Reply to Reviewer 3

General Comments:

The authors present a GPU acceleration of the source term part within the parallel context of the WW3 Framework. The work is innovative and important. However, there are some flaws within the theoretical approach and the tests that have been done when evaluating the new parallelization option.

In eq. (1) the wave action equation we have certain terms that are local parts (e.g. source terms and spectral advection), and global parts, which is the geographical advection. The geo. advection needs some parallel exchange either for the CD or the DD approach.

Now, when expanding this asymptotically using Amdahl's law for the given problem it can easily be seen that ultimatievly for a infinite number of computational cores the only cost that remains would be the parallell exchange since all other workload tends to zero.

Introducing now the communicators to the GPU, this would remain as well an overhead and add up to the global exchange of the advection part itself.

Now the most important question is how does the scheme scale for various constellation. Since the GPU layer was introduced the scalability analysis becomes twodimensional in terms of number of GPU and CPU. This question remain open in this paper even if the authors have sufficient acces to the needed computational resources. In the sense of the above also the quantification of the computation cost of the source terms is rather linked to give testcase constellation investigated in this paper.

I conclusion I think that the work is interesting and the implemenation is a important topic for GMD but the work lacks in depth scalability analyis and therefore no final statment can be made in terms of efficiency. Especially, in the context that only 8 cores have been used from possible NSPEC cores within the CD approach.

In conclusion I think that much more work must be done to evaluate the performance of this approach before the given conclusion can be made and the general contribution of the work can be evaluated, which is now not the case.

We wish to thank the reviewer for their comments and suggestions regarding the manuscript. The focus of this study is to accelerate WW3 for global modeling and prepare the wave model for the E3SM exascale regime. In light of this, we focused on the current WW3 configurations in E3SM and only at the most computationally intensive region, which is the source term (Fig. 1). This is true for any propagation schemes or grid parallelization concepts at least when using ST4 source term. Other parts of WW3 can benefit from GPU, but as the first GPU work in spectral wave modeling, our aim is to focus on the expensive region first. As suggested by the reviewer, we have also included a scaling analysis up to 5 nodes (i.e. 210 CPU cores and 30 GPUs) in the manuscript. Table 4 in the manuscript shows the results of scaling the 228K mesh size over multiple nodes. We used the full GPU configuration here by launching 7 MPI ranks on each GPU. In the results, it can be seen that the speedup is relatively uniform across multiple nodes.

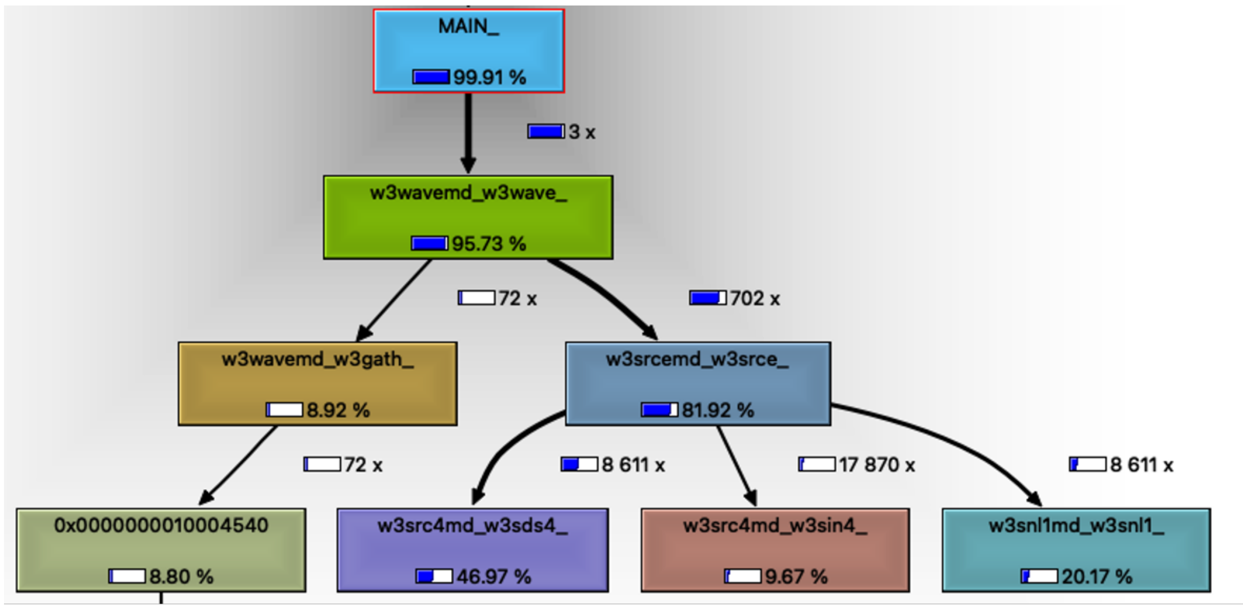

Figure 1: The Callgraph obtained by profiling WAVEWATCH III with 300 MPI ranks. Each box includes the name of the subroutine and its relative execution time as a percentage

Specific Comments:

1. 94: "is the most computationally intensive part of WW3" Can this be quantified? Moreover, I can not agree on this, since this would be rather linked to the given configuration and the used schemes. For implicit schemes and various other constellations with high resolution geographical space this must not be true. I think that too much general statements have been derived by the given configuration and testcase. I could imagine that with high spatial resolution and a lot of computational cores, which have unfortunately not been used here, the communication itself will take more time than the computation of the source terms itself, as explained in the general comments above.

The authors are using CD for their parallelization strategy, but I do not see why one could not use domain decomposition in combination with GPU. How would the scheme perform with DD approach? Why so few computation cores. Why so much development and such little evaluation of the performance?

The authors have limited their simulation to just 8 processors, however, I do not see why it could not be used on say 1000 processors if each of those processors has access to a GPU. This cannot be repeated often enough. It remains the major concern of this work. The Kodiak and Summit supercomputer have several thousand processors. Why is the limit to 8?

The interaction between explicit and implicit computations is not considered. I would think that the code can be used for explicit and implicit and the scalability should be evaluated for both.

Further code moving to the GPU could be the frequency shifting and refraction in explicit mode, has this been considered as well?

We focused on the current WW3 grid parallelization method in E3SM. The number of processors used in profiling does not change the fact that the source term is the most expensive region. Figure 1 shows the callgraph obtained by profiling WW3 with 300 MPI ranks. The source term is still the most computationally intensive region. We have updated Fig. 4 in the revised manuscript. Since this is the first attempt to run WW3 GPU, future studies would focus more on how different WW3 setup (grid parallelization method, source term switches, propagation schemes, e.t.c.) impacts GPU speedup over the CPU.

We have conducted sensitivity studies (performed on only CPUs, at this point) to determine the potential speed up due to domain decomposition (as opposed to card deck decomposition) parallelization methods as well as the impact of explicit computations (as opposed to implicit). We found domain decomposition and implicit computations perform significantly worse (fewer Simulated Years Per Day (SYPD)) than explicit, card deck computations for the global E3SM grid used in this study (results shown in Figures 2 and 3 below). We believe this is due to the fact that the E3SM mesh considered here is smaller and has a less drastic difference between the minimum and maximum resolution compared to the mesh used in Abdolali et al. 2020.

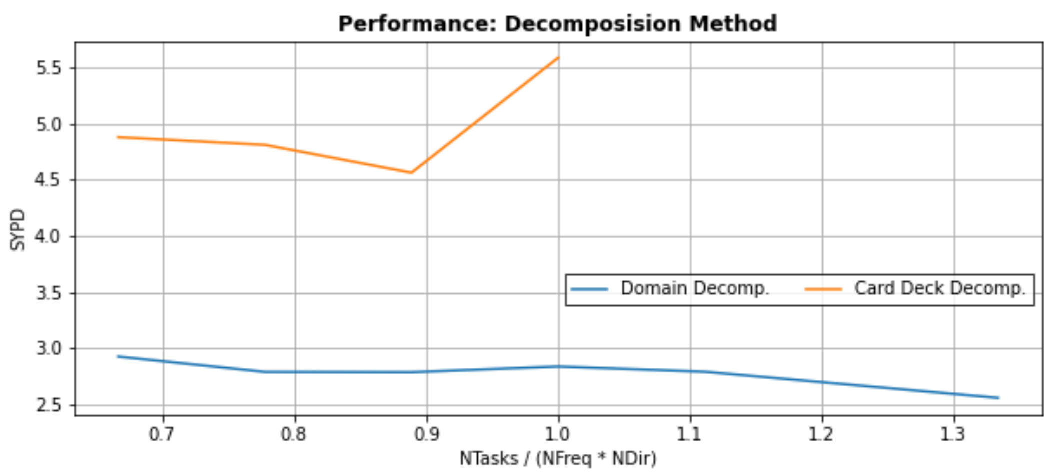

Figure 2. Simulated Years Per Day (SYPD) of Explicit, Domain Decomposition (blue line) and Explicit, Card Deck Decomposition (orange line) simulations of WW3 plotted as a function of the number of tasks. For direct comparison, the simulations are plotted by the number of tasks normalized by the spectral resolution. Note that the number of tasks used by Explicit, Card Deck Decomposition is limited by the spectral resolution (the maximum tasks used cannot exceed the spectral resolution). Explicit, Card Deck Decomposition clearly outperforms Explicit, Domain Decomposition.

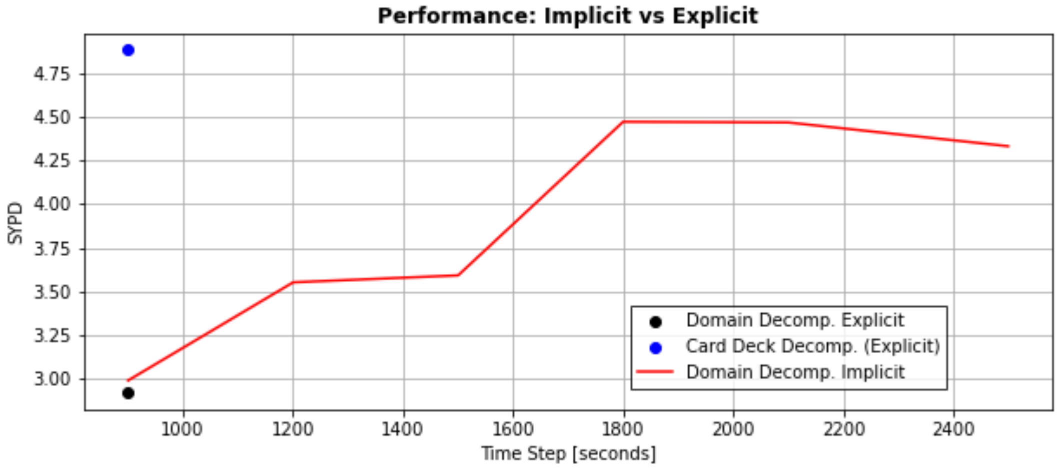

Figure 3. Simulated Years Per Day (SYPD) of Explicit, Card Deck Decomposition (blue point), and Explicit, Domain Decomposition (black point) and Implicit, Domain Decomposition as a function of time step (red line). The explicit simulations use a time step of 900 seconds, while the time step of the implicit simulations varies between 900 seconds and 2600 seconds. Increasing the time step of the implicit simulations, increases the SYPD of WW3 until about 1800 seconds, beyond which, the time step does not speed up (or increase the SYPD) of WW3.

Technical issues:

- There is a typo in Equation (1), the "+" is missing

  We have corrected the typo.

  The literature reference in the paper is not done properly. Here are two examples:

- 19: Either referer to the paper with „e.g." or put the original reference, which is not Hasselmann 1991. It was Gelci et al. 1957 if i remember right.
- 146: Again, Chawla et al. Was referenced but those authors did not derive the scaling of 1.1. The original publication should be cited Hasselmann & Hasselmann, 1981 or it can be add „e.g.". However, the latter I do not find very useful since we want to honor and cite the original work. This should be cleaned throughout the paper.

  We are attributing our entire spectral mesh settings to Chawla and not the 1.1 factor. We have modified the sentence

Missing reference:

- 104: Brus et al., 2021 reference is missing.

  We have corrected the missing reference (Line 111)

All references have been checked.