EGU21-11551
https://doi.org/10.5194/egusphere-egu21-11551
EGU General Assembly 2021

GPU acceleration of the FESOM-2 ocean and sea-ice model

Gijs van den Oord[1], Alessio Sclocco[1], Georges-Emmanuel Moulard[2], David Guibert[2], Dmitry Sidorenko[3], Nikolay Koldunov[3], Ben van Werkhoven[1], Erwan Raffin[2], and Natalja Rakowski[3]
[1]Netherlands eScience Center, Amsterdam, Netherlands
[2]Atos Center for Excellence in Performance Programming, Rennes, France
[3]Alfred Wegener Institute, Bremen, Germany

FESOM-2 is a finite volume ocean circulation and sea ice model developed by the Alfred Wegener Institute (AWI). It solves the primitive equations using the hydrostatic and Bousinessq approximations on an unstructured grid, allowing seamless mesh resolution increase towards eddy-resolving scales in regions of high variability or along coast lines. FESOM-2 is a highly optimized MPI-parallel Fortran code that displays excellent scaling to tens of thousands of cores. In the context of ESiWACE-2 services, we have explored the benefits of GPU acceleration of FESOM-2 in a six-month engineering effort. We have determined the flux-corrected tracer transport, and in particular the advection of temperature and salinity, to be a dominant factor in the application profile and we have ported this routine to GPUs using both OpenACC and CUDA-C. We conclude that the memory access patterns in FESOM-2 are suitable to map onto GPU accelerators and that both strategies are viable options, giving significant speedups for tracer advection in high-resolution mesh configurations. We have benchmarked the ported application on Nvidia Kepler, Volta and Ampere architectures and observe that our tuned kernels can approach the peak memory bandwidth, and we also see that OpenACC offers a competitive performance with less development and maintenance effort. We conclude that an expansion of the OpenACC directives is the most promising road to utilize upcoming GPU-equipped exascale machines for FESOM-2.