Critically Missing Pieces in Heterogeneous Accelerator Computing

Pavan Balaji, Argonne National Laboratory

Heterogeneous architectures play a massive role in architecting the largest systems in the world. However, much of the interest in these architectures is an artifact of the hype associated with them. For such architectures to truly be successful, it is important that we look beyond this hype and learn what these architectures provide and what is critically missing. This continuing BoF series brings together researchers working on aspects of accelerator architectures---including data management, resilience, programming, tools, benchmarking, and auto-tuning---to identify critical gaps in the accelerator ecosystem.

We had a successful BoF at SC13 this year. We had a number of representatives from universities, labs and industry give their perspective on different missing aspects in accelerator computing --- busting the hype building around accelerators and allowing attendees to reason about what their state is, and what needs to be improved for their broader adoption. Jean-Francois Mehaut from Universite Joseph Fourier initiated a discussion on non-GPU accelerators, such as the Monte Blanc project, and challenges involved in such ecosystems. Franck Cappello provided various insights into faults and error modeling issues on GPUs. Milind Chabbi focused on the current state of performance tools for accelerators and what is missing in the accelerator ecosystem to better support such tools. Vinod Tipparaju from AMD and Luiz DeRose from Cray gave indications on where the industry is moving on accelerators and where their primary focus is, on resolving some of the missing parts in accelerators.

There was plenty of time at the end for audience participation. There were many questions and comments from the audience, which were addressed by the speakers.