HPC Lab of IAC-CNR in the context of some European projects and of the Italian National Center on HPC, Big Data and Quantum Computing.

WEDNESDAY, 1 MARCH

TEATRO 11:30 AM - 1:00 PM

MAIN TRACK.

CHAIR: Karl-Erwin Grosspietsch (Euromicro)

- Karthick Panner Selvam and Mats Brorsson: **Performance Analysis and Benchmarking of a Temperature Downscaling Deep Learning Mode**l
- Shoichi Hirasawa and Michihiro Koibuchi: **An Auto-Tuning Method for High-Bandwidth Low-Latency Approximate Interconnection Networks**
- Raffaele Montella, Diana Di Luccio, Ciro Giuseppe De Vita, Gennaro Mellone, Marco Lapegna, Gloria Ortega, Livia Marcellino, Enrico Zambianchi, and Giulio Giunta: **A highly scalable high-performance Lagrangian transport and diffusion model for marine pollutants assessment**
- Marcelo Koji Moori, Hiago Mayk G. de A. Rocha, Matheus Almeida Silva, Janaína Schwarzrock, Arthur Lorenzon, and Antonio Carlos Schneider Beck Filho: **Automatic CPU-GPU Allocation for Graph Execution**

Performance Analysis and Benchmarking of a Temperature Downscaling Deep Learning Model

Karthick Panner Selvam and Mats Brorsson

We are presenting here a detailed analysis and performance characterization of a Statistical temperature downscaling application used in the MAELSTROM EuroHPC project. This application uses a deep learning methodology to convert low-resolution atmospheric temperature states into high-resolution. We have performed in-depth profiling and roofline analysis at different levels (Operators, Training, Distributed Training, Inference) of the downscaling model on different hardware architectures (Nvidia V100 \& A100 GPUs). Finally, we compare the training and inference cost of the downscaling model with various cloud providers. Our results identify the model bottlenecks which can be used to enhance the model architecture and determine hardware configuration for efficiently utilizing the HPC. Furthermore, we provide a comprehensive methodology for in-depth profiling and benchmarking of the deep learning models.

An Auto-Tuning Method for High-Bandwidth Low-Latency Approximate Interconnection Networks

Shoichi Hirasawa and Michihiro Koibuchi

The next-generation interconnection networks, such as 400GbE specification, impose Forwarding Error Correction (FEC) operation, such as RS-FEC (544,514), to incoming packets at every switch. The significant FEC latency increases the end-to-end communication latency that degrades the application performance in parallel computers. To resolve the FEC latency problem, a prior work presented error-prone high-bandwidth low-latency networks that do not perform the FEC operation. They enable high-bandwidth approximate data transfer and low-bandwidth perfect data transfer to support various kinds of parallel applications subject to different levels of probability of bit-flip occurrence. As the number of approximate data transfers increases, the parallel applications can obtain a significant speedup of their execution at the expense of the moderate degraded quality of results (QoRs). However, it is difficult for users to identify whether each communication should be approximate or not, so as to obtain the shortest execution time with enough QoRs for a given parallel application. In this study, we apply an auto-tuning framework for approximate interconnection networks; t automatically identifies whether each communication should be approximate data transfer or not, by attempting thousands executions of a given parallel application. An auto-tuning attempts a large number of program executions by varying the possible communication parameters to find out the best execution configuration of the program. The multiple executions would generate different positions of bit flips on communication data that may provide different qualities of results even if the same parameters are taken. Although this uncertainty introduces difficulties in the optimization of the auto-tuning, many offline trials lead to a high probability of the program's success execution. Evaluation results show that high-performance MPI applications with our auto-tuning method result in 1.30 average performance improvement on error-prone high-performance approximate networks.

A highly scalable high-performance Lagrangian transport and diffusion model for marine pollutants assessment

Raffaele Montella, Diana Di Luccio, Ciro Giuseppe De Vita, Gennaro Mellone, Marco Lapegna, Gloria Ortega, Livia Marcellino, Enrico Zambianchi, and Giulio Giunta:

While using High-Performance Computing (HPC) for precise and accurate air quality forecasts is a common issue, similar services devoted to marine pollution in coastal areas remain challenging. This paper presents Water quality Community Model Plus Plus (WaComM++) leveraging a parallelization schema enabling the users to run it on heterogeneous parallel architectures. We evaluated the proposed model under several execution approaches using a real-world application for pollutants forecast in the Gulf of Napoli (Campania, Italy). As a result, WaComM++ has produced results 657K times faster than the sequential run (taking into account the Particles' Outer Cycle and not considering the particle domain distribution) when using distributed and shared memory with multi-GPUs dealing with about 25 million particles.

Automatic CPU-GPU Allocation for Graph Execution

Marcelo Koji Moori, Hiago Mayk G. de A. Rocha, Matheus Almeida Silva, Janaína Schwarzrock, Arthur Lorenzon, and Antonio Carlos Schneider Beck Filho:

Although advances in modern GPUs have accelerated the execution of heavy data processing applications, speeding up graph processing on these systems is not a trivial task: graph applications are characterized by their high volume of irregular memory access that varies with the graph structure so that they do not reach their peak performance when executing on GPUs in many times. In these cases, the CPU execution is more suitable. Given that graph structures can be identified through high-level metrics (e.g., diameter and average clustering coefficient), they may assist the designer in deciding whether to execute a given input graph (GPU or CPU). Based on that, in this work, we propose GraCo: a graph processing framework to help the decision-making on where to process a batch of graph applications. Whenever a new batch is submitted to the target HPC system, GraCo decides the best machine to execute each application based only on the available high-level features, precluding any additional applications' execution. Our experimental results comparing GraCo with three other strategies executed on an HPC system comprised of 4 CPUs and 3 GPUs showed that GraCo outperforms the other strategies by at least 34.94x, 13.59x, and 492.31x in total execution time, energy, and energy-delay product.

WEDNESDAY, 1 MARCH

SALA PROCIDA 11:30 AM - 1:00 PM

SPECIAL SESSION. Big Data Convergence: From Sensors to Applications (BDCSA2023)

CHAIR: Jesus Carretero (Universidad Complutense de Madrid)

- Jesus Carretero and Cristhian Martinez: **Blockchain-based schemes for continuous verifiability and traceability of IoT data**
- Javier Garcia Blas, Cosmin Octavian Petre, Genaro Juan Sanchez Gallegos, and Jesus Carretero: **Network accelerated in-memory ad-hoc file system for data-centric and high-performance applications**
- David E. Singh, Alvaro Arbe Milara, and Jesus Carretero: **Energy-aware malleable scheduling techniques**
- Paula Ferrero-Roza, José A. Moríñigo, and Filippo Terragni: **Scaling of the SVD Algorithm for HPC Science: A PETSc-based Approach**

Blockchain-based schemes for continuous verifiability and traceability of IoT data

Jesus Carretero and Cristhian Martinez

This paper will present a continuous delivery/continuous verifiability (CD/CV) framework for IoT dataflows in edge–fog–cloud. I this framework a CD model based on extraction, transformation, and load (ETL) mechanism as well as a directed acyclic graph (DAG) construction, enable end-users to create efficient schemes for the continuous verification and validation of the execution of applications in edge–fog–cloud infrastructures. This framework also provides tools for continuous verification and validation (CV) of predefined execution sequences and the integrity of digital assets using blockchain. CV model converts ETL and DAG into business model, smart contracts in a private blockchain for the automatic and transparent registration of transactions performed by each application in workflows/pipelines created by CD model without altering applications nor edge–fog–cloud workflows. This framework ensures that IoT dataflow delivers verifiable information for organizations to conduct critical decision-making processes with certainty. A containerized parallelism approach solves portability issues and reduces/compensates the overhead produced by CD/CV operations. The talk will also present evaluation results of the CD/CV framework based on a case study where user mobility information is used to identify interest points, patterns, and maps. The experimental evaluation results show the feasibility of CD/CV to register transactions performed in IoT dataflows through edge–fog–cloud in a private blockchain network.

Network accelerated in-memory ad-hoc file system for data-centric and high-performance applications.

Javier Garcia Blas, Cosmin Octavian Petre, Genaro Juan Sanchez Gallegos and Jesus Carretero:

New data-intensive applications hugely increase the demands on the performance and capacity of HPC storage systems. This growing need for data processing is accompanied by many improvements in storage technologies like the introduction of NVMe and persistent memory to support highest performance requirements. Nevertheless, putting all data on ultra-fast storage devices is not cost efficient, so that it becomes necessary to build multi-tier storage hierarchies and to move data between tiers according to their hotness. One approach to overcome the aforementioned problem is ah-doc file systems. Ad-hoc file systems utilise compute node's available storage (i.e., memory, persistent storage) to create a temporary storage system according to the application behaviour in the HPC environment. This work presents the high efficient communication model introduced in Hercules, an in-memory ad-hoc file system. The initial prototype takes advantage of the Unified Communication X framework (UCX), which enables a generic interface layer for network access. The current implementation exploits the capabilities offered by Infiniband protocol such as RDMA and zero-copy transfers. Preliminary evaluation results demonstrate a high grade of network utilisation.

Energy-aware malleable scheduling techniques.

David E. Singh, Alvaro Arbe Milara and Jesus Carretero

This work introduces an application scheduler that is able to determine the application configuration that minimizes its energy consumption. This scheduler is used in combination with DVFS techniques and FlexMPI, a runtime that provides malleable and monitoring capabilities to MPI applications. The work describes the framework that integrates all these components, including two scheduling algorithms and one application modeling methodology. The first scheduling algorithm determines the application minimum energy consumption, while the second one uses a cost equation for balancing performance and energy consumption. The application modeling is performed in run-time and allows us to characterize the energy profile of the executing applications. These models are used to support the scheduler's decision-making process. All these methods and algorithms are evaluated in a real platform. Results show that according to the application profile (CPU-intensive, I/O-intensive or communication-intensive) it is possible to determine the best application configuration (number of processes and DVFS value) that optimizes the memory consumption.

Scaling of the SVD Algorithm for HPC Science: A PETSc-based Approach.

Paula Ferrero-Roza, José A. Moríñigo and Filippo Terragni

The Singular Value Decomposition (SVD) algorithm is ubiquitous in many fields of science and technology. It may be used embedded into other advanced algorithms, solvers or data processing chains. In those scenarios dealing with large data volumes expressed as a huge matrix, there is the need of a parallel SVD version to process it efficiently. We present some ideas and results obtained within the PETSc framework which provides promising HPC scalable solvers. The analysed SVD implementations have been taken from the SLEPc library, plugged into PETSc to extend its functionality. A preliminary assessment and performance comparison is given.

WEDNESDAY, 1 MARCH

AWS ACADEMY 11:30 AM - 1:00 PM

SPECIAL SESSION. Scalable Algorithms, Libraries and Tools for Computational Science and Machine Learning on new Heterogeneous HPC Systems (HSALTCSMLNHPC2023)

CHAIR: Salvatore Cuomo (University of Naples Federico II)

- Raúl Marichal, Guillermo Toyos, Ernesto Dufrechou, and Pablo Ezzatti: **Evaluation of architecture-aware optimization techniques for Convolutional Neural Networks**

- Bruno Galluzzi, Stefano Izzo, Fabio Giampaolo, Salvatore Cuomo, Marco Vanoni, Lilia Alberghina, Chiara Damiani, and Francesco Piccialli: **Coupling constrained-based flux sampling and clustering to tackle cancer metabolic heterogeneity**

- Kashif Qureshi, Noman Arshad, and Thomas Newe: **Intrusion Detection Systems for Cyber Attacks Detection in Power Line Communications Networks**

- Jia-Hao Syu, Jerry Chun-Wei Lin, Marcin Fojcik, and Rafal Cupek: **HTPS: Heterogeneous Transferring Prediction System for Healthcare Datasets**

Evaluation of architecture-aware optimization techniques for Convolutional Neural Networks.

Raúl Marichal, Guillermo Toyos, Ernesto Dufrechou and Pablo Ezzatti

The growing need to perform Neural network inference with low latency is giving place to a wide spectrum of heterogeneous devices with deep learning capabilities. Therefore,obtaining the best performance from each device and choosing the most suitable platform for a given problem has become challenging. This paper evaluates multiple inference platforms using architecture-aware optimizations for convolutional neural networks. Specifically, we use TensorRT and OpenVINO frameworks for hardware optimizations on top of the platform-aware NetAdapt algorithm. The experimental evaluation shows that on MobileNet and AlexNet, using NetAdapt with TensorRT or OpenVINO can improve latency up to 10× and 5.3× respectively. Moreover, a throughput test using different batch sizes showed variable performance improvement on the different devices. The discussion of the experimental results can be useful to guide the selection of devices and optimizations for different AI solutions.

Coupling constrained-based flux sampling and clustering to tackle cancer metabolic heterogeneity.

Bruno Galluzzi, Stefano Izzo, Fabio Giampaolo, Salvatore Cuomo, Marco Vanoni, Lilia Alberghina, Chiara Damiani and Francesco Piccialli

Characterizing the heterogeneity of cancer metabolism requires the knowledge of metabolic fluxes in different tumor types. These fluxes cannot be directly determined, especially at a sub-cellular level. Still, they can be obtained numerically through constraint-based steady-state models after integrating other high-throughput -omics data, such as transcriptomics. In this work, we proposed to study cancer metabolism through data analysis and machine learning methodologies. To this aim, we considered transcriptomics profiles for a large set of cancer cells. Using a core metabolic network as a scaffold, we generated many feasible flux distributions for each cancer cell. Then, we used cluster analysis to analyze these data. This preliminary analysis revealed three well-separated clusters having different metabolic behaviors.

Intrusion Detection Systems for Cyber Attacks Detection in Power Line Communications Networks.

Kashif Qureshi, Noman Arshad and Thomas Newe.

Power Line Communication (PLC) is categorized into wired and wireless technologies to distribute the power and transmit the data at different frequency ranges. System administration is one of the significant areas in these networks to manage communication processes. Security is one of the significant concerns which leads to slow and unavailability of services, false and altered instructions, malfunctioning, and abnormal behavior of systems. Intrusion Detection Systems (IDS) is used to handle security attacks and assure the security of the system. However, the existing IDS systems have limited capabilities to handle the new attacks. This paper proposes a Machine Learning (ML) algorithm for IDS systems used in PLC networks to improve the overall system performance and detect the vulnerabilities of the system. The proposed system can detect the latest assaults and protect the systems from unauthorized and malicious activities. The proposed IDS system is assessed by using a virtual environment using the latest dataset and compared with existing traditional systems. The experiment results indicated the better performance of the proposed system to handle the new assaults and protect the systems.

HTPS: Heterogeneous Transferring Prediction System for Healthcare Datasets.

Jia-Hao Syu, Jerry Chun-Wei Lin, Marcin Fojcik and Rafal Cupek.

The medical internet of things leads to revolutionary improvements in medical services, also known as smart healthcare. With the big healthcare data, data mining and machine learning can assist wellness management and intelligent diagnosis, and achieve the P4-medicine. However, healthcare data has high sparsity and heterogeneity, which is not a trivial task to handle it. In this paper, we propose a Heterogeneous Transferring Prediction System (HTPS). We first apply the feature engineering mechanism to transform the dataset into sparse and dense feature matrices. After that, the autoencoders in the embedding networks then not only embed features but also transfer knowledge from heterogeneous datasets. Experimental results showed that the proposed HTPS outperforms the benchmark models on various prediction tasks and datasets, and ablation studies present the effectiveness of each designed mechanism. Results also demonstrated the negative impact of heterogeneous data on benchmark models and the high transferability of the proposed HTPS.

WEDNESDAY, 1 MARCH

TEATRO 3:30 PM - 5:30 PM

MAIN TRACK.

CHAIR: Massimo Torquati (University of Pisa)

- Lucas Leandro Nesi, Vinícius Garcia Pinto, Lucas Mello Schnorr, and Arnaud Legrand: **Summarizing task-based applications behavior over many nodes through progression clustering**

- Adriano Vogel, Marco Danelutto, Dalvan Griebler, and Luiz Fernandes: **Revisiting self-adaptation for efficient decision-making at run-time in parallel executions**

- Franz Biersack, Kilian Holzinger, Henning Stubbe, Thomas Wild, Georg Carle, and Andreas Herkersdorf: **Priority-aware Inter-Server Receive Side Scaling**

- Pasqua D'Ambra, Fabio Durastante, S M Ferdous, Salvatore Filippone, Mahantesh Halappanavar, and Alex Pothen: **AMG Preconditioners based on parallel hybrid coarsening and multi-objective graph matching**

Summarizing task-based applications behavior over many nodes through progression clustering.

Lucas Leandro Nesi, Vinícius Garcia Pinto, Lucas Mello Schnorr and Arnaud Legrand.

Visualization strategies are a valuable tool in the performance evaluation of HPC applications. Although the traditional Gantt charts are a widespread and enlightening strategy, it presents scalability problems and may misguide the analysis by focusing on resource utilization alone. This paper proposes an overview strategy to indicate nodes of interest for further investigation with classical visualizations like Gantt charts. For this, it uses a progression metric that captures work done per node inferred from the task-based structure, a time-step clustering of those metrics to decrease redundant information, and a more scalable visualization technique. We demonstrate with six scenarios and two applications that such a strategy can indicate problematic nodes more straightforwardly while using the same visualization space. Also, we provide examples where it correctly captures application work progression, showing application problems earlier and as an easy way to compare nodes. At the same time that traditional methods are misleading.

Revisiting self-adaptation for efficient decision-making at run-time in parallel executions.

Adriano Vogel, Marco Danelutto, Dalvan Griebler and Luiz Fernandes.

Self-adaptation is a potential alternative to provide a higher level of autonomic abstractions and run-time responsiveness in parallel executions. However, the recurrent problem is that self-adaptation is still limited in flexibility and efficiency. For instance, there is a lack of mechanisms to apply adaptation actions and efficient decision-making strategies to decide which configurations should be conveniently enforced at run-time. In this work, we are interested in providing and evaluating potential abstractions achievable with self-adaptation transparently managing parallel executions. Therefore, we provide a new mechanism to support self-adaptation in applications with multiple parallel stages executed in multi-cores. Moreover, we reproduce, reimplement, and evaluate in our scenario an existing decision-making strategy. The observations from the results show that the proposed mechanism for self-adaptation can provide new parallelism abstractions and autonomous responsiveness at run-time. On the other hand, there is a need for more accurate decision-making strategies to enable efficient executions of applications in resource-constrained scenarios like multi-cores.

Priority-aware Inter-Server Receive Side Scaling.

Franz Biersack, Kilian Holzinger, Henning Stubbe, Thomas Wild, Georg Carle and Andreas Herkersdorf.

Next-generation automotive networks will be characterized by a high number of interconnected sensors, actuators and applications on electronic control units communicating with each other over a high-speed Ethernet backbone network. As these applications have various criticalities, high volumes of fluctuating traffic with different priorities will have to be processed in a reliable and efficient manner. To cope with these challenges, we present Priority-aware Inter-Server Receive Side Scaling (prioRSS), a new SmartNIC-based hardware accelerator designed for automotive compute nodes. prioRSS builds upon Receive Side Scaling and introduces priority-awareness into an intra- and inter-node load balancer. It uses a priority-partitioned indirection table within which flows of the same priority are bundled. Low-latency reconfigurations issued by a Network Health Monitoring software allow for adapting the table content to changing network conditions. Simulative evaluations and comparisons to a priority-unaware version of our design show that prioRSS enables per-priority resource assignments without degrading end-to-end packet latencies while using the same table memory space. Paired with a priority-aware scheduler, end-to-end latencies of high priority flows can be notably reduced compared to average packet latencies, at the expense of lowest priority traffic. The best results are acquired when partitioning the table proportionally to the associated traffic share.

AMG Preconditioners based on parallel hybrid coarsening and multi-objective graph matching.

Pasqua D'Ambra, Fabio Durastante, S M Ferdous, Salvatore Filippone, Mahantesh Halappanavar and Alex Pothen.

We describe preliminary results from a multi-objective graph matching algorithm, in the coarsening step of an aggregation-based Algebraic MultiGrid (AMG) preconditioner, for solving large and sparse linear systems of equations on high-end parallel computers. We have two objectives. First, we wish to improve the convergence behavior of the AMG method when applied to highly anisotropic problems. Second, we wish to extend the parallel package \emph{Omitted for blind review} to exploit multi-threaded parallelism at the node level on multi-core processors. Our matching proposal balances the need to simultaneously compute high weights and large cardinalities by a new formulation of the weighted matching problem combining both these objectives using a parameter λ. We compute the matching by a parallel $2/3-\varepsilon$-approximation algorithm for maximum weight matchings. Results with the new matching algorithm show that for a suitable choice of the parameter λ we compute effective preconditioners in the presence of anisotropy, i.e., smaller solve times, setup times, iterations counts, and operator complexity.

WEDNESDAY, 1 MARCH

SALA PROCIDA 3:30 PM - 5:30 PM

SPECIAL SESSION. Big Data Convergence: From Sensors to Applications (BDCSA2023)

CHAIR: Katzalin Olcoz (Universidad Complutense de Madrid)

- Javier Campoy, Ignacio-Iker Prado-Rujas, José L. Risco-Martín, Katzalin Olcoz, and María S. Pérez: **Distributed training and inference of deep learning solar energy forecasting models**

- Alvaro Cuartero-Montilla and Rafael Mayo-García: **Application of advanced Artificial Intelligence methodologies for the development of a gene therapy for the primary Hyperoxaluria**

- Tommaso Marinelli, José Ignacio Gómez Pérez, Christian Tenllado, and Francky Catthoor: **Efficiency-Aimed Pattern Analysis and Data Mapping in Hybrid Cache-SPM Architectures**

- Elías Del-Pozo-Puñal, Felix Garcia-Carballeria, and Diego Camarmas-Alonso: **ENIGMA: A Scalable Simulator for IoT and Edge Computing**

Distributed training and inference of deep learning solar energy forecasting models.

Javier Campoy, Ignacio-Iker Prado-Rujas, José L. Risco-Martín, Katzalin Olcoz and María S. Pérez.

Different accurate predictive models have been developed to forecast the amount of solar energy produced in a given area. These models are usually run in a centralized manner, considering irradiance inputs taken from a set of sensors that are deployed in that area. CAIDE is a framework that supports the deployment and analysis of solar plants following Model Based System Engineering (MBSE) and Internet of Things (IoT) methodologies. However, the current solution performs the training and inference phases of the solar energy forecasting models in a central way, not taking advantage of the distributed environment modeled by means of CAIDE. This work presents an extension of CAIDE that allows us to distribute the training and inference phases, obtaining performance improvements, and achieving a greater adaptation to the inherently distributed topology of the deployment of the sensors.

Application of advanced Artificial Intelligence methodologies for the development of a gene therapy for the primary Hyperoxaluria.

Alvaro Cuartero-Montilla and Rafael Mayo-García.

The objective of this work has been the application of deep learning techniques in the analysis of Primary Hyperoxaluria through histological images taken from kidneys of a mouse model of the disease at the Biomedical Innovation Unit of CIEMAT. For this purpose, Transformers type models have been used to classify the histologies depending on the degree of severity of the disease. On the other hand, it has been studied the feasibility of using generative models, specifically models of diffusion, for the synthetic generation of images of renal tissue. The general results have been satisfactory, with models capable of classifying images with around 80% accuracy and models capable of generating images positively validated by experts. These results will help in the final efficacy evaluation of this gene therapy protocol to fight against Primary Hyperoxaluria.

Efficiency-Aimed Pattern Analysis and Data Mapping in Hybrid Cache-SPM Architectures.

Tommaso Marinelli, José Ignacio Gómez Pérez, Christian Tenllado and Francky Catthoor.

Integrated memory plays an extremely important role in modern CPUs. It enables the creation of local copies of main memory data, to reduce their access cost in terms of latency and energy. Multiple levels of cache are typically present, which operate transparently to the user and exploit locality principles to maximize their utilization. Despite being effective in many cases, not all applications make a fair use of the caches, and some improvement could be achieved with a modified memory hierarchy. Specifically, adding one or more scratchpad memories (SPMs), which have a simpler design since they lack automatic hardware management, can help reduce the overall energy consumption of the system where caches fail to shine. In this work, a methodology to profile the memory accesses of applications is introduced, which allows to detect potentially cache-inefficient patterns. Performed simulations show that a different placement of data in a hybrid cache-SPM system can have a positive impact on the overall energy consumption.

Elías Del-Pozo-Puñal, Felix Garcia-Carballeria and Diego Camarmas-Alonso.

Tommaso Marinelli, José Ignacio Gómez Pérez, Christian Tenllado and Francky Catthoor.

In recent years, the devices that make up the Internet of Things (IoT) paradigm have been increasing in number and complexity as different layers of network computing, such as Cloud, Edge and Fog Computing, have emerged. Thanks to these latest paradigms and their different types of communications, it has been possible to reduce and distribute the computational load on the network. However, to build these infrastructures and reduce the cost, it is necessary to use a simulation platform to model these environments and analyze their behavior (power consumption, CPU usage, bandwidth, etc.) beforehand. An essential aspect of simulators is scalability, the ability to add new components and simulate large infrastructures without compromising the performance of the simulator. Many existing simulators, as will be discussed in this article, cannot scale adequately as new elements are added to the system. Furthermore, in these environments, it is useful to simulate mobile devices and to know the location of these devices for the development and analysis of new algorithms. To solve these problems, this article presents and describes the ENIGMA simulator. ENIGMA is a scalable simulator of Edge, Fog and Cloud computing infrastructures, which allows to efficiently simulate a large number of devices and elements and to analyze different characteristics (CPU usage, power consumption, network bandwidth, application execution time, etc.). Its objective is to analyze new infrastructures and algorithms. ENIGMA also includes support for including mobile devices and an API to integrate their visualization in graphical maps. The article presents an evaluation to compare the scalability of ENIGMA with other state-of-the-art simulators and different use cases that illustrate its operation.

WEDNESDAY, 1 MARCH

AWS ACADEMY 3:30 PM - 5:30 PM

SPECIAL SESSION. Scalable Algorithms, Libraries and Tools for Computational Science and Machine Learning on new Heterogeneous HPC Systems (HSALTCSMLNHPC2023)

CHAIR: Francesco Piccialli (University of Naples Federico II)

- Nicolo Romandini, Carlo Mazzocca, and Rebecca Montanari: **Federated Learning Meets Blockchain: a Power Consumption Case Study**
- Maria Pia De Rosa, Fabio Giampaolo, Francesco Piccialli, and Salvatore Cuomo: **Modelling the COVID-19 infection rate through a Physics-Informed learning approach**
- Kevin Crampon, Alexis Giorkallos, Stéphanie Baud, and Luiz Angelo Steffenel: **Convolutional graph neural network training scalability for molecular docking**
- Jie Lei, José Flich, and Enrique S. Quintana-Ortí: **Toward Matrix Multiplication for Deep Learning Inference on the Xilinx Versal**
- Tao Tao: **Synchronization Efficient Scheduling of Fine-grained Irregular Programs**

Federated Learning Meets Blockchain: a Power Consumption Case Study.

Nicolò Romandini, Carlo Mazzocca and Rebecca Montanari.

Federated learning (FL) is emerging as the most promising approach to collaboratively train a machine learning (ML) model on a common task without centralizing data. During each FL round, participants locally train a partial model with its on-premises data. Such models are subsequently aggregated to derive a global one. How these partial models are combined is a primary concern. Traditional approaches usually rely on a parameter server that introduces many weaknesses such as single point of failure, lack of trustworthiness among unknown participants, and incapacity to handle the traffic generated from millions of devices. Thus, to overcome such concerns, blockchain has recently been proposed as a valuable solution to improve the robustness of FL approaches. The full-blown benefits of using blockchain enable tackling the limits of centralized servers. However, energy consumption is still one of the significant factors inhibiting its widespread due to the current discussions on climate change and sustainability. Recently, a growing number of research works have been focusing on integrating FL and blockchain, nevertheless, adequate analysis and estimate of their energy and power consumption are often lacking. This paper presents an estimate of the power consumption of FlowChain, an architecture that integrates FL with blockchain to simplify the use of FL. Experimental results demonstrate that the overall power consumption significantly depends on the ML model adopted.

Modelling the COVID-19 infection rate through a Physics-Informed learning approach.

Maria Pia De Rosa, Fabio Giampaolo, Francesco Piccialli and Salvatore Cuomo.

Over the past two years, the COVID-19 pandemic has been one of the most frequently and hotly debated social topic. From lockdowns and restrictions to radical changes in the way of working and socializing due to social distancing and masks wearing, the still ongoing pandemic has had a strong impact on people's life and psychological health. The well-known infection rate (Rt) has been the main tool used by national and local governments worldwide to handle the pandemic, especially for defining containment policies (lockdowns, social distancing, intermittent regional strategies, etc) which have definitely affected the social life. In the present paper we propose an artificial intelligence (AI) approach for the modelling of the COVID-19 infection rate Rt, by exploiting the novel methodology of the Physics-Informed Neural Networks (PINNs) to compute the (susceptible-infected-dead-recovered) SIDR model. The predictions regarding the susceptible, infected, dead and recovered of the next 30 days has been done to test the Rt prediction. Lastly, a new methodology to compute the infection rate are proposed for future works.

Convolutional graph neural network training scalability for molecular docking.

Kevin Crampon, Alexis Giorkallos, Stéphanie Baud and Luiz Angelo Steffenel.

Deep learning use is growing in many numerical simulation fields, and drug discovery does not escape this trend. Indeed, before proceeding with in vitro and then in vivo experiments, drug discovery now relies on in sillico techniques such as molecular docking to narrow the number of experiments and identify the best candidates. This method explores the receptor surface and the ligand's conformational space providing numerous ligand-receptor poses. All these poses are then scored and ranked by a scoring function allowing to predict the best poses among all, then compare different ligands regarding a given receptor or different targets regarding a given ligand. Since the 2010s, numerous deep learning methods have been used to tackle this problem. Nowadays, there are two significant trends in deep learning for molecular docking: (i) the augmentation of available structural data and (ii) the use of a new kind of neural network: the graph convolutional neural networks (GCNs). In this paper, we propose the study of training scalability of a GCN—a molecular complex scoring function—on an increasing number of GPUs and with a variety of batch sizes. After a hyperparameter analysis, we achieve an 80% reduction in the training time, but this improvement sometimes involves a performance metrics degradation that the final users must ponder.

Toward Matrix Multiplication for Deep Learning Inference on the Xilinx Versal.

Jie Lei, José Flich and Enrique S. Quintana-Ortí.

The remarkable positive impact of Deep Neural Networks on many Artificial Intelligence (AI) tasks has led to the development of various high performance algorithms as well as specialized processor architectures and accelerators for this purpose. In this paper we address this scenario by demonstrating that the optimization principles underlying the modern realization of the general matrix multiplication (GEMM) in conventional processor architectures, are also valid to achieve high performance for the type of operations that arise in deep learning (DL) on an exotic accelerator such as the AI Engine (AIE) tile embedded in many Xilinx Versal platforms. In particular, our experimental results with a prototype implementation of the GEMM kernel, on a Xilinx Versal VCK190, delivers performance close to 78% of the theoretical peak that can be expected on an AIE tile, for 8-bit integer operands, when using Xilinx high level library.

Synchronization Efficient Scheduling of Fine-grained Irregular Programs.

Tao Tao.

This paper discusses the theory behind the global rebalancing policy, a new emerging paradigm of scheduling dynamic, irregular programs. According to this policy, task workload is distributed in rebalancing sessions enabled by global thread barriers, while traditional approaches such as work stealing rely on localized, concurrent deque operations. We show in theory that the parallel execution model based on the global rebalancing policy has an amortized running time bound of $\mathcal O(T_1/P+T_\infty)$, including all synchronization overhead. Based on this result, we further conclude that the global rebalancing policy asymptotically outperforms traditional work stealing as long as the input program is sufficiently parallel. Compared to the state of the art, the global rebalancing policy enjoys the advantage of having a worst-case running time guarantee, while the state of the art has no well-defined worst-case behavior. The worst-case guarantee enables important applications in deadline-sensitive and mission-critical applications.

THURSDAY, 2 MARCH

TEATRO 11:00 AM - 12:30 PM

MAIN TRACK

CHAIR: Marco Aldinucci (University of Torino)

- Keisuke Sugiura and Hiroki Matsutani: **An Efficient Accelerator for Deep Learning-based Point Cloud Registration on FPGAs**
- Midia Reshadi and David Gregg: **Dynamic Resource Partitioning for Multi-Tenant Systolic Array Based DNN Accelerator**
- Jorge Villarrubia, Luis Costero, Francisco D. Igual, and Katzalin Olcoz: **Improving inference time in multi-TPU systems with profiled model segmentation**
- Alberto Ottimo, Gabriele Mencagli, and Marco Danelutto: FSP: **a Framework for Data Stream Processing Applications targeting FPGAs**

An Efficient Accelerator for Deep Learning-based Point Cloud Registration on FPGAs.

Keisuke Sugiura and Hiroki Matsutani.

Point cloud registration is the basis for many robotic applications such as odometry and Simultaneous Localization And Mapping (SLAM), which are increasingly important for autonomous mobile robots. The limitation of computational resources and power budgets on such robots motivates us to study the resource-efficient registration method on low-cost edge devices. In this paper, we propose an FPGA-based novel pipeline for 3D point cloud registration built upon a recent deep learning-based method, PointNetLK. Based on the profiling results, we focus on the PointNet feature extraction as it becomes a major bottleneck; we improve its scalability and memory-efficiency by consuming each input point one-by-one in a pipelined manner instead of processing the whole point cloud at once. We then design a fully-parallelized and pipelined accelerator consisting of a custom PointNet IP core, which fits within both low-cost and mid-range FPGAs (e.g., Avnet Ultra96v2 and Xilinx ZCU104). Experimental results show that our proposed pipeline achieves up to 21.34x and 69.60x faster registration speed than the vanilla PointNetLK and ICP, respectively, while only consuming 722mW and maintaining the same level of accuracy.

Dynamic Resource Partitioning for Multi-Tenant Systolic Array Based DNN Accelerator.

Midia Reshadi and David Gregg.

Deep neural networks (DNN) have become a significant applications in both cloud-server and edge devices. Meanwhile, the growing number of DNNs on those platforms raises the need to execute multiple DNNs on the same device. This paper proposes a dynamic partitioning algorithm to perform concurrent processing of multiple DNNs on a systolic-array-based accelerator. Sharing an accelerator's storage and processing resources across multiple DNNs increases resource utilization and reduces computation time and energy consumption. To this end, we made a minor modification to the processing element. We also propose a dynamic partitioning algorithm to divide the execution units of a weight-stationary systolic array between two or more DNNs. We evaluate the energy consumption and computation time with both heavy and light workloads. Simulation results show a 35% and 62% improvement in energy consumption and 56% and 44% in computation time under heavy and light workloads, respectively, compared with single tenancy.

Improving inference time in multi-TPU systems with profiled model segmentation.

Jorge Villarrubia, Luis Costero, Francisco D. Igual and Katzalin Olcoz.

In this paper, we systematically evaluate the inference performance of the Edge TPU by Google for neural networks with different characteristics. Specifically, we determine that, given the limited amount of on-chip memory on the Edge TPU, accesses to external (host) memory rapidly become an important performance bottleneck. We demonstrate how multiple devices can be jointly used to alleviate the bottleneck introduced by accessing the host memory. We propose a solution combining model segmentation and pipelining on up to four TPUs, with remarkable performance improvements that range from 6× for neural networks with convolutional layers to 46× for fully connected layers, compared with single-TPU setups.

FSP: a Framework for Data Stream Processing Applications targeting FPGAs.

Jorge Villarrubia, Luis Costero, Francisco D. Igual and Katzalin Olcoz.

FPGA architectures are becoming popular because of their high performance-to-energy ratio. Nonetheless, their effective exploitation is often counterbalanced by a high programming effort, since most of the modern hardware description languages provide only low-level programming abstractions. This paper proposes FSP, a framework to productively support the development of Data Stream Processing applications on CPU+FPGA System-on-Chip devices (SoCs). By exploiting a code generation approach starting from a high-level DSL in Python, FSP generates an efficient OpenCL skeleton implementation of the parallel pipeline on FPGA and the library to be used by host programs to transfer inputs and collect results to/from the FPGA program. The experimental results showcase the effectiveness of FSP on an SoC equipped with an Intel Arria 10 FPGA by running two streaming benchmark applications.

THURSDAY, 2 MARCH

SALA PROCIDA 11:00 AM - 12:30 PM

SPECIAL SESSION: High Performance Computing in Modelling and Simulation (HPCMS 23)

CHAIR: William Spataro (University of Calabria)

- Andrea Giordano, Donato D'Ambrosio, Davide Macrì, Rocco Rongo, William Spataro, Gladys Utrera, and Marisa Gil: **OpenCAL++: An object-oriented architecture for transparent Parallel Execution of Cellular Automata models**

- Luca Barillaro, Giuseppe Agapito, and Mario Cannataro: **Using Edge-based Deep Learning Model for Early Detection of Cancer**

- Lorella Bottino, Marzia Settino, and Mario Cannataro: **Distributed ICT solutions for scoliosis management**

- Alessio De Rango, Luca Furnari, Alfonso Senatore, Giuseppe Mendicino, Andrea Giordano, Davide Macrì, Gladys Utrera, and Donato D'Ambrosio: **Performance Analysis and Optimization of the CUDA Implementation of the Three-Dimensional Subsurface XCA-Flow Cellular Automaton**

OpenCAL++: An object-oriented architecture for transparent Parallel Execution of Cellular Automata models.

Andrea Giordano, Donato D'Ambrosio, Davide Macrì, Rocco Rongo, William Spataro, Gladys Utrera and Marisa Gil.

Cellular Automata (CA) models, initially studied by John von Neumann, have been developed by numerous researchers and applied in both academic and scientific fields. Thanks to their local and independent rules, simulations of complex systems can be easily implemented based on CA modelling on parallel machines. However, due to the heterogeneity of the components - from the hardware to the software perspective-the various possible scenarios running parallelism in today's architectures can pose a challenge in such implementations, making it difficult to exploit. This paper presents OpenCAL++, a transparent and efficient object-oriented platform for the parallel execution of cellular automata models. The architecture of OpenCAL++ ensures the modeller a fully transparent parallel execution and a strong "separation of concerns" between the execution parallelism issues and the model implementation. The code implementing the Cellular Automata model remains the same whether the execution performs in a shared-, distributed-memory or a GPGPU context, irrespective of the optimizations adopted. To this aim, the object-oriented paradigm has been intensely exploited. As well as the OpenCAL++ architecture, we present the description of a simple Cellular Automata model implementation for illustrative purposes.

Using Edge-based Deep Learning Model for Early Detection of Cancer.

Luca Barillaro, Giuseppe Agapito and Mario Cannataro.

Cancer is one of the most frequent causes of death in the world. Usually, cancer can be easily diagnosed if characteristic symptoms occur. However, many people who are suffering from cancer have no symptoms.Early diagnosis of tumors is essential to contrast their progression, helping to define more effective treatments to provide long-term survival. Early cancer detection is effective if sensible data can be investigated through high-performance technologies like edge computing. Edge computing is a new paradigm for analyzing data as close to the source as possible, avoiding exporting them outside. Hence, edge-based deep learning models can be applied to improve early cancer detection. This paper provides an use case of a classification task on tumor-related data based on the famous UCI machine learning data sets repository using a deep learning approach based on edge computing. In addition, the manuscript provides an overview of the edge computing paradigm, highlighting its advantages and usability. We also described a small experiment with real tumor data to characterize performance considerations. Moreover, the presented model can be used with different data types, such as images, EGC, and ECC signals.

Distributed ICT solutions for scoliosis management.

Lorella Bottino, Marzia Settino and Mario Cannataro.

Scoliosis is a curvature of the spine often found in adolescents. Commonly the management of patients with scoliosis is done through manual methods. The use of smartphone applications with integrated sensors allows a scoliosis managememt remotely. Scoliosis management moves from a centralized system to a decentralized system, with obvious benefits for both the doctor and the patient.

Performance Analysis and Optimization of the CUDA Implementation of the Three-Dimensional Subsurface XCA-Flow Cellular Automaton.

Alessio De Rango, Luca Furnari, Alfonso Senatore, Giuseppe Mendicino, Andrea Giordano, Davide Macrì, Gladys Utrera and Donato D'Ambrosio.

We present the results of a performance assessment and optimisation work regarding the CUDA implementation of the three-dimensional XCA-Flow subsurface Extended Cellular Automata model. To this end, we have considered a ten days long simulation already considered in previous works, characterised by a constant infiltration rate and a heterogeneous hydraulic conductivity field, as the benchmark. We ran the experiments on the Nvidia V100 high-performance many-core device. We have analysed essential aspects of the XCA-Flow model by updating its kernels. We applied classical tiling/shared memory techniques to the stencil-based and reduction kernels in the first step. Results suggested applying a thorough analysis of the model. Both theoretical and experimental assessments have driven this analysis, which has pointed out the need to increase the achieved warp occupancy to speed up the computation. The resulting general redesign of the application allowed for a 20.3\% mean performance gain (over the CUDA block configurations considered). We also performed two Roofline analyses to characterise the kernels of the original and improved implementations in terms of arithmetic intensity and performance. Besides the improved performance, we have obtained meaningful insights about the CUDA implementation of the XCA-Flow model that could, in principle, allow for further optimisations.

THURSDAY, 2 MARCH

AWS ACADEMY 11:00 AM - 12:30 PM

TUTORIAL: FlexMPI: Malleability Techniques and Applications in High-Performance Computing

E4 & ADMIRE

The current static usage model of HPC systems is becoming increasingly inefficient. This is driven by the continuously growing complexity and heterogeneity of system architectures, in combination with the increased usage of coupled applications, the need for strong scaling with extreme scale parallelism, and the increasing reliance on complex and dynamic workflows. As a consequence, we see a rise in research on malleable systems, middleware software and applications, which can adjust resource usage dynamically in order to extract a maximum of efficiency. Malleability allows systems to dynamically adjust the computation and storage needs of applications, on the one side, and the global system on the other. Such malleable systems, however, face a series of fundamental research challenges, including: who initiates changes in resource availability or usage? How is it communicated? How to compute the optimal usage? How can applications cope with dynamically changing resources? What should malleable programming models and abstractions look like? How to design resource management frameworks for malleable systems? What should be the API for applications? This tutorial will provide an in-depth presentation of emerging software designs to achieve malleability in high-performance computing, high-level parallel programming models, and programmability techniques to improve applications' malleability. The main part of the tutorial will be devoted to showing and demonstrating FlexMPI, a framework for HPC malleability, and Limitless, an HPC monitoring system to get information from applications and systems and the usage of AI and ML techniques to steer malleability in systems and applications. Finally, we will show how to apply the solutions presented to two use cases: Wacom++ and Nek5000.

Presenters:

- Jesús Carretero (Universidad Carlos III de Madrid)
- Alberto Cascajo (Universidad Carlos III de Madrid)
- David Expósito-Singh (Universidad Carlos III de Madrid)
- Raffaele Montella (University of Napoli "Parthenope")

THURSDAY, 2 MARCH

TEATRO 3:00 PM - 4:30 PM

CHAIR: Giuliano Laccetti (University of Naples Federico II)

MAIN TRACK

- Federica Uccello, Salvatore D'Antonio, Roberto Nardone, and Nicola Russo: **A Tamper-Resistant Storage Framework for Smart Grid security**

- Marco Danelutto, Paolo Palazzari, Alberto Ottimo, Gabriele Mencagli, and Francesco Iannone: **FastFlow targeting FPGAs**

- Ciro Giuseppe De Vita, Gennaro Mellone, Aniello Florio, Catherine Alessandra Torres Charles, Diana Di Luccio, Guido Benassai, Marco Lapegna, Giorgio Budillon, and Raffaele Montella: **Parallel and hierarchically-distributed Shoreline Alert Model (SAM)**

- Giuseppe Coviello, Kunal Rao, Gennaro Mellone, Ciro Giuseppe De Vita, Priscilla Benedetti, and Srimat Chakradhar: **Content-aware auto-scaling of stream processing applications on container orchestration platforms**

A Tamper-Resistant Storage Framework for Smart Grid security.

Federica Uccello, Salvatore D'Antonio, Roberto Nardone and Nicola Russo.

In the past few years, the energy sector has been among the most targeted by cyber-criminals. Due to the strong reliance of Critical Infrastructures on energy distribution, and the strategic value of such systems, the impact of intrusions and data breaches cannot be underestimated. In this scenario, data constitutes a critical asset to protect, especially as the latest technological development has led to interconnected intelligent systems, named smart grids. The consequences of data tampering, exposure or loss can range from disruption of essential services, to serious risks for environment, economy and people safety. Data provenance, as the documentation of the origin of data and the processes and methodology that led to it, can bring support when facing the aforementioned attacks. The present work aims to address security issues in the energy domain, by proposing the Advanced Tamper-Resistant Storage (ATRS), a novel framework for data provenance recording based on blockchain technology. The ATRS allows for the creation and storage of data provenance records, whose reliability is ensured by the tamper-resistance feature enabled through the combination of blockchain and TLS-based communication. The framework, tailored and tested for the smart grid domain, can easily be customized for different critical use cases.

FastFlow targeting FPGAs.

Marco Danelutto, Paolo Palazzari, Alberto Ottimo, Gabriele Mencagli and Francesco Iannone.

FastFlow is a structured parallel programming environment mainly targeting shared memory multi core machines. It currently provides some kind of support to include GPU offloading of data parallel computations. In this work we discuss an extension of FastFlow version 3.0 to include support for the offloading of task computation to FPGAs. In particular, we will show how FPGA task offloading to Xilinx boards is seamlessly integrated in a standard FastFlow application exploiting available FPGA kernels compiled using the Vitis tool-chain. Preliminary results demonstrating the feasibility of the integration as well as the absence of overhead with respect to hand coded pure OpenCL host applications using the very same FPGA kernels will be shown. Finally, some possibilities for FPGA kernel exploitation strictly tied to the necessity to use parallel design patterns in FastFlow applications will be discussed, outlining possibilities that otherwise require quite a programming effort by expert FPGA programmers.

Parallel and hierarchically-distributed Shoreline Alert Model (SAM).

Ciro Giuseppe De Vita, Gennaro Mellone, Aniello Florio, Catherine Alessandra Torres Charles, Diana Di Luccio, Guido Benassai, Marco Lapegna, Giorgio Budillon and Raffaele Montella.

In this paper, the Shoreline Alert Model (SAM) is presented as a component of a computation platform based on workflows dedicated to extreme weather/marine event simulation. This model aims to mitigate the effects of global change, giving the decision-makers, the scientists, and the engineers a novel, next-generation toolset for facing extreme weather events to implement the related management or emergency responses. In this paper, we present Shoreline Alert Model (SAM) leveraging a parallelization schema enabling the users to run it on heterogeneous parallel architectures. SAM has produced results ~24 times faster than the baseline when using shared memory with distributed memory dealing with about 20K transects along the Campania coastline.

Content-aware auto-scaling of stream processing applications on container orchestration platforms.

Giuseppe Coviello, Kunal Rao, Gennaro Mellone, Ciro Giuseppe De Vita, Priscilla Benedetti and Srimat Chakradhar.

Modern applications are designed as an interacting set of microservices, and these applications are typically deployed on container orchestration platforms like Kubernetes. There are several attractive features in Kubernetes that make it a popular choice for the deployment of applications, and automatic scaling is one such feature. The default horizontal scaling technique in Kubernetes is the Horizontal Pod Autoscaler (HPA). It scales each microservice independently while ignoring the interactions among the microservices in an application. In this paper, we show that ignoring such interactions by HPA leads to inefficient scaling, and the optimal scaling of different microservices in the applica- tion varies as the stream content changes. To automatically adapt to variations in stream content, we present a novel system called DataX AutoScaler that leverages knowledge of the entire stream processing application pipeline to efficiently auto-scale different microservices by taking into account their complex interactions. DataX AutoScaler runs as a Kubernetes operator to efficiently auto-scale the stream processing application pipelines. Unlike HPA, DataX AutoScaler does not require specification of any scaling-related configuration parameters for each microservice. Furthermore, DataX AutoScaler performs better scaling than HPA by balancing resources allocated to different microservices in the application pipeline such that there is minimal wastage of computing resources. Through experiments on real-world video analytics applications like face recognition and pose classification, we show that DataX AutoScaler is able to automatically adapt to variation in stream content and achieve up to 43% improvement in overall application performance when compared to a baseline system that uses HPA.

THURSDAY, 2 MARCH

SALA PROCIDA 3:00 PM - 4:30 PM

CHAIR: Giuseppe Trunfio (University of Sassari)

SPECIAL SESSION: High Performance Computing in Modelling and Simulation (HPCMS 23)

- Luca Barillaro, Giuseppe Agapito, and Mario Cannataro: **High performance deep learning libraries for biomedical applications**

- Natiele Lucca, Claudio Schepke, and Gabriel Dineck Tremarin: **Parallel Directives Evaluation in Porous Media Application: A Case Study**

- Arianna Anniciello and Elio Masciari: **A Judgment Aggregation Method For Fuzzy Multi Criteria Decision Making**

- Marjan Firouznia, Pietro Ruiu, and Giuseppe A. Trunfio: **Robust feature selection for high-dimensional datasets using a GPU-accelerated ensemble of cooperative coevolutionary optimizers**

High performance deep learning libraries for biomedical applications.

Luca Barillaro, Giuseppe Agapito and Mario Cannataro.

This paper aims to discuss about some libraries specifically designed to provide convenient high performance computing (HPC) oriented deep learning support to biomedical applications. Deep learning approaches are a topic of growing interest since they can achieve high precision in machine learning tasks and may be useful in several scenarios, while high performance computing is one of the driving factor for deep learning applications since they require massive computational power. One of these scenarios is related to biomedical context since the massive growth of data generated by several medical procedures. Deep learning applications on these data may be useful both for medical procedures both for further knowledge discovery in a particular field. Therefore the importance to have a deep learning library tailored for these task is evident. In this paper we describe two libraries developed inside a European project, the Deep Health Project, to support both deep learning basic operations and computer vision tasks oriented to a distributed computing fashion and with some special features for manage biomedical data.

Parallel Directives Evaluation in Porous Media Application: A Case Study.

Natiele Lucca, Claudio Schepke and Gabriel Dineck Tremarin.

High-performance computing provides the acceleration of scientific applications through the use of parallelism. Applications of this type usually demand a lot of computation time for a version with a single code execution stream. The adoption of different models of parallel programming is for developing concurrent code. In general, it has chosen the data processing division in scientific applications. On the other hand, task parallelism is an approach that different and independent perform computations could use. In this sense, this paper evaluates parallel interfaces and their programming models. Therefore, as a case study, we will be evaluated an application of porous media that provides the simulation of grain drying using OpenMP (loop, sections, tasks, target, and teams approach) and OpenACC programming interfaces. The results show a reduction in processing time in all test cases. The total parallel simulation time for a multicore architecture (16 physical cores) was 6.14 times less using loops, 5.96 using target, and 8.5 using teams. The reduction using a single GPU (Quadro M5000) was 8.38. We also contribute with some collected traces, identifying the parallel steps and synchronization time.

A Judgment Aggregation Method For Fuzzy Multi Criteria Decision Making.

Arianna Anniciello and Elio Masciari.

In most decision-making scenarios, the decision relies upon a multiplicity of factors/criteria which need to be selected and prioritized counting on the knowledge of a group of experts, with the aim of reaching a global decision. Through judgment aggregation models we seek a rational mechanism to reconcile individual judgments into a collective prioritization of requirements. In this paper we present a fuzzy multi-criteria group decision-making model, and experiment its application in credit rating determination, applying Majority Judgment as a method to consolidate the evaluations of a bank's board of direction into a shared and rational decision on SME rating

Robust feature selection for high-dimensional datasets using a GPU-accelerated ensemble of cooperative coevolutionary optimizers.

Marjan Firouznia, Pietro Ruiu and Giuseppe A. Trunfio.

Feature selection is an increasingly important step in the application of machine learning and knowledge discovery techniques to high-dimensional datasets. However, the growing complexity and size of datasets have made feature selection increasingly challenging, as selecting an optimal subset of features can be computationally very expensive, especially when a robust solution is required. To address this issue, we present an approach based on ensembles of cooperative coevolutionary optimisers and its parallelisation for hybrid multi-core CPU and GPU computation. The application of the developed algorithm to some typical high-dimensional datasets is discussed in the paper. According to the preliminary results, the proposed framework represents a valuable tool for addressing the computational challenges faced in feature selection, and it can be potentially applied to a wide range of machine learning and knowledge discovery tasks.

FRIDAY, 3 MARCH

TEATRO 11:00 AM - 12:30 AM

MAIN TRACK

CHAIR: Marco Lapegna (University of Naples Federico II)

- Aymar Cublier Martínez, Alejandro Álvarez Isabel, Jesús Carretero, and David E. Singh: **Fine-grained parallel social modelling for analyzing the COVID-19 propagation**
- Iker Martín Álvarez, José Ignacio Aliaga, Maribel Castillo, and Sergio Iserte: **Configurable synthetic application for studying malleability in HPC**
- Paulo Souza, Carlos Kayser, Lucas Roges, and Tiago Ferreto: Thea – **a QoS, Privacy, and Power-aware Algorithm for Placing Applications on Federated Edges**
- Gennaro Mellone, Ciro Giuseppe De Vita, Dante Domizzi Sánchez-Gallegos, Diana Di Luccio, Gaia Mattei, Francesco Peluso, Pietro Patrizio Ciro Aucelli, Angelo Ciaramella, and Raffaele Montella: **A containerized distributed processing platform for autonomous surface vehicles: preliminary results for marine litter detection**

Fine-grained parallel social modelling for analyzing the COVID-19 propagation.

Aymar Cublier Martínez, Alejandro Álvarez Isabel, Jesús Carretero and David E. Singh.

Agent-based epidemiological simulators have been proven to be one of the most successful tools for the analysis of the COVID-19 propagation. The ability of these tools to reproduce the behavior and interactions of each single individual leads to accurate and detailed results, that can be used to model fine-grained health-related policies like selective vaccination campaigns or immunity waning. One characteristic of these tools is the large amount of input data and computational resources and that they require. This relies on the development of parallel algorithms and methodologies for generating, accessing and processing large volumes of data from multiple data sources. This work presents a parallel workflow for extending the social modelling of OMITTED, an agent-based simulator. We have included two novel parallel social generation stages -that provide detailed and realistic social model- and one new visualization stage. The work presents a description of the algorithms used in each stage and a practical evaluation on a real platform. Results show that this contribution can be efficiently executed in parallel architectures and increases the simulation detail level, representing a significant advance in the simulator scenario modelling.

Configurable synthetic application for studying malleability in HPC.

Iker Martín Álvarez, José Ignacio Aliaga, Maribel Castillo and Sergio Iserte.

Nowadays, the throughput improvement in large clusters of computers recommends the development of malleable applications. Thus, during the execution of these applications in a job, the resource management system (RMS) can modify its resource allocation, in order to increase the global throughput. There are different alternatives to complete the different steps in which the reallocation of resources is decomposed. To find the best alternatives, this paper introduces a configurable synthetic iterative MPI malleable application capable of modifying, in execution time, the number of MPI processes according to several parameters. The application includes a performance module to measure stages time within steps, from processes management to data redistribution. In this way, the analysis of different scenarios will allow to conclude how the reconfiguration of application has to be made in different circumstances. At the same time, this tool can be used to create workloads that will allow to analyse the impact of malleability on a system and the work in progress.

Thea – a QoS, Privacy, and Power-aware Algorithm for Placing Applications on Federated Edges.

Paulo Souza, Carlos Kayser, Lucas Roges and Tiago Ferreto.

Federations between Edge Computing infrastructure providers represent a promising approach for improving the applications' Quality of Service (QoS) and the infrastructure's resource usage. At the same time, federated edges impose particular provisioning challenges, as data protection policies implemented by certain providers within a federation may conflict with the privacy requirements of services carrying out sensitive information (e.g., databases). In addition, the popularization of complex software architectures (e.g., composite applications) sets strict latency requirements that narrow the provisioning possibilities even further. Previous research efforts targeting federated edges have focused either on coupling with end-user performance requirements (e.g., latency and privacy) or on satisfying infrastructure providers' objectives (e.g., power consumption reduction), but none on balancing both. This paper presents Thea, a novel approach for provisioning composite applications on federated edges which optimizes applications' latency and privacy while reducing the infrastructure's power consumption. Simulated experiments show that Thea can achieve near-optimal results, reducing application latency and privacy issues by 50% and 42.11% and the infrastructure's power consumption by 18.95% compared to state-of-the-art approaches.

A containerized distributed processing platform for autonomous surface vehicles: preliminary results for marine litter detection.

Gennaro Mellone, Ciro Giuseppe De Vita, Dante Domizzi Sánchez-Gallegos, Diana Di Luccio, Gaia Mattei, Francesco Peluso, Pietro Patrizio Ciro Aucelli, Angelo Ciaramella and Raffaele Montella.

Autonomous Surface Vehicles and their management represent one of the significant challenges in coastal and offshore surveying. Although the development of this kind of data acquisition device has skyrocketed in the last few years, line guides and technological solutions still need to come. On the other hand, this kind of robotic vessel's true potential has yet to be explored. This paper presents ArgonautAI, a containerized distributed processing platform for autonomous surface vehicles. The proposed ArgonautAI architecture leverage a cluster of single-board computers with diverse and different characteristics (computing power, CUDA GPUs, FPGAs, GPIOs, PWMs, specialized I/O) orchestrated using Kubernetes and a customized programming interface. Furthermore, the proposed solution introduces two different types of containers: 1) the platform containers hosting the software life support for the platform and 2) the mission containers defined to support the survey mission-specific scopes. The firsts manage the vehicle's instruments (position, attitude, environment, depth), the data storage, the vessel-to-shore communication, and so on; the latter host mission-specific software components. Finally, as proof of concept of the proposed platform, we present an AI-based marine litter detection application using a hierarchical computer vision approach on heterogenic onboard computing resources.

FRIDAY, 3 MARCH

SALA PROCIDA 11:00 AM - 12:30 PM

SPECIAL SESSION: Cloud Computing on Infrastructure as a service and its Applications (CCIaaSA2023)

CHAIR: Attila Kettesz (University of Hungary)

- Lucía Pons, Salvador Petit, Julio Pons, Maria E. Gomez, Chaoyi Huang, and Julio Sahuquillo: Stratus: **A Hardware/Software Infrastructure for Controlled Cloud Research**

- Sezar Jarrous-Holtrup, Sergei Gorlatch, Michael Dey, and Folker Schamel: **Multi-Cloud Container Orchestration for High-Performance Real-Time Online Applications**

Stratus: A Hardware/Software Infrastructure for Controlled Cloud Research.

Lucía Pons, Salvador Petit, Julio Pons, Maria E. Gomez, Chaoyi Huang and Julio Sahuquillo.

Cloud systems deploy wide variety of shared resources and host a large amount of tenant applications. To perform cloud research, a small experimental platform is commonly used, which hides the huge system complexity and provides flexibility. In spite of being simpler, this platform should include the main cloud system components (hardware and software) in order to provide representative results. A wide set of platforms have spread on the recent years, however, most of them only include a major cloud component or lack the deployment of virtual machines (VMs) to provide isolation. This paper presents Stratus, an experimental platform that is currently being used to carry out cloud research. To the best of our knowledge, Stratus is the only platform that jointly provides three main features: uses VMs to provide isolation of tenant applications, deploys the three types of cloud nodes (server, client and storage) and manages all main shared system resources (CPUs, LLC space, memory, network and disk bandwidth). Moreover, Stratus implements a software manager to ease the research and aid the design of QoS-aware policies. The manager integrates three main functionalities: management and control of the execution of VMs and running applications, monitoring of hardware performance counters and system resource utilization, and partitioning of the main shared system resources by using technologies available in commercial processors.

Multi-Cloud Container Orchestration for High-Performance Real-Time Online Applications.

Sezar Jarrous-Holtrup, Sergei Gorlatch, Michael Dey and Folker Schamel.

We develop a novel multi-cloud container orchestration architecture for high-performance Real-Time Online Interactive Applications (ROIA), with use cases including product configurators, multiplayer online gaming, e-learning and -training. Running the core components of ROIA, e.g., real-time 3D rendering, on a multi-cloud enables access to high-performance resources and prevents proprietary ‘vendor lock-in’. Our container orchestration addresses several challenges: (1) satisfying the strict Quality of Service (QoS) requirements, (2) secure communication between cluster nodes from different clouds, (3) automatic scalability, and (4) resource usage optimization. We improve previous work by using session slots that set a limit on the number of concurrent user sessions for a service instance without loss of QoS. Our implementation provides a vendor-independent, OpenVPN-based interconnection between cloud nodes, both Linux and Windows, possibly located in different LANs of a multi-cloud. We experimentally evaluate our orchestration approach on a Kubernetes-based cluster with a prototype of an interactive car configurator.

FRIDAY, 3 MARCH

AWS ACADEMY 11:00 AM - 12:30 PM

SPECIAL SESSION: Compute Continuum (CC 2023)

CHAIR: Maria Fazio (University of Messina)

- Francesco Martella, Valeria Lukaj, Maria Fazio, Antonio Celesti, and Massimo Villari: **On-Demand and Automatic Deployment of Microservice at the Edge Based on NGSI-LD**

- Gabriele Russo Russo, Valeria Cardellini, and Francesco Lo Presti: **Serverless Functions in the Cloud-Edge Continuum: Challenges and Opportunities**

- Yasir Arfat, Gianluca Mittone, Iacopo Colonnelli, Fabrizio D'Ascenzo, Roberto Esposito, and Marco Aldinucci: **Pooling critical datasets with Federated Learning**

- Loris Belcastro, Fabrizio Marozzo, Alessio Orsino, Domenico Talia, and Paolo Trunfio: **Using the Compute Continuum for Data Analysis: Edge-cloud Integration for Urban Mobility**

On-Demand and Automatic Deployment of Microservice at the Edge Based on NGSI-LD.

Francesco Martella, Valeria Lukaj, Maria Fazio, Antonio Celesti and Massimo Villari.

This paper presents the innovative concept of ”virtual sensors” operating in smart environments, which are abstracted components able to map different behaviours on the same IoT-based infrastructures according to the needs of the high-level application. To realize ”virtual sensors”, it is necessary to codify user requests in an automation process for the deployment at the Edge of the microservices (MSs) that satisfy such requests. We present a solution that implements all the necessary functionalities to bind the user application with the Edge device in charge to execute the "virtual sensors". In particular, the solution we proposed is based on the FIWARE NGSI-LD information model, which helps us to standardize the communication among the different entities involved in the process. The paper describes the reference architecture we designed, provides the implementation details of our first prototype and reports the results of our evaluation tests.

Serverless Functions in the Cloud-Edge Continuum: Challenges and Opportunities.

Gabriele Russo Russo, Valeria Cardellini and Francesco Lo Presti.

The Function-as-a-Service (FaaS) paradigm is increasingly adopted for the development of Cloud-native applications, which especially benefit from the seamless scalability and attractive pricing models of serverless deployments. With the continuous emergence of latency-sensitive applications and services, including Internet-of-Things and augmented reality, it is now natural to wonder whether and how the FaaS paradigm can be efficiently exploited in the Cloud-Edge Continuum, where serverless functions may benefit from reduced network delay between their invoking users and the FaaS platform. In this paper, we illustrate the key challenges that must be faced to effectively deploy serverless functions in the Cloud-Edge Continuum and review recent contributions proposed by the research community towards overcoming those challenges. We also discuss the key issues that currently remain unsolved and highlight a few research opportunities for better support of FaaS in the Compute Continuum.

Pooling critical datasets with Federated Learning.

Yasir Arfat, Gianluca Mittone, Iacopo Colonnelli, Fabrizio D'Ascenzo, Roberto Esposito and Marco Aldinucci.

Federated Learning (FL) is becoming popular in different industrial sectors where data access is critical for security, privacy and the economic value of data itself. Unlike traditional machine learning, where all the data must be globally gathered for analysis, FL makes it possible to extract knowledge from data distributed across different organizations that can be coupled with different Machine Learning paradigms. In this work, we replicate, using Federated Learning, the analysis of a pooled dataset (with AdaBoost) that has been used to define the PRAISE score, which is today among the most accurate scores to evaluate the risk of a second acute myocardial infarction. We show that thanks to the extended-OpenFL framework, which implements AdaBoost.F, we can train a federated PRAISE model that exhibits comparable accuracy and recall as the centralised model. We use the federated PRAISE score as a running example to critically review the FL methods and tools and their run-time performance on distributed infrastructures.

Using the Compute Continuum for Data Analysis: Edge-cloud Integration for Urban Mobility.

Loris Belcastro, Fabrizio Marozzo, Alessio Orsino, Domenico Talia and Paolo Trunfio.

More and more in recent years, IT companies have adopted edge-cloud continuum solutions to efficiently perform analysis tasks on data generated by IoT devices. As an example, in the context of urban mobility, the use of edge solutions can be extremely effective in managing tasks that require real-time analysis and low response times, such as driver assistance, collision avoidance and traffic sign recognition. On the other hand, the integration with cloud systems can be convenient for tasks that require a lot of computing resources for accessing and analyzing big data collections, such as route calculations and targeted advertising. Designing and testing such hybrid edge-cloud architectures are still open issues due to their novelty, large scale, heterogeneity, and complexity. In this paper, we analyze how the compute continuum can be exploited for efficiently managing urban mobility tasks. In particular, we focus on a case study related to taxi fleets that need to find locations where they are more likely to find new passengers. Through a simulation-based approach, we demonstrate that these solutions turn out to be effective for this class of problems, especially as the number of connected vehicles increases.

FRIDAY, 3 MARCH

TEATRO 3:00 PM - 4:30 PM

MAIN TRACK

CHAIR: David E. Singh (Universidad Carlos III de Madrid)

- Julen Galarza, Javier Navaridas, Jose A. Pascual, Juan L Muñoz, Ibon Bustinduy, and Txomin Romero: Parallelizing Multipacting **Simulation for the Design of Particle Accelerator Components**

- Ryota Yasudo: **Bandit-based Variable Fixing for Binary Optimization on GPU Parallel Computing**

- Thomas Jakobs, Sebastian Kratzsch, and Gudula Ruenger: **Analyzing Data Reordering of a combined MPI and AVX execution of a Jacobi Method**

- Adriano Marques Garcia, Dalvan Griebler, Claudio Schepke, André Sacilotto Santos, Jose Daniel Garcia, Javier Fernandez Muñoz, and Luiz Gustavo Fernandes: **A Latency, Throughput, and Programmability Perspective of GrPPI for Streaming on Multi-cores**

Parallelizing Multipacting Simulation for the Design of Particle Accelerator Components.

Julen Galarza, Javier Navaridas, Jose A. Pascual, Juan L Muñoz, Ibon Bustinduy and Txomin Romero.

Particle trajectory and collision simulation is a critical step of the design and construction of novel particle accelerator components. However it requires a huge computational effort which can slow down the design process. We started from a sequential simulation program which is used to study an event called “Multipacting”. Our work explains the physical problem that is simulated and the implications it can have on the behavior of the components. Then we analyze the original program’s operation to fnd the best options for parallelization. We frst developed a parallel version of the Multipacting simulation and were able to accelerate the execution up to 35× with 48 or 56 cores. In the best cases, parallelization efciency was maintained up to 16 cores (95%) and the speed-up plateaus at around 40 to 48 cores. When this frst parallelization effort was tried for multi-power simulations, we found that parallelism was severely limited with a maximum of 20× speed-up. For this reason, we introduced a new method to improve the parallelization efciency for this second use case. This method uses a shared processor pool for all simulations of electrons (OnePool). OnePool improved scalability by pushing the speed-up to over 32×.

Bandit-based Variable Fixing for Binary Optimization on GPU Parallel Computing.

Ryota Yasudo.

The combination of different parallel programming environments can be used to exploit all heterogeneous levels of parallel hardware, which might lead to an optimization of application programs. An exemplary combination is the use of the Message Passing Interface (MPI) together with vectorization based on the Advanced vector extensions (AVX), which is investigated in this article. A special emphasis lies on MPI data orderings and their influence on AVX vectorization strategies. The Jacobi method is used as case study for which several parallel program version have been implemented and analyzed.

Analyzing Data Reordering of a combined MPI and AVX execution of a Jacobi Method.

Thomas Jakobs, Sebastian Kratzsch, and Gudula Ruenger.

The combination of different parallel programming environments can be used to exploit all heterogeneous levels of parallel hardware, which might lead to an optimization of application programs. An exemplary combination is the use of the Message Passing Interface (MPI) together with vectorization based on the Advanced vector extensions (AVX), which is investigated in this article. A special emphasis lies on MPI data orderings and their influence on AVX vectorization strategies. The Jacobi method is used as case study for which several parallel program version have been implemented and analyzed.

Summarizing task-based applications behavior over many nodes through progression clustering.

Lucas Leandro Nesi, Vinícius Garcia Pinto, Lucas Mello Schnorr and Arnaud Legrand.

Visualization strategies are a valuable tool in the performance evaluation of HPC applications. Although the traditional Gantt charts are a widespread and enlightening strategy, it presents scalability problems and may misguide the analysis by focusing on resource utilization alone. This paper proposes an overview strategy to indicate nodes of interest for further investigation with classical visualizations like Gantt charts. For this, it uses a progression metric that captures work done per node inferred from the task-based structure, a time-step clustering of those metrics to decrease redundant information, and a more scalable visualization technique. We demonstrate with six scenarios and two applications that such a strategy can indicate problematic nodes more straightforwardly while using the same visualization space. Also, we provide examples where it correctly captures application work progression, showing application problems earlier and as an easy way to compare nodes. At the same time that traditional methods are misleading.