HCW 2022 Call for Papers

In conjunction with IPDPS 2022, May 30, 2022, Lyon, France
Sponsored by the IEEE Computer Society
through the Technical Committee on Parallel Processing (TCPP)

Most modern computing systems are **heterogeneous**, either for *organic reasons* because components grew independently, as it is the case in desktop grids, or *by design* to leverage the strength of specific hardware, as it is the case in accelerated systems. In any case, all computing systems have some form of ***hardware or software heterogeneity*** that must been managed, leveraged, understood, and exploited. The Heterogeneity in Computing Workshop (HCW) is a venue to *discuss and innovate in all theoretical and practical* aspects of heterogeneous computing: design, programmability, efficient utilization, algorithms, modeling, applications, etc. The 2021 HCW is the 30th annual gathering of this workshop.

TOPICS
Topics of interest include but are not limited to the following areas:

!!! SPECIAL TOPIC 1 !!! **Heterogenous Integration of Quantum Computing:** Future of computers will lead to a system that consists of both classical and Quantum computers for accelerated performance. Design, exploration, and analysis of architectures and software frameworks that enables and needs the heterogeneous integration of classical computing and Quantum computing. (e.g., heterogeneous quantum computers, error correction, heterogeneous applications that use both classical and quantum logic, benchmarks for heterogeneous quantum computers.)

!!! SPECIAL TOPIC 2 !!! **Heterogeneity and Interoperability in SW & Data systems**: Design, exploration, and analysis of architectures and software frameworks for interoperability in SW and Data systems. (e.g., semantic based frameworks, interoperability for heterogeneous IoT systems, model-driven frameworks.)

Heterogeneous multicore systems and architectures: Design, exploration, and experimental analysis of **heterogeneous computing systems** such as GPGPUs, heterogeneous systems-on-chip (SoC), accelerator systems (e.g., Intel Xeon Phi, AI chips such as Google's TPUs), FPGAs, big.LITTLE, and application-specific architectures.

Heterogeneous parallel and distributed systems: Design and analysis of computing grids, cloud systems, hybrid clusters, datacenters, geo-distributed computing systems, and supercomputers.

Deep-memory hierarchies: Design and analysis of memory hierarchies with SRAM, DRAM, Flash/SSD, and HDD technologies; NUMA architectures; cache coherence strategies; novel memory systems such as phase-change RAM, magnetic (e.g., STT) RAM, 3D Xpoint/crossbars, and memristors.

On-chip, off-chip and heterogeneous network architectures: Network-on-chip (NoC) architectures and protocols for heterogeneous multicore applications; energy, latency, reliability, and security optimizations for NoCs; off-chip (chip-to-chip) network architectures and optimizations; heterogeneous networks (combination of NoC and off-chip) design, evaluation, and optimizations; large scale parallel and distributed heterogeneous network design, evaluation, and optimizations.

Programming models and tools: Programming paradigms and tools for heterogeneous systems; middleware and runtime systems; performance-abstraction tradeoff; interoperability of heterogeneous software environments; workflows; dataflows.

Resource management and algorithms for heterogeneous systems: Parallel algorithms for solving problems on heterogeneous systems (e.g., multicores, hybrid clusters, grids or clouds); strategies for scheduling and allocation on heterogeneous 2D and 3D multicore architectures; static and dynamic scheduling and resource management for large-scale and parallel heterogeneous systems.

Modeling, characterization, and optimizations: Performance models and their use in the design of parallel and distributed algorithms for heterogeneous platforms, characterizations and optimizations for improving the time to solve a problem (e.g., throughput, latency, runtime), modeling and optimizing electric consumption (e.g., power, energy); modeling for failure management (e.g., fault tolerance, recovery, reliability); modeling for security in heterogeneous platforms.

Applications on heterogeneous systems: Case studies; confluence of Big Data systems and heterogeneous systems; data-intensive computing; deep learning; scientific computing.

IMPORTANT DATES

Paper submission: February 11, 2022 (final extension)
Author notification: ~March 1, 2022
Camera Ready: ~March 14, 2022

PAPER SUBMISSIONS

- Papers are to be submitted through https://ssl.linklings.net/conferences/ipdps/

- Submissions for the Special Topic Session on DSAs: Please add (Special Topic Submission) to your paper title during the submission process.

- Papers submitted to HCW 2021 should not have been previously published or be under review for a different workshop, conference, or journal.

- It is required that all accepted papers will be presented at the workshop by one of the authors.

WORKSHOP ORGANIZATION

General Chair: Ryan D. Friese, Pacific Northwest National Laboratory, USA

Program Chair: Jong-Kook Kim, Korea University, Korea

Steering Committee:
Behrooz Shirazi, Washington State University, USA **(Chair)**
H. J. Siegel, Colorado State University, USA **(Past Chair)**
John Antonio, University of Oklahoma, USA
David Bader, New Jersey Institute of Technology, USA
Anne Benoit, École Normale Supérieure de Lyon, France
Jack Dongarra, University of Tennessee, USA
Alexey Lastovetsky, University College Dublin, UK
Sudeep Pasricha, Colorado State University, USA
Viktor K. Prasanna, University of Southern California, USA
Yves Robert, École Normale Supérieure de Lyon, France
Erik Saule, University of North Carolina at Charlotte, USA
Uwe Schwiegelshohn, TU Dortmund University, Germany

Technical Program Committee:
Jong-Kook Kim, Korea University, Korea (**TPC Chair**)
Mohsen Amini, University of Louisiana Lafayette, USA
David Bader, New Jersey Institute of Technology, USA
Burcu Mutlu, Pacific Northwest National Laboratory, USA
Louis-Claude Canon, Université de Franche-Comté, France
Daniel Cordeiro, University of São Paulo, Brazil
Matthias Diener, University of Illinois at Urbana-Champaign, USA
Diana Göhringer, Technische Universität Dresden, Germany
Mahzabeen Islam, AMD, USA
Krishna Kavi, University of North Texas, USA
Georgios Keramidas, Aristotle University, Greece
Joongheon Kim, Korea University, Korea
Sung Il Kim, SK Hynix, Korea
Hyoukjun Kwon, Facebook, USA
Alexey Lastovetsky, University College Dublin, Ireland
Hatem Ltaief, KAUST, Saudi Arabia
Tirthak Patel, Northeastern University, USA
Dana Petcu, West University of Timisoara, Romania
Laercio Lima Pilla, CNRS, France
Sridhar Radhakrishnan, University of Oklahoma, USA
Srishti Srivastava, University of Southern Indiana, USA
Achim Streit, Karlsruhe Institute of Technology, Germany
Samuel Thibault, LaBRI, Université Bordeaux, France
Cheng Wang, Microsoft, USA

Questions may be sent to the program chair: Jong-Kook Kim at jongkook@korea.ac.kr