

Panel

Accelerator Interconnects

Abstract

Nearly all the emerging, high-value applications, especially those residing on the edge, use accelerators rather than CPUs for their essential functions. Moreover, in a cluster of any size, including at the edge, multiple accelerators are employed to meet the dynamic demands of such applications. These accelerators can be of the same type, as a pool, or of different types. But as with the case of CPUs, the performance of the system is almost always limited by the I/O or interconnect. Therefore we want to explore how the interconnections of accelerators not only avoid being the bottleneck but more importantly make it possible for the accelerators to work in concert to enable new applications and revolutionize the architecture of computing. This panel will explore the most innovative approaches to accelerator interconnects to take fullest advantage of the capability of these amazing accelerators.

Moderator

Fabrizio Petrini, *Intel*

Panelists

Darius Bunandar, *Lightmatter*

Brad Burres, *Intel*

Larry Dennison, *NVIDIA*

Frank Helms, *AMD*

Torsten Hoefler, *ETH Zurich*

Bithika Khargharia, *Meta*

Tina Tsou, *Arm*