Inverse Problems, Kronecker Products and Mixed Precision Computations

James Nagy

Abstract

The gaming industry, machine learning (ML), and artificial intelligence (AI) are areas that require substantial computational resources and/or require very fast computations, but do not always require high accuracy in certain computational problems. This has motivated GPU vendors, such as NVIDIA, Google and AMD to manufacture hardware that can perform computations using low precision 16-bit floating-point formats [4]. Two examples are bfloat16 and FP16. In comparison, IEEE single precision uses a 32-bit floating-point format, and double precision (e.g., the default in MATLAB) uses a 64-bit floating-point format. The use of 16-bit format can result in a $4 \times$ speedup compared to double precision, and certain hardware accelerators (called Tensor Cores) can further accelerate performance for operations such as matrix-vector multiplications [4].

The potential for much faster computations has fueled a growing interest in the last decade to use powerful GPU servers for scientific applications, and in particular to use mixed precision algorithms for problems that require high accuracy; that is, when possible, use low precision for speed, but mix in some high precision computations to improve accuracy. Recent previous work for solving general, well-conditioned linear systems, including iterative refinement [1, 2, 7], Cholesky factorization and least squares problems [1, 5], QR factorization [8], and GMRES [6].

Relatively little work has been done to exploit mixed precision computations for inverse problems, where the aim is to compute approximations of x from measured data, b, where

$$b = Ax + e \,. \tag{1}$$

A is assumed to be a large, severely ill-conditioned matrix, and e represents unknown noise and other data measurement errors. In some applications A is known to high accuracy, while in other applications it may be that only an approximation of A is given, or that $A \equiv A(y)$ is given in parametric form. Even in the case when A is known to high accuracy, due to the ill-posedness of the problem, and the presence of noise in the measured data, computing accurate approximations of x is a nontrivial task; special considerations, such as regularization approaches, need to be considered for these problems [3]. In this presentation we show how Kronecker product structure can be exploited and used in mixed precision algorithms for inverse problems.

References

- [1] E. Carson, N. J. Higham, and S. Pranesh. Three-precision GMRES-based iterative refinement for least squares problems. *SIAM Journal on Scientific Computing*, 42(6):A4063–A4083, 2020.
- [2] A. Haidar, H. Bayraktar, S. Tomov, J. Dongarra, and N. J. Higham. Mixed-precision iterative refinement using tensor cores on GPUs to accelerate solution of linear systems. *Proceedings of* the Royal Society A, 476(2243):20200110, 2020.
- [3] P. C. Hansen. Discrete Inverse Problems: Insight and Algorithms. SIAM, Philadelphia, PA, 2010.
- [4] N. Higham and T. Mary. Mixed precision algorithms in numerical linear algebra. Acta Numerica, 31:347–414, 2022.

- [5] N. J. Higham and S. Pranesh. Exploiting lower precision arithmetic in solving symmetric positive definite linear systems and least squares problems. SIAM Journal on Scientific Computing, 43(1):A258–A277, 2021.
- [6] J. A. Loe, C. A. Glusa, I. Yamazaki, E. G. Boman, and S. Rajamanickam. A study of mixed precision strategies for gmres on gpus. arXiv preprint arXiv:2105.07544, 2021.
- [7] E. Oktay and E. Carson. Multistage mixed precision iterative refinement. *Numerical Linear Algebra with Applications*, 29(4):e2434, 2022.
- [8] L. M. Yang, A. Fox, and G. Sanders. Rounding error analysis of mixed precision block Householder QR algorithms. SIAM J. Sci. Comput., 43:A1723–A1753, 2021.