



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

S. Horn

horn@tropos.de

Received and published: 27 December 2011

The author thanks the reviewer for the constructive comments and will address the minor revisions in detail:

Q:Relative comparisons

A:The author agrees that relative comparisons are not suitable for scientific publications and will replace them with specific comparisons in a revised manuscript.

Q.1-13.: orthographic and grammar comments

A:The author agrees to those comments and will correct the mistakes in a revised manuscript.

C1241

Q.14: Page 2638, Line 10: split sentence. "GLSL supports functions, if-branches, and loops. The choice of these functions significantly affect the performance of the resultant code as if-branches can take XXX % longer to execute than other floating-point operations."

A:The author agrees to split the sentence, but it is not possible for him to define how much longer if-branches may take than other floating point operations, because that strongly depends on the content of the branch and the branching capabilities of the used hardware. Early GPUs for example processed both branches of an if-else statement and then saved the value only for the true condition. Modern GPUs support branching and can handle if statements much more like CPUs.

Q.15: Page 2638, Line 14: "a number of output texture framebuffers" should say how many output texture framebuffers to be more precise.

A:That differs between the calculation steps and the shader used. For example the shader for the microphysics calculates sinks and sources for 12 variables. For these three output texture buffers with four components each are used. The shader for momentum calculates sinks and sources for three variables, in this case only one texture buffer is necessary.

Q.16: Page 2638, Line 20: revise sentence so that it makes sense.

A:One of the main differences between OpenGL and OpenCL or CUDA is the fact, that with OpenGL the developer has no control on parallelization details like block sizes or the number of threads used by the GPU. The whole fine grain parallelization is done by the device driver transparent to the developer.

Q.17: Page 2638, Line 25: This sentence uses relative, qualitative comparisons (massive amount, wide bandwidth, most efficient). Those comparisons should be more precise.

A:This sentence sounds like a commercial, it will be removed.

C1242

Q.18: Page 2652, Line 6: remove "new" it is unnecessary as in a year it will be "old"

Q.19: Page 2652, Line 17: "ice-phase"

A:The author agrees.

Q.20: to add a figure tab showing the dimensions of a single cloud cell

A: Fig.1 (attached to this comment) is an illustration of the dimension of a single open cell surrounded by clouds. The green circle has a radius of 6.5 km. The author agrees to exchange Fig.5 in a revised manuscript to add this information.

Q.21: A figure or table that shows the relationship between precision/resolution gained, number of GPUs used, time of execution, and power consumed.

A: A plot for the very first scaling experiments with the values already presented in Tab.1 in the manuscript is shown in Fig.2 (attached to this comment). The first impression is, that for smaller domains which fit into the memory of one GPU, only small speedup is possible using two GPUs, most of the time is used for communication between the GPUs. I think we have some space for optimization there. For larger domains more GPUs are necessary because of the restricted amount of on-board memory. And it is hard to extrapolate scaling behavior from such small numbers of blocks, because with higher numbers of GPU threads and CPU nodes, more and more communication can happen independently from other blocks and simultaneously. So for systems with more than nine GPUs communication overhead is expected to not increase with domain size (except perhaps due data collisions in the connecting network). So there is a lot of future work that has to be done during the further development of this model approach.

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

C1243




Fig. 1. Visualisation of Cloud Field, the green circle has a radius of 6.5km

C1244




Fig. 2. First scaling experiments

C1245