

Solving linear systems efficiently using Block Low-Rank compression in mixed precision

Matthieu GEREST

EDF R&D, LIP6-Sorbonne University, France

Abstract

Many applications involve data sparse matrices whose off-diagonal blocks have low numerical ranks. Block Low-Rank (BLR) compression allows us to exploit this property, by representing these blocks as low-ranks matrices, in order to reduce the costs for both storing the matrix and computing its LU factorization.

We show that such a BLR compression can be improved by using several floating-point formats simultaneously. In fact, most of the coefficients and operations may be switched to lower precision formats, such as single precision, while preserving high accuracy. We show how this mixed precision approach can reduce both the memory peak and the computation time of the sparse solver MUMPS.