

Interactive comment on “Technical Note: Improving the computational efficiency of sparse matrix multiplication in linear atmospheric inverse problems” by Vineet Yadav and Anna M. Michalak

Anonymous Referee #1

Received and published: 21 October 2016

The authors present a way to accelerate sparse matrix-matrix multiplication which can be a major bottleneck in the course of solving inverse problems, using a specific formulation. The paper is short, well written and self-contained.

The paper only provides matrix–matrix multiplication. As it stands, I find the novelty of this paper to be quite limited. There are other parallel matrix libraries that can perform in similar computational environments (although they would might require MPI) such as SCALAPACK, PETSc, FLAME, Elemental, etc. Why should one use this package as opposed to these other packages?

I think the benefit of this package maybe the examples briefly mentioned in Section 2. This paper may be considered for publication if additional software developed is

C1

performed to construct different types of covariance matrices (and their appropriate data structures) and handle different types of forward operators. For example, how can one pre-compute the sparsity pattern in inverse problems? More insight into these issues maybe more valuable.

General Comments: The numerical comparisons are limited in depth and breadth which weakens the contributions of this paper.

1. The algorithm proposed uses OpenMP but intel-mkl libraries are multi-threaded. Why is comparison between the two methods fair since they use different software paradigms? Please provide some justification for this.
2. Timing results are only provided for only one realization of m, n and one particular sparsity pattern. Are similar gains to be found for different sparsity patterns, or problem sizes? Since this is claimed to be a general purpose algorithm for inverse problems and matrix–matrix multiplication, either there should be a detailed analysis of computational cost, or extensive numerical testing for different choices of m, n and sparsity pattern.

Minor Comments:

1. In the introduction, matrix–matrix multiplication is mentioned as a fundamental operator in Bayesian inverse problems. This is only true for one particular computational formulation of the inverse problem; an alternative approach is to solve the MAP estimate using iterative methods, for which matrix–matrix multiplication may not be necessary nor desired. Please change the introduction to reflect this.
2. It would be useful if you could provide the inputs to the routines in Figure 1.

To conclude, this paper is not convincing either in its novelty or its execution. However, there may be some benefit to users if this package is more user friendly to application

C2

specialists and careful numerical experiments are performed to illustrate the computational benefits over intel mkl and other packages.

Interactive comment on Geosci. Model Dev. Discuss., doi:10.5194/gmd-2016-204, 2016.