Review of **"Towards a real-time modeling of global ocean waves by the fully GPU-accelerated spectral wave model WAM6-GPU"**

Summary

The authors have fully ported the spectral wave model (WAM6) to GPU using OpenACC with a substantial amount of code refactoring. On a GPU cluster with 32-core Intel Xeon6326 and 8 NVIDIA A100 GPUs, the WAM6-GPU code achieved a speed-up of 37x when utilizing all the resources on a node. As a result, they achieved around 90% reduction in power consumption.

This is an important study that would enable century-long global simulations with a stand-alone wave model and also facilitate the integration of wave models into Earth system models. However, before accepting this manuscript, the authors need to address the following issues thoroughly.

Main:

1. In the abstract, the authors need to state the speed-up value based on a node comparison e.g., 32-core intel Xeon6326 and 8 NVIDIA A100.

2. Looking into the code, I saw that most of the subroutines/modules were refactored. A rough estimate of how much the original CPU code has been refactored should be discussed within the manuscript.

3. One important thing missing from this paper is the structure of the WAM code. The authors should include a skeletal code structure of both CPU and GPU versions of some parts of the code. This would greatly improve the manuscript for readers, especially for understanding the S_{NL} optimization explained in line 245-255.

4. The use of two CPU-only HPC clusters is confusing. Given that the study focuses on GPU and not the optimization of the CPU code on the CPU, I think there is no need to run the CPU code on two CPU-only HPCs. Since the NMEFC's GPU server does not have more than one node needed for scalability of the GPU code, the authors should only keep the NMEFC's HPC cluster for comparing resource usage needed to achieve the GPU execution time.

5. Fig. 7: The authors should show the spatial difference between the output parameters generated by the WAM6-GPU and the CPU version. Mean difference (Fig. 8) sometimes averages out the spatial difference between, if any.

6. Apart from running on the NVIDIA H100 GPU, are there any other further optimization strategies to improve the WAM6-GPU code on A100?

7. Just curious. Considering this study started in 2020, I wonder if the authors used P100 and V100. If so, what were the achieved speed ups?

Minor:

1. Line 11: This is a scientific dataset. Cite Cavaleri et al., 2012 as in Line 20
2. Line 13: Check citation format.
3. Line 33: The new U.S. Department of Energy (DOE) Energy Exascale Earth System Model (E3SM) has also included WW3 as part of the default component. Cite Ikuyajolu et al., 2024 and Brus et al., 2021
4. Line 228: Define all terms in the equation
5. Figure 6: Check caption for incorrect latex degree symbol

Reference:

Ikuyajolu, O. J., L. Van Roekel, S. R. Brus, E. E. Thomas, Y. Deng, and J. J. Benedict, 2024: Effects of Surface Turbulence Flux Parameterizations on the MJO: The Role of Ocean Surface Waves. J. Climate, https://doi.org/10.1175/JCLI-D-23-0490.1, in press.

Brus, S. R., Wolfram, P. J., Van Roekel, L. P., and Meixner, J. D.: Unstructured global to coastal wave modeling for the Energy Exascale Earth System Model using WAVEWATCH III version 6.07, Geosci. Model Dev., 14, 2917–2938, https://doi.org/10.5194/gmd-14-2917-2021, 2021.

Cavaleri, L., Fox-Kemper, B., and Hemer, M.:WindWaves in the Coupled Climate System, Bulletin of the American Meteorological Society, 93, 1651 – 1661, https://doi.org/10.1175/BAMS-D-11-00170.1, 2012.