

Geosci. Model Dev. Discuss., referee comment RC2
<https://doi.org/10.5194/gmd-2021-216-RC2>, 2021
© Author(s) 2021. This work is distributed under
the Creative Commons Attribution 4.0 License.



Comment on gmd-2021-216

Anonymous Referee #2

Referee comment on "MagIC v5.10: a two-dimensional MPI distribution for pseudo-spectral magneto hydrodynamics simulations in spherical geometry" by Rafael Lago et al., Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2021-216-RC2>, 2021

This paper presents the implementation of a 2D data parallelization strategy for the MagIC pseudo-spectral code. With this strategy, the three dimensional field data is distributed, using MPI, along two dimensions. As a consequence, two MPI communication stages are required during the spectral transform from physical to spectral space and vice-versa. This strategy has been presented and discussed in other publications. In this regard, the manuscript doesn't provide a novel parallelization approach but a discussion on the technical aspects of the implementation in MagIC. A detailed comparison to the hybrid MPI and OpenMP parallelization, which involves a single MPI communication stage, available in MagIC is made. Performance benchmarks are used to illustrate which strategy is best depending on the computational resources. In this light and as a development and technical paper, I recommend to accept this paper with minor changes. The questions I would like to be addressed are given below.

Considering that a large portion of today's supercomputers include accelerators, possibly making up most of their raw computational power, does the 2D strategy bring any advantages with regards to using accelerators?

The discussion of the transposition in section 3.3 is confusing. The paragraph starting at line 335 discusses the importance of the size of the queue, but the MPI algorithms (l. 325) specify a MPI communication "per scalar field". The discussion in Section 4.1 also seems to imply a single communication call per "queue". Can you clarify this?

The performance benchmarks provide insight into the behaviour of the 2D parallelization implementation but are executed on a single cluster. A more general discussion on what kind of performance to expect on another HPC cluster depending on its technical characteristics would be helpful for the wider community.

The strong scaling experiment with the 1D hybrid strategy increases the number of threads for runs with higher number of cores. This is counter intuitive as it leads to even smaller computational load per thread which I would expect to affect the scaling negatively. Can you explain this behaviour?

At high resolution, the memory footprint of a pseudo-spectral code can become important.

Is there a benefit, at the memory level, to use a 2D distribution in MagIC?

Corrections:

- l. 132: P_{lm} should be the "associated Legendre polynomials" and not the "Legendre polynomials".
- Table 3: Adding a κT_r column would make it easier to follow the discussion at the end of Section 4.6