

Workshop: Accelerating CUDA C++ Applications with Multiple GPUs

Wednesday, 11 September 2024

Copy/Compute Overlap with CUDA Streams (– Learn the key concepts for effectively performing copy/compute overlap. – Explore robust indexing strategies for the flexible use of copy/compute overlap in applications. – Refactor the single-GPU CUDA C++ application to perform copy/compute overlap. – See copy/compute overlap in the Nsight Systems visual profiler timeline.) (13:45 - 15:15)