

Interactive comment on “FAMOUS, faster: using parallel computing techniques to accelerate the FAMOUS/HadCM3 climate model with a focus on the radiative transfer algorithm” by P. Hanappe et al.

Anonymous Referee #1

Received and published: 3 August 2011

This paper presents a case study of porting a large, fortran, climate code to various multi-core and many-core architectures. The accelerated code showed very encouraging speed-ups with respect to the original version(s), after some serious (and imaginative) code restructuring. The authors' methodology of transforming the code in stages and testing after each showed good discipline (but see my first specific comment below). Their task was certainly helped by the fact that the performance of the original code was dominated by that of a handful of routines (short- and long-wave radiation). I am grateful for their frankness in including an account of the effort spent (2.5 person

C502

years), for this kind of information is very helpful for developers contemplating porting a major code to accelerator-based architectures. Overall, I liked this paper very much and have no hesitation in recommending that it be published, after making some very minor revisions, as suggested below.

Specific comments:

The description of methodology in section 4 relates how the first step in restructuring the radiation algorithm was to translate the code from Fortran to C, but does not explain why this step was deemed to be necessary. The reader deserves a full explanation as to why this step was undertaken. How much of the work could have been done in Fortran?

Minor and technical comments:

In the abstract, it is a little confusing as to what the speed-ups relate to. For example, the abstract says "... the new radiation code runs ... on graphics processors more than 2.5 times faster than the original code". This is strange, given that the original code was in fortran and hence does not run on graphics processors without some modification. My comment is a criticism about the language used in the abstract, not about the body of the paper, which does give sufficient information to clear up these ambiguities.

In the second paragraph of section 3, the last sentence does not explain why simplifying the logic of the main CPU and additional processors should help to remedy "the slowing down of the computation due to the latency of the data transfer to and from memory". Some attention should be given to the logical construction of this paragraph.

Last sentence of section 4.2. What does "quasi identical" mean?

Section 4.4, second paragraph. "Former vector machines can efficiently...". Consider re-phrasing this. Perhaps "Vector machines, formerly more widespread than today, can...". Or just drop the word "former" altogether.

C503

C504