Ninth Annual Workshop on Accelerators and Hybrid Exascale Systems (AsHES)

Description

The development of ever larger and more energy-efficient computer systems in recent years has led to more and more systems with heterogeneous computing units (CPUs, GPUs or FPGAS) and systems with heterogeneous storage systems (High Memory Bandwidth). With the rise of persistent memory, attached to the PCIe bus or to the memory DIMMs, the border between storage and memory becomes more and more fluid. Other systems offer different types of compute nodes, so that a group of nodes build the accelerator (modular supercomputing). Hierarchical storage architectures, for example using burst buffers, try to overcome the IO problems. Programming such a system can be a real challenge along with locality, scheduling, load balancing, concurrency and so on.

This workshop focuses on understanding the implications of accelerators and heterogeneous designs on the hardware systems, porting applications, performing compiler optimizations, and developing programming environments for current and emerging systems. It seeks to ground accelerator research through studies of application kernels or whole applications on such systems, as well as tools and libraries and runtime systems that improve the performance and productivity of applications on these systems.

The goal of this workshop is to bring together researchers and practitioners who are involved in application studies for accelerators and other heterogeneous systems, to learn the opportunities and challenges in future design trends for HPC applications and systems.

DOI 10.1109/IPDPSW.2019.00086