Review of "Accelerating 3D magnetotelluric forward modelling with domain decomposition and order-reduction methods" by L. Tao, A. Muixi, S. Zlotnik, F. Zyserman, J. Afonso, P. Diez

August 1, 2026

This paper develops a domain decomposition (DD) formulation for solving the 3D magnetotelluric (MT) equations, paired with a POD/Galerkin dimension reduction technique. The coupling is achieved via the non-overlapping Schwarz alternating method with Robin-Robin transmission conditions. The proposed framework allows one to couple subdomain-local finite element models, or subdomain-local projection-based reduced order models (ROMs). The resulting coupled models are termed DD and DD-POD, respectively. It is demonstrated that the DD-based models achieve significant speedups over their non-DD monolithic counterparts on two examples, a canonical toy problem, and a real-world problem.

The approach presented in the paper is timely, as DD-based methods involving ROMs have been taking off in recent years, and this topic is a particular research interest of mine. I believe it is among the first, or maybe the first, to combine DD and ROM for the MT application (caveat: I am not an expert on MT). The paper is interesting and well-written, but unfortunately there are numerous important details about the methodology and numerical results that are missing. Some important references are missing as well. I also find the results that the DD variant of the method is significantly cheaper than the non-DD extremely surprising, especially since this is a non-overlapping Schwarz coupling framework, and a relatively large number of subdomains is being coupled, I believe. More evidence is needed to justify/explain these results. My recommendation for the paper is "Major Revision". I ask that the authors revise their paper to address the following comments.

1 Major comments

1. There has been a lot of recent work on DD + ROM, including DD + ROM + Schwarz, and it is worth citing some of it. Please refer to

   - Ruan, Class, Rozza, A Structured Review of Reduced Order Modeling for Domain Decomposition Problems: State of the Art and Perspectives, 2026.

   for a comprehensive review of what has been done with DD + ROM until now. In particular, I would cite:

   - Moore, Wentland, Gruber, Tezaur, Domain Decomposition-based Coupling of Operator Inference Reduced Order Models via the Schwarz Alternating Method, 2024

- Rodriguez, C., Tezaur, I., Mota, A., Gruber, A., Parish, E., and Wentland, C. (2025). Transmission Conditions for the Non-Overlapping Schwarz Coupling of Full Order and Operator Inference Models. arXiv:2509.12228.
- Hoang, C., Choi, Y., and Carlberg, K. (2021). Domain-decomposition least-squares Petrov–Galerkin (DD-LSPG) nonlinear model reduction. Computer Methods in Applied Mechanics and Engineering, 384, 113997.
- Diaz, P., Choi, Y., and Heinkenschloss, M. (2024). A fast and accurate domain decomposition nonlinear manifold reduced-order model. Computer Methods in Applied Mechanics and Engineering.
- Chung, S. W., Choi, Y., Roy, P., Moore, T., Roy, T., Lin, T. Y., Nguyen, D. T., Hahn, C., Duoss, E. B., and Baker, S. E. (2024). Train small, model big: Scalable physics simulators via reduced order modeling and domain decomposition. Computer Methods in Applied Mechanics and Engineering, 427, 117041.

2. I am confused by the discussion on p. 7 about enforcing the BC at a single point. The relevant BC is a Robin one, not a Dirichlet one. Thus, by construction, it requires multple points along a subdomain boundary to define. Can you please clarify what is meant by the "single point enforcement"? Are you perhaps integrating over the boundary using a single integration point?

3. I was surprised by the introduction of Lagrange multipliers (LMs) within the Schwarz iteration on p. 7. One of the principal advantages of Robin-Schwarz methods is their minimal intrusiveness: existing subdomain solvers can typically be reused with only modified interface boundary conditions and data exchange. Introducing Lagrange multipliers adds interface unknowns and a saddle-point coupling, which appears to increase implementation complexity. It would be helpful if the authors clarified why this additional machinery is necessary and what advantages it provides over a purely Robin-Schwarz formulation.

4. It is also not clear to me what exactly the LMs are enforcing on p. 7.

5. Still on the same note, the motivation for introducing the Lagrange multiplier formulation and the single-point treatment of the interface conditions is not entirely clear. In Schwarz methods, the dominant computational cost typically arises from the repeated subdomain solves over multiple interface iterations, whereas the enforcement of the transmission conditions is relatively inexpensive. It would therefore be helpful for the authors to explain what computational or algorithmic benefit these additional interface treatments provide, for example in terms of reduced Schwarz iterations, improved accuracy, or reduced online cost. Did the authors ever try the method without these two "speed-ups"? What was the relative performance? If the LM makes it not that much better, I would say giving up the non-intrusiveness of regular Schwarz by introducing it is not worth it.

6. The Schwarz method can technically handle non-conformal meshes in the subdomains, but interpolation must be introduced. Is this use case of interest to you?

7. The reduced systems appear to be solved using an iterative linear solver. Since projection-based ROMs typically yield relatively small dense systems, a direct dense factorization is often more efficient and avoids iterative convergence issues. Could the authors comment on their choice of iterative solver and whether direct dense solvers were considered or compared?

8. The proposed framework appears to couple reduced models on all subdomains. Could the methodology be extended to hybrid FOM–ROM coupling, as in some of the references above, where only a subset of subdomains is reduced? Such a capability would increase flexibility by allowing reduced and full-order subdomains to coexist within the Schwarz iteration.

9. The offline training procedure is not described in sufficient detail, especially for the first problem. It would be helpful to specify the sampling strategy for the conductivity models, the number of training realizations, the snapshot collection procedure, the POD truncation criterion, and whether separate reduced bases are constructed for different frequencies. These details are important for reproducibility and for assessing the offline cost of the proposed approach.

10. How were the basis sizes selected? Was it using an energy criterion?

11. The reported speedup of the domain-decomposed ROM over the global ROM, and the domain-decomposed full order model (FOM) over the global FOM, is extremely surprising. Classical Schwarz methods, particularly non-overlapping variants, generally introduce additional interface iterations and therefore do not reduce the total computational work in serial. I agree that the subdomain problems are easier to solve. However, with nine subdomains for the first problem, one would expect the domain decomposition overhead to offset the benefit of the smaller local systems. It would be helpful if the authors could provide additional insight into the source of the observed speedup–for example, by reporting the local and global reduced dimensions, the number of Schwarz iterations, and a timing breakdown separating local solves from interface iteration. Are the comparisons in CPU times apples-to-apples, i.e., is the same amount of computational resources utilized for each set of runs?

12. The reported timing results for the domain-decomposition method are difficult to assess without the corresponding Schwarz iteration counts. Please report the number of iterations required for convergence for each numerical experiment, together with the stopping tolerance, relaxation strategy, and any variation with frequency or subdomain count. This information is essential for interpreting both the computational cost and the scalability of the method.

13. The paper appears to use a single reduced basis dimension for each polarization across all subdomains. Since the POD bases are constructed locally, did the authors consider selecting the basis dimension independently for each subdomain using a local truncation criterion? Such an approach could further reduce the online cost and better reflect the varying complexity of the local solution manifolds.

14. I am surprised the local ROMs are so much faster given that the number of modes in the local ROMs seems to be quite large (over 100 possibly). I assume the global ROM had the same number of modes? Please comment. It is not entirely clear whether the reported ROM assessments are reproductive or predictive. Please clarify whether the test conductivity models are independent of the training set and, if so, how the training and testing parameter sets were partitioned. If the assessments are primarily reproductive, this should be stated explicitly. Moreover, I would like to see some predictive results, as the predictive regime is the true regime in which the ROM should be assessed.

15. I don't believe the authors say how many subdomains are used in the realistic example, or how the DD is done. Please add this information.

16. An interesting observation is that DD-POD is more accurate than G-POD for the first example but less accurate for the second. This reversal is not discussed in the manuscript. Some insight into the factors governing when local POD bases outperform a global basis (or vice versa) would strengthen the paper and help readers understand the advantages and limitations of the proposed approach.

17. The performance of Robin-Schwarz methods can depend strongly on the choice of the Robin transmission parameter. The manuscript would benefit from additional discussion of how this parameter was selected. Was a single value used throughout, or was it tuned for each frequency and test case? Some discussion of the sensitivity of the convergence and runtime to the Robin parameter would help readers assess the robustness and reproducibility of the proposed approach.

2 Minor comments

- In the introduction, the authors say that "matrix-based approaches generally lead to ill-posed inverse problems". This statement is misleading and conflates the issue. The problem is not that the approaches are matrix-based, it is that the inverse problem is ill-posed when posed in the deterministic setting. Often regularization is used to mitigate the issue. Please reword this phrase to make it more accurate.

- The error metric in equation (12) is a bit weird to me. Normally, you would create a relative error, so you would divide simply by $\sum_j ||E_j^{(n+1)}||$. I am curious, why do the authors adopt the metric in (12) rather than the usual one?

- In Figure 1, why is the DD only done into horizontal subdomains?

- In Figure 5, please clarify what the integers next to "XY=" and "YX=" mean.

- Tables 5-6 report absolute errors, which are generally not very informative. Please switch to relative errors.

- Please change "validation" to "verification". The former requires comparing to real-world data from the field or experiments. What you are doing is effectively verification.

- At first, I was wondering about extensions to nonlinear problems, but when I looked it up, it appears the linaer MT equations are the state-of-the-art ones to use; there is not a relevant nonlinear version. You could highlight this when you present the equations or in the introduction, in case people criticize the approach for being linear.