

High Performance Multilevel Preconditioners on GPUs

We discuss several issues to obtain high performance for solvers and, especially, preconditioners on Graphics Processor Units (GPUs). Fine grained parallelism is essential, and for that reason preconditioners that are similar to a matrix-vector product achieve high performance. To also maintain good convergence, we consider several multilevel strategies. We demonstrate results for several CFD applications.

This is joint work with Katarzyna Swirydowicz, Chris Roy, Amit Amritkar, and Danesh Tafti.