Geosci. Model Dev. Discuss., 4, C530–C536, 2011
www.geosci-model-dev-discuss.net/4/C530/2011/

**Geoscientific
Model Development
Discussions**

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.

P. Hanappe et al.

hanappe@csl.sony.fr

Received and published: 12 August 2011

We would like to thank reviewer for taking the time to read the manuscript in detail and for his constructive comments. The comments will greatly help to improve the text and are integrated in the revised version of the document that we submit as a supplement to our reply. We also address the comments below.

Specific comments

RC: The abstract should point out that the optimisation work reported was carried out in the context of a version of FAMOUS that did not exploit any MPI-based domain decomposition parallelism (i.e. the base code was sequential).

AC: We reformulated the abstract as follows:

"Instead of the existing MPI-based domain decomposition, we used a task queue and a thread pool to schedule the computation of individual columns on the available processors."

RC: The time taken to undertake the porting of the algorithm is a key result being reported in this paper (authors should always be encouraged to publish such data!). I would like to see the effort mentioned in the abstract, rather than just being buried in the depths of the paper.

AC: We modified the abstract and added at the end that: *"We estimate that this project required around two and a half man-years of work."*

RC: The paper should point out that the CELL architecture was dropped by IBM in 2009. This could be done in a footnote added to the first paragraph on p.1276.

AC: There are contradictory rumors about the status of the CELL project at IBM but we did not find an official announcement that the architecture was dropped. IBM also continues to sell CELL-based rack blades. We did include the following footnote for completeness: *"At the time of writing, it is unclear whether the CELL product line will be further developed."*

RC: 4.2: Add motivation for the translation to C. At the time, this 'choice' would have been forced on the authors in order to exploit the hardware technologies (it is not such a forced choice today).

AC: We agree. This comment was also made by RC #1. We modified the first paragraph of Sect. 4.3 as follows:

"The initial hardware platform that we targeted in this project was the commercial version of the PlayStation 3 game console. Because no Fortran compiler existed for the SPEs, we were compelled to translate the radiation code to C. An additional motivation for this translation was the good support provided by most C compilers for the vector data types and SIMD instructions, as discussed in Sect. 4.5. Some support for SIMD instructions is provided by commercial Fortran compilers on other platforms but they also require significant code changes."

RC: 4.2: For the column entry, please state that Sect. 4.4 provides motivation for this reorganisation.

AC: We made the suggested change.

RC: 4.2: In the final para. on validation, please point to section 5.2 which provides some explanation. A definition of quasi-identical is required!

AC: We agree! We modified the sentence as follows: *"We ran a 120 year simulation and compared the statistical properties of the results against a reference run (see discussion in Sect. 5.2).*

RC: 4.2: Ideally, I would like to see some evidence that a FAMOUS that doesn't provide bit reproducible results to the original version of the code is useful to scientists. Bit reproducibility is a major issue for the Met Office, especially in climate work. Perhaps a reference to a paper discussing this issue (in the context of FAMOUS) could be added here.

AC: Bit reproducibility is generally also lost when the Met Office's Unified Model is ported to another platform. We did make some changes to the paragraph for clarity and added a reference to a paper that discusses the validation through bit reproducibility as it is used at the U.K. Met Office (S. M. Easterbrook and T. Johns, Engineering

the Software for Understanding Climate Change. IEEE Computing in Science and Engineering, Vol 11 (6), pp65-74. November 2009)

The new paragraph reads:

"We validated our changes in two ways. First, the binary output generated by our modified versions of FAMOUS is bit-wise identical with the binary output generated by the original version, when the code is generated without compiler optimisations. Bit reproducibility is a strong validation for code changes on the same computing platform (Easterbrook:2009). This test is not feasible for the spe version because we cannot run the original code on the SPEs and because the SPEs have a different floating-point implementation than generic CPUs. We therefore introduced the second test. We ran a 120 year simulation and compared the statistical properties of the results against a reference run (see discussion in Sect. 5.2)."

RC: 4.3: The first line states that it was 'decided' to translate the code. Please clarify whether this choice was forced. If not, please motivate the decision.

AC: We agree. This point was also mentioned by RC #1 and we have adapted the first paragraph of 4.3. (See reply to RC #1).

RC: 4.4: Para. 1, Add a sentence to the end of the para. stating that the motivation for the column-based reorganisation, and the benefits associated with it, are discussed in the rest of this section. (Otherwise the reader is left wondering why at this point.)

AC: We agree. We rephrased line 9, p. 1279 as follows: *"Because this organisation is more advantageous for the CELL processor, as discussed below, we restructured the algorithm taking the column as the guiding principle."*

RC: Section 4.6 and 4.7:

- I am slightly confused by the story of 'deciding' to develop a distributed version of the code (using a thread pool) and the subsequent, apparently independent decision to 'extend' this version to exploit the SPEs.

- It seems to me that the decision to implement a thread-pool version was required to enable the execution on multiple SPEs (having ruled out 'most of the higher-level library-based approaches').

- The point is that FAMOUS already supports MPI-based distribution of work - and MPI decomposition can obviously be used on a shared memory, multi-core machine, so why wasn't this explored instead, or as well as, the thread-pool approach? Please clarify the story here. Was it actually the case that the thread-pool approach was taken **because** it supported the exploitation of the SPEs in a more straightforward manner?

AC: The reviewer is correct that the main motivation for the use of the thread pool was to support the porting to the SPEs. We have made changes to the introduction in Section 4.6 to clarify our choices:

"We developed the multi-threaded version of the code in order to facilitate the distribution of the computation on the SPEs. Ordinarily, FAMOUS uses MPI-based domain decomposition to distribute the computation of the radiation. In this approach, it is important to distinguish between the technique of domain decomposition, used to divide the data set by the number of available computing nodes, and the MPI technology, used to distribute the computation to the computing nodes.

Domain decomposition is not well suited for the CELL processor, because, with only six SPEs, the resulting data sets are too large to be stored in the SPE's Local Store. The MPI standard would be an appropriate choice to distribute the computation to the SPEs, but no freely available implementation of MPI for the SPEs was available (see also the discussion in the next section).

These two constraints lead us to use a thread pool and a task queue to distribute the computation over several processors."

RC: p.1286, para. with line 15 in it. It would be worth stressing the point that the performance of the radiation calculations are not limited by the transfer of data between

the CPU and the device. This is a very positive point for weather and climate computations, particularly if this property turns out to be true of other physics computations. (Although the issue of data transfer will probably disappear in the future as 'devices' become integrated on chips.)

AC: We made changes to stress the positive results, but we did not venture to extend these results to other physics computations because they have very different computational patterns. We do believe these results should be reproducible for other column-based radiation algorithms. We modified the paragraph to mention this.

Technical corrections

RC: P.1280, line 8. By 'quite fast' you probably mean 'at an acceptable fraction of peak performance'... better to be quantitative if possible but 'quite' is not useful.

AC: We agree. A quantitative evaluation would require an analysis of the chache misses. Since this is out of the scope of this study, we have removed the reference to the 'quite fast' execution on modern CPUs.

RC: p.1282: A reference to the technique to handle conditional expressions would be useful.

AC: We added the reference.

RC: 5.1: I would like to see (a reference to) some discussion about the scientific acceptability of computing at single precision. A pointer to a paper confirming this is standard for the scientific uses of FAMOUS.

AC: It seems difficult to show the validity, or not, of single-point computation for scientific use in the general case. Concerning FAMOUS, the University of Oxford uses the single-precision version of FAMOUS in the grand ensemble experiments run through

the ClimatePrediction.net network. After extensive testing, there were no reason to believe that the single-precision arithmetic was insufficient to execute the algorithms. It was found, however, that several procedures on the ocean model were sensitive to rounding errors and are therefore executed with double-precision floating point arithmetic.

We have added a footnote at the end of line 23, p. 1284 stating that: *"ClimatePrediction.net also uses a single-precision version of FAMOUS."*

RC: p.1288, para. including line 20. The claim about energy reduction is out of place here, particularly given the early motivation that faster models are required to enable higher resolution and bigger ensembles! The energy case should be added to the statements in the introduction.

AC: We agree. Because it is not a central point we want to make in this paper, we have removed the claim.

RC: Fig. 5. Caption should read 'radiation as a function' rather than 'radiation in function'.

AC: We have corrected the caption.

Please also note the supplement to this comment:
http://www.geosci-model-dev-discuss.net/4/C530/2011/gmdd-4-C530-2011-supplement.pdf

---

Interactive comment on Geosci. Model Dev. Discuss., 4, 1273, 2011.