

Multi-GPU Computing and GPU MapReduce

John Owens

UC Davis

jowens@ece.ucdavis.edu

It's easy to build a multi-GPU application with simple, static communication patterns. It's more interesting to build a multi-GPU system that can provide more fully-featured communication primitives. What's most interesting is building a system that can do both. In our group we have developed an MPI-like abstraction for GPUs that gives both high performance and high flexibility, and atop that built a multi-GPU, out-of-core MapReduce framework. In my talk I will discuss this work together with the important challenges as we move forward toward widespread use of GPUs in large computing systems.