

MEETING SUMMARIES

HIGH-PERFORMANCE COMPUTING FOR CLIMATE MODELING

BY JEAN-CLAUDE ANDRÉ, GIOVANNI ALOISIO, JOACHIM BIERCAMP, REINHARD BUDICH,
SYLVIE JOUSSAUME, BRYAN LAWRENCE, AND SOPHIE VALCKE

The European Network for Earth System Modelling (ENES) is composed of those in the European scientific community who develop and apply climate models of the Earth system within the framework of its infrastructure project (IS-ENES). This workshop is the second of a series that started in mid-December of 2011 in Lecce, Italy, with a workshop devoted to “dynamical cores for climate models.” After the success of this first workshop, it was felt that there is continuous need for a place to discuss the main issues facing the European community involved in the development of climate models, and especially those related to the improvement and development of numerical models adapted to new architectures, both for computing and data management. Participation in the workshop was by invitation only, and the response was very enthusiastic, with more than

THE SECOND IS-ENES WORKSHOP ON HIGH-PERFORMANCE COMPUTING FOR CLIMATE MODELS

WHAT: Within the framework of the Infrastructure project of the European Network for Earth System Modelling (IS-ENES), more than 50 international developers and users of European Earth system models met to address the major issues facing today’s climate modelers, especially with regard to the efficient use of new super-computer architectures.

WHEN: 30 January–1 February 2013

WHERE: Toulouse, France

50 participants (see <https://verc.enes.org/ISENES2/archive/documents-1/is-enes-2nd-hpc-workshop-presentations-february-2013/list-of-participants>), including all the major climate modeling groups in Europe and some representatives from the United States (unfortunately, no one from Japan was able to travel to France to participate in the workshop).

The program was organized so as to review the context of European Union (EU) exascale projects (session 1), the main advances taking place in Europe and the United States (sessions 2 and 3) and resulting from projects running on the Partnership for Advanced Computing in Europe (PRACE; www.prace-ri.eu) platforms (session 4), and to discuss issues connected with the use of inhomogeneous—for example, general-purpose computation on graphics processing units (GPGPUs) or accelerator-based systems—architecture (session 5) and with new computing environments (session 6). The full program and the presentations are available on the IS-ENES

AFFILIATIONS: ANDRÉ—JCA Consultance and Analyse, Toulouse, France; ALOISIO—Università del Salento, and Centro Euro-Mediterraneo sui Cambiamenti Climatici, Lecce, Italy; BIERCAMP—Deutsches Klimarechenzentrum, Hamburg, Germany; BUDICH—Max-Planck-Institut für Meteorologie, Hamburg, Germany; JOUSSAUME—Laboratoire des Sciences du Climat et de l’Environnement, Gif sur Yvette, France; LAWRENCE—National Centre for Atmospheric Science, University of Reading, Reading, United Kingdom; VALCKE—Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique, Toulouse, France
CORRESPONDING AUTHOR: Jean-Claude André, JCA Consultance and Analyse, 5 Rue Saint-Antoine du T, 31000 Toulouse, France
E-mail: jc_andre@sfr.fr

DOI:10.1175/BAMS-D-13-00098.1

In final form 31 July 2013
©2014 American Meteorological Society

website (<https://verc.enes.org/ISENES2/archive/events/workshop-on-hpc-for-climate-models-january-30th-february-1st-2013-in-toulouse-france>). A seventh session was devoted to a general and strategically oriented discussion, from which recommendations for the high-performance computing community and generalized messages to supporting agencies could be prepared. The purpose of this short summary is to sum up these recommendations.

PERFORMANCE INTERCOMPARISONS.

There are now seven climate modeling groups within Europe participating in international activities, such as phase 5 of the Coupled Model Intercomparison Project (CMIP5) used in the Intergovernmental Panel on Climate Change (IPCC) assessments. The need for intercomparisons is of quickly growing importance, both to the advancement of the science and the shared exploitation of technical advances in efficiency of these various models. Without such active intercomparisons, there is a high risk that progress achieved by a particular group will not rapidly benefit others nor the community at large, and consequently, that the limited manpower available to the community will remain too scattered across the groups to achieve rapid scientific progress. The aim of such intercomparisons should be to facilitate the evaluation of both scientific and technical aspects of model code, so that best practices can be identified and shared. To be fully useful, they should be based on agreed metrics (for scaling and definition of the variables to be compared: simulated years/day, model configuration, horizontal resolution, etc.) and should also include metadata relative to the model components. In doing so, the community will be recognizing that the best practice for capacity simulations may be different than for capability simulations, both of which are needed, as emphasized in the IS-ENES strategy (Mitchell et al. 2012). In the former case, one optimizes for overall throughput, and in the latter for speed of particular simulations. Defining such metrics will require further discussion and work—this may be the theme of a third workshop.

European climate modeling groups access computers at their respective national level using so-called tier 1 computers and, for some of them, at the European level using so-called tier 0 computers operated by PRACE. The issues of accessing the high-performance computing (HPC) facilities are many, but they are linked largely by the way the computing centers and PRACE operate their computers. First of all climate simulations include both production runs (e.g., those requested by IPCC assessments)

for which tier 1, mostly national, machines are the most adequate and for which multiyear access is a requisite, and frontier runs (e.g., very high horizontal resolution runs or large ensembles of high-resolution members) for which only tier 0 machines are appropriate. The tier 0 platforms today should allow for the development, validation, and running of frontier applications that will tomorrow run operationally on what will be tier 1 systems. This raises the question of compatibility between tier 1 and tier 0 computers: if too large of a gap exists between the architectures, the time to port the codes and to achieve good science will be much too long. A necessary step to gain insight into such issues is to obtain access to the largest configurations of the most advanced tier 0 and tier 1 computers concurrently, which mandates a very good integration of tier 0 and tier 1 machines. This does not seem to be the case today, however. It is also a requirement that the peer review process for tier 0 access recognizes the necessity for large-scale large-resource development projects. It was decided at the workshop to collect from the large-scale projects already running under PRACE detailed feedback on their experience using these platforms in the current framework, in order to prepare for future interactions with PRACE and computing centers.

WHICH MODELS FOR PETASCALE AND EXASCALE?

Given the time needed for constructing a new climate model, it is crucial to assess whether new petascale and future exascale architectures will require developing models based on new principles. (It should, however, be *strongly* underscored here that most participants consider that technical efficiency is not an objective per se but is important to achieve in order to reach scientific goals; the driver for the technical developments is the climate science.) It also has been recognized that more effort should be made to better exploit the “complementarity” among climate scientists and computational scientists. Strategic approaches should be pursued to encourage and define interdisciplinary teams where computational and climate scientists can work together to address specific scientific issues. These efforts should leave climate scientists more time to work on solving the main scientific questions, to do better science, and to gain insight into key climate questions, while computational scientists can help in evaluating model performance and related strategies to improve their scalability. This approach would allow large simulations to efficiently run on high-end computing resources. Of course, this is not an easy task because different backgrounds, methodological approaches,

and goals need to be taken into account. However, these differences, which might seem a great barrier to working together, represent real added value if properly exploited.

One very big problem for climate models is to deal with the very high level of parallelism in modern computer architectures. A central part of all climate models is their “dynamical core”—the numerical representation of the model’s transport equations in the model code. The development of new dynamical cores (e.g., based on new grids) has been intensely worked on over the past 3–5 years. These new dynamical cores are presently used with success in a number of atmospheric models, especially within the United States, where runs are now possible that use up to 10^5 computational computing threads in parallel, while Europe is still a little behind. The experience from the United States shows that new dynamical cores are better able to exploit the highly parallel architecture of modern supercomputers even if some traditional codes still show good performance in a number of applications. In Europe, several groups are developing new dynamical cores for atmospheric models. The issue then no longer seems to be whether new dynamical cores are needed (in a sense that they cannot be considered as a disruptive technology anymore) but rather that their advancement and use in other parts of climate models (e.g., oceans) are continually reviewed.

Using GPGPUs for climate models has proven slightly disappointing, at least so far, with only a relatively modest increase in model performance. Issues raised by using GPGPUs, as well as by other types of (hybrid) computing architectures using a high level of parallelism on the chip, include insufficient main memory per computational task, the low available bandwidth to access the memory, multiple levels of parallelism (threads, tasks, computational units), and the silent errors, among others. Given the amount of effort necessary to solve such problems, and the current state of the supporting tooling, the community is not really enthusiastic about switching to these new types of supercomputers, at least in the short term!

General agreement was reached about the need for revisiting the model code structure, which was also recommended in the National Research Council (2012) report on climate modeling, and this could be another topic for a forthcoming workshop. Issues would be as follows:

- how to ensure more modularity in the codes (component approach) and better isolate the “science” from the underlying technical software layers

[code infrastructure; utilities for parallelization, input/output (I/O), etc.; and code superstructure, that is, the shell assembling and interconnecting the components]?

- how to separate the scientific software from underlying implementation using underlying software kernels that might utilize unfamiliar programming models?;
- how to access more efficient algorithms working with much higher parallelism (this is seen as a major disruptive technology with high positive influence on climate modeling techniques)?;
- more generally and on a longer time scale, whether we should try to converge on common code infrastructure and superstructure, and how to increase their adaptability and robustness.

THE DATA CHALLENGE. General consensus is that the exascale challenge for climate is more an *exabyte* challenge than an *exaflop* challenge! The community is likely to reach exascale with exabytes of data before it can exploit exaflop computing, and the biggest challenge today is to develop methods for handling high volume data, including active storage, dedicated data retrieval, and processing and analysis environments, customized for climate data. Models are indeed run without writing all the data produced, as selection of the data of interest for offline diagnostics and postprocessing can be easily done later. Even with such data selection, the actual volumes for storage are inadequate, both with fast storage for model products while simulations are running and for later analysis (whether fast or not). This is a clear limitation that needs to be solved: output from climate simulations is indeed of patrimonial value, and many groups are interpreting and intercomparing data from different models during a rather long period (months to a couple of years) after the simulations are completed.

It should be emphasized that technologies relating to data are currently not keeping pace with peak performance characteristics of computing systems: it is necessary to optimize the slow data flow through the numerous layers between applications and hardware. All these layers are influencing each other nonlinearly in many ways, often disadvantageous for performance. In the primary models, I/O and diagnostic processing servers have been, and are being developed to make model output asynchronous to the computation and to reduce some of the volume of model output. However, little work has been done on easy-to-use and efficient parallel data analysis tools (including optimal hardware environments) for postprocessing. As a consequence, despite significant

and increasing investments, I/O and data issues are expected to remain a problem in all parts of the simulation workflow en route to exascale.

INTERNAL AND EXTERNAL COMMUNITY COLLABORATIONS. Another point, already shortly addressed in the “Performance intercomparisons” section, concerns the collaborations that the climate modeling groups have to establish and reinforce. This is an important objective of IS-ENES and is one of the ENES infrastructure strategy recommendations.

Internal to the community, the need for more exchange is clear, either for comparing model performances, both scientifically and technically, or for sharing model components or software pieces. In this respect the developments undertaken by the various groups would all be facilitated by using open source approaches. Another issue is the possible buildup of virtual teams crosscutting the various modeling groups, in order to gather all specialists necessary to prepare and run large-scale projects, as some groups

are not of sufficient size and do not have diversity in competences and cannot always engage in all significant large-scale undertakings.

There is also a strong need to build better links to other disciplines and to establish more interdisciplinary teams, in which climate modelers would actively collaborate with applied mathematicians on the one hand (algorithms, solvers, etc.) and with computer scientists on the other hand (software environments, etc.).

REFERENCES

- Mitchell, J. F., R. Budich, S. Joussaume, B. Lawrence, and J. Marotzke, 2012: Infrastructure strategy for the European Earth system modelling community, 2012–2022. ENES, 34 pp. [Available online at <https://verc.enes.org/community/about-enes/the-future-of-enes/ENES%20foresight.pdf>.]
- National Research Council, 2012: *National Strategy for Advancing Climate Modeling*. The National Academies Press, 294 pp.