EUROPEAN COMMISSION
Community research

is-enes
INFRASTRUCTURE FOR THE EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

IS-ENES2 DELIVERABLE (D -N°: 10.1)

Documentation on benchmarking framework

(Original title: Documentation and User Guide for benchmarking framework)

File name: {IS-ENES2_Deliverable_10.1.pdf}

Author(s): Irina Fast, Joachim Biercamp, Marie-Alice Foujols, Graham Riley

Reviewer(s): Uwe Fladrich, Eric Maisonnave

Reporting period: 01/04/2013 – 30/09/2014

Release date for review: **24/02/2015**

Final date of issue: **10/03/2015**

Revision table			
Version	**Date**	**Name**	**Comments**
1	24/02/15	Irina Fast	First release for review
2	10/03/15	Irina Fast	Final version with reviewers' comments integrated

Abstract

This deliverable outlines our strategy on how to implement a technical and organisational framework for development, distribution and evaluation of application benchmarks for Earth System Modelling. The structural components of the framework are (1) Service platform "Performance Analysis for Earth System Modelling", (2) Earth System Model Benchmark Suite containing climate application benchmarks of varying complexity, (3) Set of common performance metrics, (4) Measurement and evaluation tools, and (5) Central benchmark repository.

Project co-funded by the European Commission's Seventh Framework Programme (FP7; 2007-2013) under the grant agreement n°312979		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants including the Commission Services	
RE	Restricted to a group specified by the partners of the **IS-ENES2** project	
CO	Confidential, only for partners of the **IS-ENES2** project	

Table of contents

Executive Summary

The aim of the JRA2 is to guide and coordinate the development and distribution of an Earth System Model (ESM) Benchmark Suite and to provide a full spectrum of related services. The main challenge thereby is to prepare application benchmarks tailored to meet the needs of all potential users: climate model developers, vendors, third party software developers, model users, and staff of computing centres. The IS-ENES2 project provides an opportunity for a tight collaboration across different institutes, companies, and countries and allows for the practical realisation of the technical and organisational framework for development, distribution, and evaluation of ESM benchmarks.

This report proposes how the benchmarking framework can be set up and managed and describes the basic structure of the framework. The proposed approach leverages developments and experiences gained within existing initiatives like RAPS (Real Applications on Parallel Systems) and UEBAS (Unified European Application Benchmark Suite).

In the first project phase major steps towards realisation of the core elements of the framework have been made: the service platform 'Performance Analysis for Earth System Modelling' ***[https://verc.enes.org/pa]*** *has been installed and first packaged versions of ESM benchmarks are available for vendors and IS-ENES2 project partners. A detailed description of ESM Benchmark Suite will be the subject of the following deliverables.*

1. Introduction

The main objective of the JRA2 (Joint Research Activity 2) as part of the IS-ENES2 project is to establish a comprehensive framework for assembling and distributing of Earth System Model (ESM) Benchmark Suite as well as for collecting, processing and exchange of characteristic performance data. In general, the term 'benchmark framework' refers to a software system that provides interfaces and tools in order to simplify construction, definition, configuration, execution, and evaluation of benchmarks. Since very little experiences are available with Earth System Model benchmarks the development of a unified software framework seems to be impractical at this stage due to the high development, adaptation, and maintenance cost. Indeed, the design of a flexible and easy-to-use software framework requires a lot of effort and experience related to a deep analysis of workflows for each particular model in order to identify common tasks and implement generic solutions. Therefore, we will initially allow for different level of compliance in the set-up and handling of different ESM benchmarks.

In this report we propose how to set up an organisational framework that ensures the availability and usability of and access to climate applications benchmarks to

- Facilitate and improve benchmarking of the HPC systems for procurements as well as for performance tests after machine upgrades, compiler and library updates etc.
- Provide a more realistic estimate of sustained performance for climate research applications
- Provide vendors, manufacturers of CPUs and compiler builders a better way to assess performance characteristics of weather and climate research applications
- Increase the level of cooperation and interaction between climate modelling community, vendors, and software developers thus fostering co-design and innovation in hardware and software development, establishing common standards, and introduction of unified interfaces
- Compare the performance of different ESMs on different computing systems in order to develop an understanding of factors affecting performance of some/all ESMs
- Develop and test generic solutions for principal issues of ESMs that prohibit high scalability of the models
- Foster comparison and evaluation of different coupling strategies between multiple components of an ESM
- Prepare ESMs and model infrastructure for extreme scale computing characterized by massively parallel, heterogeneous and fault-prone environment

The definition of the benchmark framework leverages the best practices and developments within existing initiatives like Real Applications on Parallel Systems (RAPS) [1] and Unified European Application Benchmark Suite (UEABS) [2]. Therefore, the first step was to assess

the current state of usage of application benchmarks, in particular climate and weather benchmarks, and identify the opportunities for improvement and further development.
Founded in the early 90s the RAPS consortium set up a number of portable and well documented benchmark codes, compiled best practices on how to prepare a portable and 'bug free' benchmark [3] and also provides regular updates of benchmark codes and scripts. However, the applications mainly cover the area of Numerical Weather Prediction (NWP); no coupled ESMs are available so far. Furthermore, low attention was paid to the gathering of the available performance results derived from benchmark runs. Therefore, we should try to ensure the collection and availability of benchmarking results to foster comparison across different climate models and HPC systems. This requires a well-defined set of common performance metrics to make the comparison meaningful. For confidentiality and licensing reasons vendors have to contact individual RAPS members to get a specific benchmark. If possible, this procedure should be made as straightforward as possible.
A number of prerequisites have been defined for application benchmarks building UEABS. These include relevance to scientific community, public availability of codes and datasets or suitable open license agreement, portability, scalability, and active support by developers. The applications building the ESM benchmark suite should be evaluated against these criteria.

2. Components of the benchmarking framework

The proposed components of the framework are as follows:

- Service platform "Performance Benchmarks for Earth System Modelling"
- ESM benchmark suite containing applications of varying complexity
- Set of common performance metrics
- Light-weight performance measurement and validation tools
- Central benchmark repository

2.1 Service platform

The service platform "Performance Benchmarks for Earth System Modelling" is intended to be used for distribution of information on climate application benchmarks and long-term hosting and update of performance data for comparison across different HPC systems and Earth System Models. Furthermore, it should provide ways to interact for model developers, vendors, software developers, model users and staff of computing centers thus maximizing the benefits of the assembled ESM benchmark suite. A prototype version of the service platform has been installed as branch of the ENES Portal [https://verc.enes.org/pa] (Figure 1).

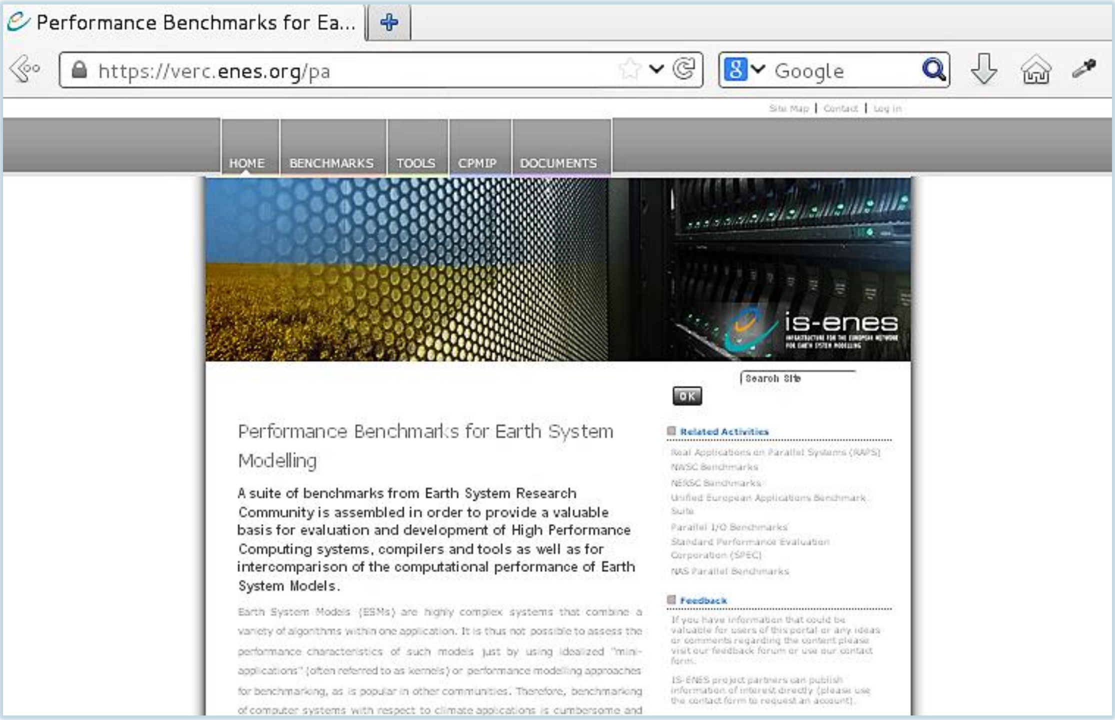

Figure 1: [https://verc.enes.org/pa] Screenshot of the prototype version of the service platform "Performance Benchmarks for Earth System Modelling".

2.2 ESM benchmark suite

The ESM benchmark suite is the core element of the framework. Due to the limited experience with coupled model benchmarks and very heterogeneous character of the suite members no unified software framework for different coupled models will be provided. However we will try to adopt the best practices defined for RAPS benchmarks with model specific benchmark setups.

Besides common application benchmark criteria like representativeness, availability, portability to different hardware architectures and operating systems, scalability, and support capabilities the following technical requirements should be met by the members of the ESM benchmark suite:

- Benchmark can be run from a self-contained set of source files, very simple build and run scripts, and input data
- Documentation on how to perform and evaluate a benchmark run and technical tuning guide are available
- Automatic documentation of the build and run environment (e.g. compiler and compiler options, include and linkage information, used MPI library, tuning variables for MPI etc.) is possible

- Automatic correctness check and reporting of benchmark results are possible
- Input data for different problem sizes (e.g. standard science run and high-resolution run) are available

The applications within benchmark suite are supposed to deal with different aspects of climate modelling. According to their complexity and intention the available benchmarks are divided into several categories:

- **Coupled Earth System Models** are the most challenging and advanced applications in climate research. They produce high compute and data workloads and stress almost all features of the computing systems: floating-point and integer performance, memory bandwidth, I/O sub-systems, network interconnects etc. The following state of the art models will be included in the benchmark suite: IPSL-CM5, MPI-ESM1, CESM-NEMO, EC-Earth.
- **Uncoupled models** allow for exclusion of coupler cost and ESM load imbalance from performance analysis. ICON, COSMO-CLM, ECHAM6, NEMO will be part of the benchmark suite.
- **Coupling technology benchmarks** will address the coupling challenges in ESMs by defining a suite of benchmarks based on simplified model components which capture the essence of the coupling while reducing scientific complexity. Benchmarks based on OASIS3-MCT, OpenPALM, and YAC coupler will be part of the benchmark suite. Further details on benchmark suite for evaluation of coupling strategies and description of test cases are available in [4].
- **Kernels** are performance sensitive key codes extracted from the NEMO and ICON models.
- **Parallel I/O benchmarks** aim to test asynchronous/parallel model output based on CDI and XIOS.
- **Fortran ISO benchmarks** provide test cases for Fortran2003 and Fortran2008 features.

Benchmarks for workflow steps related to data movement and processing are currently not a part of the suite.

2.3 Performance metrics

Efforts to define a set of metrics allowing for robust and objective performance analysis have been made within WP9/JRA1 and WP3/NA2 of the IS-ENES2 project. WP9/JRA1 activities focus on the assessment and intercomparison of real model performance. The following metrics have been introduced:

- Model resolution and number of grid points
- Complexity

- Simulated years per day
- Actual simulated years per day
- Core hours per simulated year
- Memory bloat
- ESM coupler cost
- ESM load imbalance
- Data output cost
- Data input coast
- Data intensity
- Workflow cost
- Parallelisation

Definitions of the metrics listed above are available at https://verc.enes.org/pa/cpmip/metrics. Details on how to collect them can be found in the IS-ENES2 deliverable D9.1 "HR ESM initial performance analysis" [5]. From our point of view, it is important to find ways how these metrics can be assembled on the basis of the ESM benchmarks, such that productive performance of ESMs can be inferred.

WP3/NA2 defined methodology for examination of the computational performance of parallel applications aimed at the identification of bottlenecks, inefficiencies, and optimisation potential. The following metrics proved to be useful for this purpose:

- Parallel speedup
- Execution time
- Parallel efficiency
- Karp-Flatt metric
- Serial fraction
- I/O time percentage
- Communication time
- Computing time
- Load balance
- FLOPS
- Instructions per cycle
- Cache hit rate
- Vectorised instructions

- Operational intensity
- Scaled speedup
- Iso-efficiency

For each metric a short description following the scheme "what"-"how"-"rationale" can be found in [6].

2.4 Measurement and evaluation tools

Most climate models provide their own solutions for granular run time measurements by use of wall-clock timers defined for performance relevant sections of code. However, to routinely gather performance metrics mentioned in section 2.3 a compelling need for a lightweight scalable profiling tool that does not significantly affect code performance, supports different programming paradigms, hardware architectures and operating systems is evident. A tool integrating the following capabilities is highly desirable:

- Automated and manual code instrumentation
- Usage of high-resolution timers
- Analysis of hardware performance counters
- MPI analysis
- OpenMP threads analysis
- Vector use analysis
- I/O analysis
- Report on memory usage

Below, open source profiling tools with low measurement overhead that cover some of the required capabilities are listed:

- SCT library [https://doc.redmine.dkrz.de/sct/html/index.html]
- mpiP [http://mpip.sourceforge.net/]
- GPTL [http://jmrosinski.github.io/GPTL/]
- DrHook [https://redmine.dkrz.de/collaboration/boards/2/topics/8]

The SCT (Simple Context Timer) library being developed by DKRZ within the IS-ENES2 project allows a limited subset of the desired performance metrics to be gathered and is recommended for use initially in conjunction with the benchmark suite. Future work, beyond IS-ENES2, will aim to address the development of tools providing more complete coverage.

Load balance between different model components in a coupled ESM is essential for the overall model performance. The LUCIA software [7] developed by CERFACS/CNRS for models coupled with OASIS coupler might be a helpful tool for analysis and minimization of the coupling load imbalance.

Numerically different results are expected due to usage of different hardware architectures, compilers, and compiler options. No tools for standardised correctness check of results are available currently. The following approaches have been adopted in NWP and climate benchmarks to check the validity of results:

- **Maximum error correctness check.** The maximum error is computed from saved reference norms and norms from performed benchmark run. The error is calculated for a set of selected dynamic and thermodynamic model variables at each time step and at each atmospheric level. The acceptable maximal error lies below 1% for a pre-defined number of time steps.
- **Check against valid range.** The values of selected integral model variables such as global annual mean of near surface temperature, sea ice area and volume in the Northern and Southern Hemisphere, global fraction of deserts etc. are compared to a predefined valid range. The benchmark developers can derive the valid range from ensemble minimum and maximum generated with a benchmark on an available machine. The basic assumption is that all valid extreme values occur at least once within ensembles of sufficient size. Alternatively, the valid range can be derived from corresponding long-term climate experiments.

Both approaches are based on providing model specific reference values assembled with a trustful configuration of a benchmark and are closely related.
The development of tools for benchmark correctness check is a future project task. The approaches to implementation need to be discussed among the project partners.

2.5 Central benchmark repository

The availability of a central benchmark repository might offer a number of benefits like automatic versioning of codes and scripts, transparent integration of bug fixes and code optimisations, uncomplicated code distribution etc. However, license agreements for different benchmarks and interests of vendors should not be violated.

3. Outlook

As noted in the previous section the ESM Benchmark Suite builds the central component of the proposed framework. A detailed description of the benchmark suite and comparison of performance on different hardware platforms will be subject of the following IS-ENES2 deliverables. The deliverable will document the available benchmarks in terms of scientific complexity, programming models, parallelisation and decomposition strategy, portability, execution, evaluation, scalability and performance.

References

[1] Real Applications on Parallel Systems (RAPS):
[https://redmine.dkrz.de/collaboration/projects/raps]

[2] Unified European Application Benchmark Suite (UEABS): [http://www.prace-project.eu/ueabs]

[3] Mozdzynski, G., 2012: RAPS Introduction. [https://verc.enes.org/pa/documents]

[4] Valcke, S., 2014: IS-ENES2 Milestone M10.1 “Benchmark definition for evaluation of coupling strategies”.

[5] Maisonnave, E., and U. Fladrich, 2014: IS-ENES2 Deliverable D9.1 “HR ESM Initial Performance Analysis”. [https://verc.enes.org/ISENES2/documents/deliverables]

[6] Epicoco, I., S. Mocavero, F. Macchia, and G. Aloisio, 2014: Performance Analysis. [https://redmine.dkrz.de/collaboration/documents/480]. Note: the document is currently accessible for registered project members only.

[7] Maisonnave, E., and A. Caubel, 2014: LUCIA, load balancing tool for OASIS coupled systems, Technical Report, TR/CMGC/14/63, URA CERFACS/CNRS No1875, France [http://pantar.cerfacs.fr/globc/publication/technicalreport/2014/lucia_documentation.pdf]