SPPEXA WORKSHOP (2ND EDITION) ON
PARALLEL PROGRAMMING MODELS - PRODUCTIVITY AND APPLICATIONS FOR EXASCALE AND BEYOND
UNIVERSITY OF VERSAILLES - OCTOBER 18, 2017

Program

9:30-9:40	Nahid EMAD	Welcome introduction
09:40-10:00	Benjamin UEKERMAN Technical University of Munich, Germany	***Overview of SPPEXA***
10:00-10:20	Taisuke BOKU University of Tsukuba, Japan	***Overview of MYX***
10:20-10:40	Georg HAGER Erlangen Regional Computing Center, Germany	***Overview of ESSEX II Project***
10:40-11:00	Jose GRACIA HLRS Stuttgart, Germany	***Overview of Smart-DASH: Distributed Data Structures and Parallel Algorithms***
11:00-11:30	Cafe break	
11:30-12:15	Vassil ALEXANDROV Barcelona Supercomputing Center, Spain	***Towards Advanced Hybrid Monte Carlo Methods for Linear Algebra for Extreme Scale Systems: Latest Advances and Results***
12:15-13:30	Lunch Break	
13:30-14:15	Jack DONGARRA University of Tennessee, Oak Ridge National Lab, USA & University of Manchester, UK	***An Overview of High-Performance and a Look at Energy Saving on the Intel Knights Landing for Linear Algebra Computations***
14:15-14:35	Matthias S. MUELLER RWTH Aachen University, Germany	***Correctness analysis for one-sided communication in MUST***
14:35-14:55	Tetsuya SAKURAI University of Tsukuba, Japan	***Development of a Scalable Parallel Eigensolver for Large-scale Simulations and Data Analysis***
14:55-15:25	Cafe Break	
15:25-15:45	Roger KOWALESKI Ludwig Maximilians Universität, Munich, Germany	***Incorporating Heterogeneous Memory Hierarchies in DASH***
15:45-16:05	Serge PETITON CNRS/CRIStAL/Maison de la Simulation	***A programming paradigm for extreme computational and data science***
16:05-16:25	Kengo NAKAJIMA University of Tokyo, Japan	***Preconditioned Iterative Solvers in ppOpen-HPC/pK-Open-HPC for ESSEX-II***
16:25-16:45	Joseph SCHUCHART HLRS Stuttgart, Germany	***Supporting Global Task Dependencies in DASH***
16:45-17:45		***Open Discussion***
17:45-18:00		***Closing remarks***

Li-Parad
Laboratoire d'informatique
parallélisme, réseaux, algorithmes distribués

UNIVERSITÉ DE VERSAILLES ST-QUENTIN-EN-YVELINES
université PARIS-SACLAY

MAISON DE LA SIMULATION

ABSTRACTS

Vassil ALEXANDROV, Barcelona Supercomputing Center, Spain
Towards Advanced Hybrid Monte Carlo Methods for Linear Algebra for Extreme Scale Systems: Latest Advances and Results

Taisuke BOKU, University of Tsukuba, Japan
Overview of MYX

MYX Project under SPPEXA program is a tri-nation collaborative research based on parallel program correctness check with static analysis and dynamic run-time checking based on MPI programming framework. Ths original technology is provided by MUST system developed in RWTH Aachen, and it is expanded for various parallel processing framework not limited to direct programming on MPI. In MYX Project, we focus on XcalableMP PGAS language and YML task scheduling system based on MPI. In this talk, the overview of MYX Project and several sub-issues to support it are introduced.

Jack DONGARRA, University of Tennessee, Oak Ridge National Lab (USA) & University of Manchester (UK)
An Overview of High-Performance and a Look at Energy Saving on the Intel Knights Landing for Linear Algebra Computations

In this talk we will look at the current state of high performance computing and look to the future toward exascale. In addition, we will examine some issues that can help in reducing the power consumption for linear algebra computations.

Jack Dongarra holds an appointment at the University of Tennessee, Oak Ridge National Laboratory, and the University of Manchester. He specializes in numerical algorithms in linear algebra, parallel computing, use of advanced-computer architectures, programming methodology, and tools for parallel computers. He was awarded the IEEE Sid Fernbach Award in 2004; in 2008 he was the recipient of the first IEEE Medal of Excellence in Scalable Computing; in 2010 he was the first recipient of the SIAM Special Interest Group on Supercomputing's award for Career Achievement; in 2011 he was the recipient of the IEEE Charles Babbage Award; and in 2013 he received the ACM/IEEE Ken Kennedy Award. He is a Fellow of the AAAS, ACM, IEEE, and SIAM and a foreign member of the Russian Academy of Science and a member of the US National Academy of Engineering.

Jose GRACIA, HLRS Stuttgart, Germany
DASH: Distributed Data Structures and Parallel Algorithms

This talk provides an overview on DASH (www.dash-project.org), a C++ template library that offers distributed data structures and parallel algorithms and provides a PGAS (partitioned global address space) abstraction. DASH data structures are modeled after the familiar C++ STL container classes but can utilize the memory of multiple compute nodes. DASH allows convenient global access to any data element using one-sided remote memory operations and also supports a strong notion of data locality and makes it easy to realize the owner-computes idiom. DASH also offers parallel variants of many STL algorithms that harness the computing power of multiple compute nodes.

Georg HAGER, Erlangen Regional Computing Center, Germany
Overview of the ESSEX-(II) project

ESSEX investigates programming concepts and numerical algorithms for scalable, efficient and robust iterative sparse matrix applications on exascale systems. Starting with successful blueprints and prototype solutions identified in ESSEX-I, the second phase project ESSEX-II aims at delivering a collection of broadly usable and scalable sparse eigenvalue solvers with high hardware efficiency for the computer architectures to come. The talk will highlight the state of the project and interesting results in the fields of sparse building blocks and fault tolerance.

Kengo NAKAJIMA, University of Tokyo, Japan
Preconditioned Iterative Solvers in ppOpen-HPC/pK-Open-HPC for ESSEX-II

"ppOpen-HPC" is an open source infrastructure for development and execution of optimized and reliable simulation code on post-peta-scale (pp) parallel computers based on many-core architectures, and it consists of various types of libraries, which cover general procedures for scientific computation."ppOpen-HPC" is part of a five-year project (FY.2011-2015) funded by JST-CREST. In 2016, the team of ppOpen-HPC joined ESSEX-II (Equipping Sparse Solvers for Exascale) project, which is funded by "Software for Exascale Computing" (SPPEXA) under Japan (JST)-Germany (DFG) collaboration untilFY.2018. In ESSEX-II, we develop pK-Open-HPC (extended version of ppOpen-HPC, framework for exa-feasible applications), and preconditioned iterative solvers for quantum sciences. In the presentation, we will show various achievements of ppOpen-HPC, ESSEX-II, and pK-OpenHPC, focusing on those in development of parallel preconditioned iterative solvers for very ill conditioned problems.

Roger KOWALEWSKI, Ludwig Maximilians Universität, Munich, Germany
Incorporating Heterogeneous Memory Hierarchies in DASH

Emerging technologies such as non-volatile and 3D-stacked memory significantly change the design of future high performance computing systems. This talk discusses how programming models can abstract the complexity of the resulting heterogeneous memory hierarchy, while still giving explicit control to domain experts. We propose the concept of memory spaces to model a heterogeneous memory hierarchy and integrate it into DASH. An experimental evaluation with a series of benchmarks, conducted on an Intel KNL platform, reveals that using this concept to support data placement strategies achieves significant speedup.

Matthias S. MUELLER, RWTH Aachen University, Germany
Correctness analysis for one-sided communication in MUST

One-sided communication is important for PGAS-style parallel programming approaches. A big challenge for correctness analysis for this communication pattern is how to exchange information for analysis in time. In this work we present an approach of analysis for MPI one-sided communication and how this is also applicable to XMP programs and how this fits into the MUST correctness-checking framework.

Serge PETITON, CNRS/CRIStAL/Maison de la Simulation, France
A programming paradigm for extreme computational and data science

Future supercomputers are expected to have highly hierarchical architectures with nodes composed by lot-of-core processors and accelerators. The different programming levels will generate new difficult algorithm issues. New intelligent applications would mix computational and data sciences. Moreover, methods have to be redesigned and new ones introduced or rehabilitated, in particular in terms of communication optimizations and data distribution. Then, new languages and frameworks should be defined and evaluated with respect to modern state-of-the-art of scientific methods. In this talk, we first present a solution based on graph of components programming experimented on several machines. Hence, we focus on YML with its high level language allowing to automate and delegate the managements of dependencies between loosely coupled clusters of processors to a specialized tool controlling the execution of the application using parallel tasks. We present some methods for scientific computation well-adapted for such programming paradigm and we discuss new researches mixing components and containers associated with data science computation.

Tetsuya SAKURAI, Yasunori FUTAMURA, Akira IMAKURA:
Development of a Scalable Parallel Eigensolver for Large-scale Simulations and Data Analysis

In this talk, we present the parallel eigensolver z-Pares, based on the Sakurai-Sugiura method (SSM) for large-scale linear/nonlinear eigenvalue problems. A hierarchical structure of the method enables efficient utilizationof massively parallel computing resources. We show parallel performances of the method on K-Computer and Oakforest-PACS.

Joseph SCHUCHART, HLRS Stuttgart, Germany
Supporting Global Task Dependencies in DASH

Task-based programming in a global address space exposes a higher degree of concurrency than the traditional two-level combination of message passing and loop-based work sharing. To provide task synchronization across process boundaries we propose to extend the concept of OpenMP task data dependencies to the PGAS programming model. We evaluate our approach using a structured benchmark application, showing a significant improvement over other PGAS-based variants and outperforming a traditional MPI-based implementation in terms of strong-scaling characteristics.