Geosci. Model Dev. Discuss.,
https://doi.org/10.5194/gmd-2020-158-RC1, 2020

Geoscientific Model Development Discussions

Interactive comment on "Development and performance optimization of a parallel computing infrastructure for an unstructured-mesh modelling framework" *by* Zhuang Liu et al.

Anonymous Referee #1

Received and published: 23 October 2020

The paper is well written and clearly describes what work the authors carried out. The explanation of the framework, how it works, and all of the components related to the paper's content are clear and easy to understand. It is also pretty easy to understand from this paper that this framework took a lot of effort and it's important not to underestimate the contribution that a modeling framework provides. The paper explores many avenues for performance improvement, such as examining local mesh reordering techniques to improve cache reuse.

While the paper is well written, it is difficult to understand why some of the choices were made. The paper heavily references another modeling framework, MPAS, which

Printer-friendly version

Discussion paper

does largely the same thing as GRIST. The overarching discretization, mesh choices, decomposition methods, and parallel I/O strategies seem almost identical to those used in MPAS. It would be good to have some additional paragraphs in the paper describing why this new modeling framework is an addition to the community on top of the existing open-sourced modeling frameworks. Two suggestions for this would be:

A new paragraph in the introduction explaining why this framework is necessary given other frameworks that already exist in an open-sourced format. I think this would help avoid any assumptions of redundancy that could be confusing to readers familiar with other efforts in this space. A new paragraph in the conclusion comparing and contrasting this new framework with other existing frameworks. This could also just be a table or whatever format is easy for the authors.

Aside from that, the paper positions itself as if it is going to look at many sweeping avenues of performance optimization, and while something like local mesh reordering is potentially important (especially if the users primarily focus on low processor count runs) it can easily be out performed by studying the on-node performance impacts, or even looking at inefficiencies from the MPI implementation. In future work (because it is a lot of effort, and probably outside of the scope of this paper) I'd highly recommend looking at how well the model vectorizes, and performs on a single node before looking at things like cache reuse. Especially since the data provided (and previous studies on this subject) show that you can negate the impact of reordering by strong scaling your model out further. This also means you could simply over-decompose your problem to avoid having to reorder data. While the authors argue that their framework shows a good scaling efficiency, it's hard to accurately estimate how good the strong scaling effects of the framework are without some serial baseline numbers. As a result, I think it would be beneficial for the authors to try and add some serial (and lower processor count) data to figure 3. It looks like the authors stop exploring their strong scaling space at around 250 processes, and it's hard to judge what the loss in parallel efficiency is without these numbers. In the end, I would suggest three major revisions elements,

Printer-friendly version

Discussion paper

mainly to help emphasize the contributions that this framework brings to the community, and to help showcase how efficient the framework is as a whole. These are listed below: A new introductory paragraph as mentioned above A new concluding paragraph or table as mentioned above Adding lower processor data to the strong scaling plots. Ideally all the way down to a single process.

Outside of those major revision elements, here are some minor points: While the paper as a whole recognizes prior contributions by frameworks such as MPAS, the abstract does not. It would be useful to add something to that effect into the abstract. Based on the usage of METIS, I'm assuming that the authors use the offline capability (specifically in METS) and not the online capability in ParMETIS, but this could be clarified Line 40: "which reduces the" -> "which can reduce the". Line 80: "establishment of numerical modelling" -> "establishment of a numerical modelling" Section 2.3: You never mention that the G8 performance actually is degraded by using the BFS reordering strategy. This is also mentioned again on lines 370-372. This is important to notice, because the reordering strategies become unnecessary once the problems are strong scaled out to a certain point. It could also be useful to give a description of what this point is as a function of number of cells / process. Section 4.2: This essentially describes NCAR's PIO library, which is available on github. Though no reference to it appears in this section, and no discussion of why this was rewritten instead of used wholesale appears. This should definitely be added. It could be as simple as describing why it is not the same as NCAR's PIO library, and moving on from there though. Figure 4: The title on the top left sub-plot is mirrored

Referee criteria: Scientific significance: 3-4 (Fair - Poor) Scientific Quality: 2 (Good) Scientific Reproducibility: 2 (Good) Presentation Quality: 1 (Excellent)

Interactive comment on Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2020-158, 2020.