



Interactive comment on “ASAMgpu V1.0 – a moist fully compressible atmospheric model using graphics processing units (GPUs)” by S. Horn

Anonymous Referee #2

Received and published: 5 December 2011

This paper mostly describes the scientific development of a new weather model specifically to run on GPUs. However, the paper does not contain sufficient depth in the GPU work, or identify new scientific aspects of the model that are new or significant. If the science is new or innovative, the author does not identify these areas. If the GPU elements are significant, the author does not have sufficient detail to be useful, or even to evaluate the quality of the work.

He did not describe the hardware that was used or its performance - where did the 10TFlops peak performance come from? Was that the system being used? There were no performance numbers, no comparisons to CPU results, or scaling numbers for serial and multi-GPU runs. What percentage of peak did you get? Can it run on CPUs? How do the results differ? If the author chose OpenGL for portability reasons, did he

C1144

compare performance and accuracy on different platforms? Was it truly portable (no code changes)? There are no indications of communications costs in the multi-GPU environment. The parallelization strategy was vague with little discussion of the code design, efficiency, readability, etc.

There are no references indicating the author knows about other GPU parallelization efforts for weather and climate models. There is good work being done in the community. There were also no references to CUDA, OpenCL, GPU compilers, or an indication why OpenGL was chosen. Was it because the author was using an older GPU which did not support the HPC languages?

The scientific research seems promising though I am not a meteorologist so it's hard to evaluate merit. However, the scope seems limited to success running some idealized test cases. Is there a CPU version of the model so results can be evaluated? If not, can it be evaluated to other models of similar scale? When will real data be used?

The author also did not research other scientific efforts to develop and run non-hydrostatic models of this type. There are plenty of models already developed that are at the scale indicated and years of work have been done in this area - so again, the scientific content does not appear new or innovative.

While this topic is of great interest to the community, I cannot see any benefit to publishing this paper without major revisions.

Interactive comment on Geosci. Model Dev. Discuss., 4, 2635, 2011.

C1145