High Performance Incomplete Sparse Factorizations on GPUs

Tobias Ribizel[1]   Hartwig Anzt[2]

Incomplete sparse factorizations like ILU(0), together with Jacobi preconditioners, are the most commonly used type of general-purpose preconditioners in iterative linear solvers. They can provide good numerical properties, but their computation exposes only a limited amount of parallelism. Modern hardware architectures like GPUs thus require innovative approaches to computing these factorizations. We discuss some of these approaches, which combine asynchronous fixed-point iterations with commonly available sparse linear algebra kernels, and can outperform classical incomplete factorization algorihms in runtime or accuracy, sometimes even both.

[1]Karlsruhe Institute of Technology (KIT)
  `ribizel@kit.edu`

[2]Karlsruhe Institute of Technology (KIT)
  `hartwig.anzt@kit.edu`