EURO-PAR
CONFERENCE 2020

Program

Show abstracts

MONDAY 24.08.2020

9:00 - 10:30

ParaMo (International Workshop on Parallel Programming) (A)

Performance Evaluation of Java/PCJ Implementation of Parallel Algorithms on the Cloud [Slack channel]
Marek Nowicki, Łukasz Górski and Piotr Bała [video]

Cloud resources are more often used for large scale computing and data processing. However, the usage of the cloud is different than traditional High-Performance Computing (HPC) systems and both algorithms and codes have to be adjusted. This work is often time-consuming and performance is not guaranteed. To address this problem we have developed the PCJ library (Parallel Computing in Java), a novel tool for scalable high-performance computing and big data processing in Java. In this paper, we present a performance evaluation of parallel applications implemented in Java using the PCJ library. The performance evaluation is based on the examples of highly scalable applications that run on the traditional HPC system and Amazon AWS Cloud. For the cloud, we have used Intel x86 and ARM processors running Java codes without changing any line of the program code and without the need for time-consuming recompilation. Presented applications have been parallelized using the PGAS programming model and its realization in the PCJ library. Our results prove that the PCJ library, due to its performance and ability to create simple portable code, has great promise to be successful for the parallelization of various applications and run them on the cloud with a similar performance as for HPC systems.

Experiments using a Software-Distributed Shared Memory, MPI and 0MQ over Heterogeneous Computing Resources
Loïc Cudennec and Kods Trabelsi [video]

Distributed heterogeneous computing systems escalate the problem of choosing the appropriate programming model. Programming models such as message passing are efficient but require low-level management of communications. Higher level of programming such as shared memory are convenient for the application design but they usually have performance issues. With the recent development of distributed heterogeneous systems and new protocols to access remote memories, there is an opportunity for distributed shared memory systems to offer a satisfying level of abstraction while not giving up on performance. In this paper a video processing application is written using MPI, 0MQ and an in-house software-distributed shared memory (S-DSM) backend and deployed over a set of heterogeneous computing boards. Results show that 0MQ implementation is the most efficient but at the price of writing the application with the targeted platform in mind. The S-DSM implementation runs up to 2 times faster than the pure OpenMPI implementation and competes with 0MQ when the data granularity is small.

Improving Existing WMS for reduced Makespan of Workflows with Lambda
Ali Al-Haboobi [video]

Scientific workflows are an increasingly important area in complex scientific applications. Recently, Function as a Service (FaaS) has emerged as a powerful platform for processing background tasks such as web applications. FaaS can play an important role in processing scientific workflows such as AWS Lambda and Google Cloud Functions. A number of works have demonstrated their ability to process small- and large-scale workflows. However, some issues were identified when workflows executed on cloud functions due to their limits as well as they are stateless. For example, more data dependencies transfer occur during the execution between object storage and the FaaS invocation environment, leading to more communication costs. DEWE v3 is one of the Workflow Management Systems (WMSs) that provides three different execution modes: (traditional cluster, cloud functions, and hybrid mode). In this paper, we have modified the job dispatch algorithm of DEWE v3 on a function environment to reduce data dependencies transfer. The modified algorithm schedules jobs with precedence constraints to be executed in a single function invocation. Therefore, successor jobs can utilise output files generated from their predecessor job in the same invocation. This will speed up the makespan of workflow execution. We have tested the improved scheduling algorithm and the original algorithm with small- and large-scale Montage workflows. The experimental results show that the improved algorithm can reduce the overall makespan in contrast to DEWE v3 in most cases.

FPDAPP (3rd International Workshop on Future Perspective of Decentralized Applications) (B)

Welcome from the chairs [Slack channel]

Trustless, Censorship-Resilient and Scalable Votings in the Permission-based Blockchain Model
Marco Lewandowsky and Sebastian Gajek [video]

Voting systems are the tool of choice when it comes to settle an agreement of different opinions. We propose a solution for a trustless, censorship-resilient and scalable electronic voting platform. By leveraging the blockchain together with the functional encryption paradigm, we fully decentralize the system and reduce the risks that a voting provider, like a corrupt government, does censor or manipulate the outcome.

A digital voting system for the 21th century
Davide Casaleggio, Vincenzo Di Nicola, Michele Marchesi, Sebastiano Missineo and Roberto Tonelli [video]

We present Terminus, a voting system based on blockchain technology, specially addressed to non-binding and political polls. Terminus relies on technology solutions pioneered by Monero, a privacy-focused Blockchain, and operational procedures: this guarantees full anonymity of the vote and addresses several concerns of digital voting systems. Terminus was tested at an event of an Italian political movement, and will be used to carry out polls to drive some of the political decisions of this movement. We also introduce an evaluation framework for DLT voting systems, and use it to compare existing systems.

P2T: pay to transport
Fadi Barbara and Claudio Schifanella [video]

Abstract. We present Pay To Transport (P2T), a protocol that lets customers buy an item remotely in an atomic, privacy preserving and trustless manner. P2T needs only basic features of a blockchain scripting language and does not need any tracking systems, arbitrator or deposit to preserve its security properties. For this reason the protocol can be implemented on any permissionless blockchain, regardless of its scripting language, without additional trust. Merchants' and transporters' addresses are public, but in P2T the parties never pay those addresses directly. Therefore P2T maintains the privacy of customers, merchant and transporters.

10:30 - 11:00

Break

11:00 - 12:30

ParaMo (International Workshop on Parallel Programming) (A)

Parallelizing Automatic Temporal Cognitive Tool for Large-scale Online Learning Analytics [Slack channel]
Tianrui Jiang, Wenjun Wu and Yanjun Pu [video]

With the advent of Massive Online Open Courses (MOOCs), the data scale of student learning behavior has significantly increased. In order to analyze these datasets efficiently and present on-the-fly intelligent tutoring to online learners, it is necessary to improve existing learning analytics tools in a parallel and automatic way. We introduce Automatic Temporal Cognitive (ATC) model to describe temporal progress of online learners and evaluate their mastery of course knowledge. As a complex dynamic Bayesian network model, it often causes high computational overhead of training the ATC model via Probabilistic Programming tools. The time-consuming Monte Carlo sampling adopted by the mainstream implementations renders parameter fitting for the model a slow execution process. To address the issue, this paper proposes to transform the ATC model into the form of nonlinear Kalman filter and presents a new parallel ATC tool based on the Spark framework with the method of Unscented Kalman Filter (UKF). This tool improves the ATC model by using a parallel UKF method with the capability of automatically estimating the parameters in the whole sequential process. Experimental results demonstrate that this tool can achieve the fast execution speed and greatly improve the robustness of training parameters on different sizes of real educational data sets.

On the Provenance Extraction Techniques from Large Scale Log Files: A Case Study for the Numerical Weather Prediction Models
Alper Tufek and Mehmet Aktaş [video]

Day by day, severe meteorological events increasingly highlight the importance of fast and accurate weather forecasting. There are various Numerical Weather Prediction (NWP) models worldwide that are run on either a local or a global scale to predict the future weather. NWP models typically take hours to finish a complete run, however, depending on the input parameters and the size of the forecast domain. Provenance information is of central importance for detecting unexpected events that may develop during the course of model execution, and also for taking necessary action as early as possible. In addition, the need to share scientific data and results between researchers or scientists also highlights the importance of data quality and reliability. In this study, we develop a framework for tracking The Weather Research and Forecasting (WRF) model and for generating, storing and analyzing provenance data. We develop a machine-learning-based log parser in order to enable the proposed system to be dynamic and adaptive so that it can adapt to different data and rules. The proposed system enables easy management and understanding of numerical weather forecast workflows by providing provenance graphs. By analyzing these graphs, potential faulty situations that may occur during the execution of WRF can be traced to their root causes. Our proposed system has been evaluated and has been shown to perform well even in a high-frequency provenance information flow.

CHAOSS (Challenges and Opportunities of HPC Storage Systems) (A)

HugeMap: Optimizing Memory-mapped I/O with Huge Pages for Fast Storage [Slack channel]
Ioannis Malliotakis, Anastasios Papagiannis, Manolis Marazakis and Angelos Bilas [video]

Memory-mapped I/O (mmio) is emerging as a viable alternative for accessing directly-attached fast storage devices compared to explicit I/O with system calls. Mmio removes the need for costly lookups in the DRAM I/O cache for cache hits, as they are handled in hardware via the virtual memory mechanism. In this work we present HugeMap, a custom mmio path in the Linux kernel that uses huge pages for file-backed mappings to accelerate applications with sequential I/O access patterns or large I/O operations.HugeMap uses huge pages to reduce CPU processing in the kernel I/O path compared to regular mmap. We explore the benefits and trade-offs of huge pages in HugeMap using microbenchmarks, IOR, and an in-house persistent key-value store designed for mmio. Our experiments show up to 3.7× higher throughput and up to 4.76× lower system time, compared to regular page configurations.

FPDAPP (3rd International Workshop on Future Perspective of Decentralized Applications) (B)

Blockchain Utility in Use Cases: Observations, Red Flags, and Requirements [Slack channel]
Tommy Koens and Erik Poll [video]

On a global scale blockchain is persistently used in thousands of use cases by corporates, governments, and academics.However, there is a lack of systematic evaluation of these use cases and the utility of blockchain In this work we systematically evaluate fifteen use cases that use blockchain. Based on our evaluation we observe six recurring problems in these use cases. These problems either relate to the utility of blockchain in the use case, or to how well-documented a use case description is. We point out four red flags that, whenever they occur in a use case description, signal that blockchain may be a sub-optimal solution for that use case. Notably, one of these red flags indicates that there are no clear requirements in the use case descriptions that warrant the use of blockchain. We address this by proposing a set of requirement templates for any use case that includes a transaction system.

Next Generation Blockchain-Based Financial Services
Roberto Moncada, Enrico Ferro, Alfredo Favenza and Pierluigi Freni [video]

This paper explores the transition towards a paradigm in which centralization and decentralization systems coexist in the provision of financial services. The application of blockchain technology to the financial industry is giving birth to Decentralised Finance (DeFi). In the first place, we discuss the main implications determined by a shift in the balance between a centralized and a decentralized management of financial services, outlining the main novelties introduced by DeFi. Subsequently, we provide an introductory investigation of blockchain technology and the consequent emergence of tokenomics as a new strand of research. The study proceeds with the analysis of eight blockchain infrastructures to present the technical scenarios within which the DeFi ecosystem has proliferated in the last years. This exploratory study allows to observe the predominance of the Ethereum blockchain and the emergence of new and more efficient infrastructures within the DeFi environment, discussing the main differences between the financial services provided on-chain and off-chain.

Ants-Review: A Protocol For Open Anonymous Peer-Reviews
Bianca Trovò and Nazzareno Massari [video]

Peer-review is a necessary and essential quality control step for scientific publications. However, the process, which is very costly in terms of time investment, not only is not remunerated but it's also not recognized by the academic community as a relevant scientific output for a researcher. Therefore, scientific dissemination is affected. Here, to solve this issue, we propose a blockchain-based incentive protocol that rewards scientists also for their contributions to other scientists' work and that builds up a reputational system. We designed a Protocol of smart contracts called Ants-Review that allows any author to issue a call for peer-reviewing their scientific publication. If requirements are met, peer-reviews will be accepted and payed by the Issuer. To promote ethical behaviour the system will implement an incentive mechanism on AntsReview.

Closing

12:30 - 13:30

Break

13:30 - 15:00

PDCLifeS (Workshop on Parallel and Distributed Computing for Life Science: Algorithms, Methodologies and Tools) (A)

Keynote: Protein sequence-structure-dynamics-function relationships: efficient tools for mining experimental and simulated data [Slack channel]
Elodie Laine, Laboratory of Computational and Quantitative Biology (LCQB), Sorbonne Université, CNRS, IBPS

Analysis of Genome Architecture Mapping data with a Machine Learning and Polymer-Physics-based tool
Luca Fiorillo, Mattia Conte, Andrea Esposito, Francesco Musella, Francesco Flora, Andrea Maria Chiariello and Simona Bianco [video]

Understanding the mechanisms driving the folding of chromosomes in nuclei is a major goal of modern Molecular Biology. Recent technological ad-vances in microscopy (FISH, STORM) and sequencing approaches (Hi-C, GAM, SPRITE) enabled to collect quantitative data about chromatin 3D architecture, revealing a non-random and highly specific organization. To transform such tre-mendous amount of data into valuable insights on genome folding, heavy com-putational analyses are required. Here, we study the performances of PRISMR, a computational tool based on Machine Learning strategies and Polymer Physics principles, to explore genome 3D structure from Genome Architecture Mapping (GAM) data. Using such data, we show that PRISMR can successfully recon-struct the 3D structure of real genomic regions at various length scales, from mega-base sized loci to whole chromosomes. Importantly, the inferred structures are validated against independent Hi-C data. Finally, we show how PRISMR can be effectively employed to explore differences between experimental methods.

15:00 - 15:30

Break

15:30 - 16:00

PDCLifeS (Workshop on Parallel and Distributed Computing for Life Science: Algorithms, Methodologies and Tools) (A)

A New Parallel Methodology for the Network Analysis of COVID-19 data. [Slack channel]
Giuseppe Agapito, Marianna Milano and Mario Cannataro [video]

Coronavirus disease (COVID-19) outbreak started at Wuhan, China, and it has rapidly spread across the world. In this article, we present a new methodology for network-based analysis of Italian COVID- 19 data. The methodology includes the following steps: (i) a parallel methodology to build similarity matrices that represent similar or dissimilar regions with respect to data; (ii) the mapping of similarity matrices into networks where nodes represent Italian regions, and edges represent similarity relationships; (iii) the discovering communities of regions that show similar behaviour. The methodology is general and can be applied to world-wide data about COVID-19. Experiments was performed on real datasets about Italian regions, and they although the limited size of the Italian COVID-19 dataset, a quite linear speed-up was obtained up to six cores.

TUESDAY 25.08.2020

9:00 - 10:30

HeteroPar (18th International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms) (A)

Welcome from Steering Committee and Program Chair [Slack channel]
Roman Wyrzykowski

Session 1: Algorithms and languages for heterogeneous computing

Chairs: Alexey Lastovetsky

Scientific keynote 1: Opportunities for Approximate vs Transprecision Computing in Sparse Linear Solvers for GPUs
Prof. Enrique S. Quintana-Orti, Technical University of Valencia

HighPerMeshes -- A Domain-Specific Language for Numerical Algorithms on Unstructured Grids
Samer Alhaddad, Jens Förstner, Stefan Groth, Daniel Grünewald, Yevgen Grynko, Frank Hannig, Tobias Kenter, Franz-Josef Pfreundt, Christian Plessl, Merlind Schotte, Thomas Steinke, Jürgen Teich, Martin Weiser and Florian Wende

Solving partial differential equations on unstructured grids is a cornerstone of engineering and scientific computing. Nowadays, heterogeneous parallel platforms with CPUs, GPUs, and FPGAs enable energy-efficient and computationally demanding simulations. We developed the HighPerMeshes C++-embedded Domain-Specific Language (DSL) for bridging the abstraction gap between the mathematical and algorithmic formulation of mesh-based algorithms for PDE problems on the one hand and an increasing number of heterogeneous platforms with their different parallel programming and runtime models on the other hand. Thus, the HighPerMeshes DSL aims at higher productivity in the code development process for multiple target platforms. We introduce the concepts as well as the basic structure of the HighPerMeshes DSL, and demonstrate its usage with three examples, a Poisson and monodomain problem, respectively, solved by the continuous finite element method, and the discontinuous Galerkin method for Maxwell's equation. The mapping of the abstract algorithmic description onto parallel hardware, including distributed memory compute clusters is presented. Finally, the achievable performance and scalability are demonstrated for a typical example problem on a multi-core CPU cluster.

10:30 - 11:00

Break

11:00 - 12:30

HeteroPar (18th International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms) (A)

Session 2: Software engineering for heterogeneous parallel systems [Slack channel]

Chairs: Tal El-Nun

An Open-Source Virtualization Layer for CUDA Applications
Niklas Eiling, Stefan Lankes and Antonello Monti

GPUs have achieved widespread adoption for High-Performance Computing and Cloud applications. However, the closed-source nature of CUDA has hindered the development of otherwise commonly used virtualization techniques. In this paper, we evaluate the easibility of building a GPU virtualization layer that isolates the GPU and CPU parts of CUDA applications to achieve better control of the interactions between applications and the CUDA libraries. We present our open-source tool that transparently intercepts CUDA library calls and executes them in a separate process using remote procedure calls. This allows the execution of CUDA applications on machines without a GPU and provides a basis for the development of tools that require fine-grained control of the GPU resources, such as checkpoint/restore and job schedulers.

Implementation and evaluation of CUDA-Unified memory in Numba
Lena Oden

Python as a programming language is increasingly gaining importance, especially in data science, scientific, and parallel programming. With the Numba-CUDA, it is even possible to program GPUs with Python using a CUDA like programming style. However, NUMBA is missing CUDA-unified memory, which can help to simplify programming even more and allows dynamic work distribution between GPUs and CPUs. In this work, we implement and evaluate the support for unified memory in Numba. As expected, the performance of unified memory is worse than using explicit data transfers but can outperform the performance of the implicit methods provided by Numba. Additionally, using unified memory can help to reduce the Python interpreter overhead and therefore help to improve the performance of small Problem sizes. The use of system-wide atomic can help to improve the work distribution between GPU and CPU, but when using more CPU threads the performance suffers under the Python global interpreter lock.

Preparing Ginkgo for AMD GPUs -- A Testimonial on Porting CUDA Code to HIP
Yuhsiang M. Tsai, Terry Cojean, Tobias Ribizel and Hartwig Anzt

With AMD reinforcing their ambition in the scientific high performance computing ecosystem, we extend the hardware scope of the GINKGO linear algebra package to feature a HIP backend for AMD GPUs. In this paper, we report and discuss the porting effort from CUDA, the extension of the HIP framework to add missing features such as cooperative groups, the performance price of compiling HIP code for AMD architectures, and the design of a library providing native backends for NVIDIA and AMD GPUs while minimizing code duplication by using a shared code base.

12:30 - 13:15 (HeteroPar)
12:30 - 13:30 (Resilience)

Break

13:15 - 15:00 (HeteroPar)
13:30 - 15:00 (Resilience)

HeteroPar (18th International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms) (A)

Session 3: Heterogeneous computing and machine learning/AI algorithms [Slack channel]

Chairs: Enrique S. Quintana-Orti

Scientific keynote 2: Stateful Dataflow Multigraphs: A Data-Centric Approach for Performance Portability on Heterogeneous Architectures
Tal El-Nun, ETHZ Zurich

Management of heterogeneous cloud resources with use of the Proximal Policy Optimization
Paweł Koperek, Wlodzimierz Funika and Jacek Kitowski

Reinforcement learning has been recently a very active eld of research. Thanks to combining it with Deep Learning, many newly de- signed algorithms improve the state of the art. In this paper we present the results of our attempt to use the recent advancements in Reinforce- ment Learning to automate the management of heterogeneous resources in an environment which hosts a compute-intensive evolutionary pro- cess. We describe the architecture of our system and present evaluation results. The experiments include autonomous management of a sample workload and a comparison of its performance to the traditional au- tomatic management approach. We also provide the details of training of the management policy using the Proximal Policy Optimization algo- rithm. Finally, we discuss the feasibility to extend the presented approach to other scenarios.

An Edge Attribute-wise Partitioning and Distributed Processing of R-GCN using GPUs
Tokio Kibata, Mineto Tsukada and Hiroki Matsutani

R-GCN (Relational Graph Convolutional Network) is one of GNNs (Graph Neural Networks). The model tries predicting latent information by considering directions and types of edges in graph-structured data, such as knowledge bases. The model builds weight matrices to each edge attribute. Thus, the size of the neural network increases linearly with the number of edge types. Although GPUs can be used for accelerating the R-GCN processing, there is a possibility that the size of weight matrices exceeds GPU device memory. To address this issue, in this paper, an edge attribute-wise partitioning is proposed for R-GCN. The proposed partitioning divides the model and graph data so that RGCN can be accelerated by using multiple GPUs. Also, the proposed approach can be applied to sequential execution on a single GPU. Both the cases can accelerate the R-GCN processing with large graph data, where the original model cannot be fit into a device memory of a single GPU without partitioning. Experimental results demonstrate that our partitioning method accelerates R-GCN by up to 3.28 times using four GPUs compared to CPU execution for a dataset with more than 1.6 million nodes and 5 million edges. Also, the proposed approach can accelerate the execution even with a single GPU by 1.55 times compared to the CPU execution for a dataset with 0.8 million nodes and 2 million edges.

Resilience (Workshop on Resiliency in High Performance Computing with Clouds, Grids, and Clusters) (B)

Opening: Resilience Workshop Organizers [Slack channel]

Keynote: Towards Resilient EU HPC Systems: A Blueprint
Petar Radojkovic, Barcelona Supercomputing Center. [video]

Predicting Hard Disk Failures in Data Centers using Temporal Convolutional Neural Networks
Alessio Burrello, Daniele Jahier Pagliari, Andrea Bartolini, Luca Benini, Enrico Macii and Massimo Poncino [video]

In modern data centers, storage system failures are major contributors to downtimes and maintenance costs. Predicting these failures by collecting measurements from disks and analyzing them with machine learning techniques can effectively reduce their impact, enabling timely maintenance. While there is a vast literature on this subject, most approaches attempt to predict hard disk failures using either classic machine learning solutions, such as Random Forests (RFs) or deep Recurrent Neural Networks (RNNs). In this work, we address hard disk failure prediction using Temporal Convolutional Networks (TCNs), a novel type of deep neural network for time series analysis. Using a real-world dataset, we show that TCNs outperform both RFs and RNNs. Specifically, we can improve the Fault Detection Rate (FDR) of ≈ 7.5% (FDR = 89.1%) compared to the state-of-the-art, while simultaneously reducing the False Alarm Rate (FAR = 0.052%). Moreover, we explore the network architecture design space showing that TCNs are consistently superior to RNNs for a given model size and complexity and that even relatively small TCNs can reach satisfactory performance.

15:00 - 15:30

Break

15:30 - 17:00

HeteroPar (18th International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms) (A)

Session 4: Heterogeneous parallel computing [Slack channel]

Chairs: Leonel Sousa

High-Performance GPU and CPU Signal Processing for a Reverse-GPS Wildlife Tracking System
Yaniv Rubinpur and Sivan Toledo

We present robust high-performance implementations of signal-processing tasks performed by a high-throughput wildlife tracking system called ATLAS. The system tracks radio transmitters attached to wild animals by estimating the time of arrival of radio packets to multiple receivers (base stations). Time-of-arrival estimation of wideband radio signals is computationally expensive, especially in acquisition mode (when the time of transmission of not known, not even approximately). These computation are a bottleneck that limits the throughput of the system. The paper reports on two implementations of ATLAS's main signal-processing algorithms, one for CPUs and the other for GPUs, and carefully evaluates their performance. The evaluations indicates that the GPU implementation dramatically improves performance and power-performance relative to our baseline, a high-end desktop CPU typical of the computers in current base stations. The performance improvements by more than 50X on a high-end GPU and more than 4X with a GPU platform consumes almost 5 times less power than the CPU platform. Performance-per-Watt ratios also improve (by more than 16X), and so do the price-performance ratios.

Parallelization of the k-means algorithm in a spectral clustering chain on CPU-GPU platforms
Guanlin He, Stephane Vialle and Marc Baboulin

k-means is a standard algorithm for clustering data. It constitutes generally the final step in a more complex chain of high quality spectral clustering. However this chain suffers from lack of scalability when addressing large datasets. This can be overcome by applying also the k-means algorithm as a pre-processing task. We describe parallel optimization techniques for the k-means algorithm on CPU and GPU. Experimental results on synthetic dataset illustrate the numerical accuracy and performance of our implementations.

High Performance Portable Solver for Tridiagonal Toeplitz Systems of Linear Equations
Beata Dmitruk and Przemyslaw Stpiczynski

The aim of this paper is to show that recently developed divide and conquer parallel algorithm for solving tridiagonal Toeplitz systems of linear equations can be easily and efficiently implemented for a variety of modern multicore and manycore architectures. Our new portable implementation uses OpenACC to ensure that it can be executed on both CPU-based and GPU-accelerated parallel systems. We consider both column-wise and row-wise storage formats for two dimensional arrays and show how to efficiently convert between these two formats using cache memory, as well as discuss which format is more suitable for CPU-based or GPU-accelerated architectures. Numerical experiments performed on Intel CPUs and Nvidia Kepler, Turing, and Volta GPUs show that our new implementation achieves good performance on these architectures.

Resilience (Workshop on Resiliency in High Performance Computing with Clouds, Grids, and Clusters) (B)

A comparison of several fault-tolerance methods for the detection and correction of floating-point errors in matrix-matrix multiplication [Slack channel]
Valentin Le Fèvre, Thomas Herault, Julien Langou and Yves Robert [video]

This paper compares several fault-tolerance methods for error detection and correction A comparison of several fault-tolerance methods for error detection and correction of floating-point errors in matrix-matrix multiplication in matrix-matrix multiplication. These methods include replication, triplication, Algorithm-Based Fault Tolerance (ABFT) and residual checking (RC). Error correction for ABFT can be achieved either by solving a small-size linear system of equations, or by recomputing corrupted coefficients. We show that both approaches can be used for RC. We provide a synthetic presentation of all methods before discussing their pros and cons. We have implemented all these methods with calls to optimized BLAS routines, and we provide performance data for a wide range of failure rates and matrix sizes. In addition, with respect to the literature, this paper consider relatively high error rates.

On the Detection of Silent Data Corruptions in HPC Applications Using Redundant Multi-Threading
Diego Perez, Thomas Ropars and Esteban Meneses [video]

This paper studies the use of Redundant Multi-Threading (RMT) to detect Silent Data Corruptions in HPC applications. To understand if it can be a viable solution in an HPC context, we study two software optimizations to reduce RMT performance overhead by reducing the amount of data exchanged between the replicated threads. We conduct experiments with representative HPC workloads to measure the performance gained obtained through these optimizations, and the error detection coverage that they can achieve. In the best case, when running on a processor that features Simultaneous Multi-Threading, our results show that the overhead can be as low as 1.4X without reducing much the ability to detect data corruptions.

Q&A Discussion (led by Resilience Workshop Organizers) / Closing remarks