

# **COLOC - Workshop on Data Locality**

# Workshop on Data Locality (COLOC)

## Workshop Description

A well-known handicap for HPC applications running on modern highly parallelized and heterogeneous HPC platforms is that an increasing amount of time is spent in communication and data transfers; thus, it is necessary to design, implement and validate new approaches to optimize process placement and data locality management. COLOC is a forum for exposing contribution from HPC application developers interested in exploring new ways to optimize their code; HPC centers and clusters managers to enhance cluster usage and application efficiency; Academics and researchers in scientific computing.

The different areas or research interest include, but are not limited to:

- Modeling node topology
- Modeling network and communication
- Performance analysis of applications to understand affinity
- Affinity metrics
- Runtime support for extracting affinity from application
- Code analysis in order to understand communication pattern
- Algorithm to improve locality
- Language, abstraction and compiler support for data locality
- Data structure and library support to better manage memory access
- Runtime-system and dynamic locality management
- System-scale locality optimization
- Validating locality optimization at thread or process level
- Memory management
- Locality management in large-scale application

We have received 6 submissions and we have accepted 5. All of them are published in this proceedings. The workshop also included the invited talk *Why don't we have data close to the computation? Let's understand and optimize data locality problem* from Fabio Baruffa, Intel.

The workshop also featured the SPPEXA <http://www.sppexa.de> Poster Session on Data Locality. Karl F rlinger (LMU, M nchen) was the poster chair. 8 posters were accepted and presented by young researchers:

- Ari Rasch, WWU M nster: Utilizing Data Locality on Multi- and Many-Core Devices via Multi-Dimensional Homomorphisms
- Roger Kowalewski, LMU M nchen: Scalable Hybrid Sorting on Distributed Many-Core Architectures using PGAS
- Huihui Sun, WWU M nster: Improving Vectorization in the Presence of Aggregated Conditions
- Richard Schulze, WWU M nster: Exploiting Data Locality for High-Performance BLAS Routines on Multi- and Many-Cores

- Yusuke Tanimura, AIST Japan: Towards Faster and Secure Data Staging From/To Object Storage for AI and Big Data Analytics
- Jannis Klinkenberg, RWTH Aachen: Assessing Task-to-Data Affinity in the LLVM OpenMP Runtime
- Pascal Jungblut, LMU München: Increasing locality by interleaving host and on-device patterns
- Florian Schmaus, FAU Erlangen: Efficient Micro-Parallelism Using Work-Stealing with Affinity Hints

## Organization

### Program Chair

Emmanuel Jeannot      Inria, France

### Program Committee

George Bosilca	UTK, USA
Florina Ciorba	University of Basel, Switzerland
Matthias Diener	UIUC, USA
Anshu Dubey	Argonne Natl Lab, USA
Karl Furlinger	LMU, München, Germany
Brice Goglin	Inria, France
Aleksandar Ilic	INESC-ID/IST, Univ. de Lisboa, Portugal
Vitus Leung	Sandia National Laboratories, USA
Hatem Ltaief	KAUST, Saudi Arabia
Farouk Mansouri	DDN, France
Naoya Maruyama	LLNL, USA
Hartmut Mix	Technische Universität Dresden, Germany
Marc Perache	CEA, France
Eric Petit	Intel, France
Didem Unat	Koç University, Turkey