Geoscientific Model Development Discussions

Geosci. Model Dev. Discuss., author comment AC1
https://doi.org/10.5194/gmd-2021-216-AC1, 2021

Reply on RC1

Rafael Lago et al.

Author comment on "MagIC v5.10: a two-dimensional message-passing interface (MPI) distribution for pseudo-spectral magnetohydrodynamics simulations in spherical geometry" by Rafael Lago et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2021-216-AC1, 2021

We'd like to thank the referee for her/his very thorough and constructive review and for the suggestions and corrections which were proposed. In the following we individually address all points raised by the Referee #1:

> First, I would like to point out that the present scaling tests are not
> in the practical range for the productive runs. Dynamo simulations in
> a rotating spherical shell are performed a few million time steps or
> even tens thousand steps. The minimum elapsed time is still more than
> one second in Figure 7, so it suggests that the present model needs
> approximately 12 days for one million steps. I guess that the practical
> problem size would be the half of the horizontal resolution for the
> productive runs. I recommend that the author describe the reason how
> they choose the spatial resolution and target elapsed time for the
> productive runs.

The value of δt (and thus, the total number of time steps) depends not only on the grid resolution, but also on the nature of the physical problem and its control parameters (Ekman, Prandtl and Rayleigh numbers). As an example, convection in thin spherical shells will require much larger spatial resolution but will converge in a significantly smaller number of iterations, than, say, geodynamo models geared to model reversals of the Earth magnetic field. As such, the total number of iterations is not a really meaningful measure.

However, as explained in lines 436 and 467 of the revised text, the times are measured for 100 timesteps. In the largest run (24,000 cores), the main application time was 6.4 seconds, thus one million timesteps would require 18 hours in average. For the the largest "recommended" run (12,000 cores, where the parallel efficiency remains within acceptable levels), one million time steps are estimated to require 20 hours in average.

Finally, the development of the 2d-MPI version is justified mostly for large grids. The chosen resolution is comparable to current state-of-the-art geodynamo models (e.g. Schaeffer et al. 2017). For smaller grids, the 1d-hybrid version of the code should deliver the optimal performance.

> Another question is that the authors perform data communications for
> each radial layer and each scalar component in the 2D parallelization.

> Calypso and Rayleigh perform these communications with single
> MPI_ISEND/IRECV or MPI_ALLREDUCEV, respectively. Can authors discuss
> the advantage of the present communications from Calypso or Rayleigh's
> approach?
As noted, a single communication involving all fields and all radii could be performed. In the early stages of the project we have performed tests using such a strategy. This has the following consequences:
(1) allows much large message sizes
(2) fewer synchronization points are needed
(3) memory requirement grows with the number of radial points (since all fields and buffers for each radial point must be stored imultaneously for a single communication).

Our Figure 3 shows that "packing" more messages past 1,750KiB (queue of length 9) does not provide any visible performance benefit for the hardware we used in our tests, voiding (1). Concerning (2), some of the early experiments showed an equivalent or slightly superior performance of the queue algorithm when many radial points were used. Moreover, for simulations involving many radial points, (3) is exacerbated.

We agree that the combining all messages in a single communication could be useful in several scenarios and we may revisit this part of MagIC in future releases, enabling the user to choose between strategies.

> And, SHTns is used in the present study. I wonder if the authors calculate
> the Legendre polynomials at Gauss-Legendre points in the initialization, or
> calculate during each Legendre transforms. I remember that SHTns has both
> feature, so it would be helpful which approach is chosen and why the authors
> choose one.
The "on-the-fly" strategy is used. This is the recommendation of SHTns' author for $\ell_{max}>32$ (see (Schaeffer, 2013) for more details).

> Lastly, the authors defined Tt. However, I can't find any information for Tt
> in Table 3. Can I find the data from the other table? So, I lost a direction
> to figure out the following discussion using Tt in page 26. Please provide
> how to figure out Tt.
T_t and t_t are given in Table 2 under the column "time". In the revised text, we added these columns to Table 3 as well.

> And, these are some minor suggestions:
> In line 2, can 'magnetohydrodynamics' be one word?
We modified all occurrences throughout the text to "magnetohydrodynamics".

> In line 5, I think "parallelization" would be more explicit than "implementation".
We changed all occurrences of "hybrid implementation" to "hybrid parallelization" throughout the text, but we kept the term "1d-hybrid implementation", since it refers to a "version" of the code. We hope that this change meets the expectations.

> In line 27, "mag" would be capitalized as "MAG" or "Mag"
This has been corrected as requested.

> In line 65, Full name should be represented for "Non-Uniform Memory
> Access (NUMA)" first.
This has been modified, but on line 59 where it appears first.

> In line 117, It would be better to add dimensionless "self" gravity.
This has been modified as requested.

> In equation (14) and (15), the diffusion term appears in the left hand side
> and right hand side, respectively. I would like to show this term in the same
> side in the both equations.

This has been modified as requested.

> In line 202, I would like to show the equation using Mlδt and Blmt below of
> the equation (17). I looked for the definition of Mlδt and Blmt for a while.

We added a sentence in the text to clarify the definition of Mlδt and Blmt.

> in line 329 and 330, I prefer to say "component" instead of "field", if the
> scalar "field" includes toroidal and toroidal components for vectors.

We fear that using the word “components” may erroneously lead the reader to believe that we are solving for vector components. To avoid this confusion we prefer to refer to them “scalar fields”.