

# Revisiting the idea of multiprecision block-Jacobi preconditioning - where do we stand 2 years later?

*Tuesday, 5 November 2019 16:20 (30 minutes)*

At the PACO workshop 2017, we presented the idea of decoupling the memory precision from the arithmetic precision, and storing a block-Jacobi preconditioner such that the precision format of each diagonal block is optimized to the numerical characteristics. The idea is to reduce the pressure on the memory bandwidth while preserving regularity of the preconditioner and the convergence of the top-level iterative solver. This is expected to render attractive resource savings as the performance and energy footprint of memory-bound applications strongly correlates to the data transfer volume. Two years later, we review the effectiveness of this idea by evaluating a sophisticated high performance implementation of the adaptive precision block-Jacobi preconditioner for GPU architectures that is implemented in the Ginkgo numerical linear algebra library. We present performance results revealing the attractiveness of the approach, with the runtime savings even exceeding the expectations. We also consider the use of non-standard precision formats, and the problem-specific trade-off between preconditioner application cost and preconditioner effectiveness.

**Primary authors:** ANZT, Hartwig (Karlsruhe Institute of Technology); COJEAN, Terry (Karlsruhe Institute of Technology); GRÜTZMACHER, Thomas (Karlsruhe Institute of Technology); QUINTANA-ORTI, Enrique (Polytechnic University of Valencia)

**Presenter:** ANZT, Hartwig (Karlsruhe Institute of Technology)

**Session Classification:** Day I