

Porting, Testing, Optimising, and Configuring the Community Earth System Model (CESM2) on ARCHER2

Technical Report for ARCHER2-eCSE03-6

eCSE ID:	ARCHER2-eCSE03-6
eCSE Title:	Porting, Testing, Optimising, and Configuring the Community Earth System Model (CESM2) on ARCHER2
PI Name:	Dr Ioana Colfescu (NCAS)
Author of this document:	
Name of technical staff on the project:	Dr Christopher C. Symonds (University of Leeds), Dr David H Case (University of Reading), Michael Bareford (EPCC)
List of names to acknowledge with the project (by default we will use the PI, Co-Is and Technical staff members):	Dr Ioana Colfescu (NCAS), Dr Mike Mineter (University of Edinburgh), Dr Massimo Bollasina (University of Edinburgh), Dr. Wuhu Feng (NCAS), Dr Mark Richardson (University of Leeds), Dr Maria Val Martin (University of Sheffield), Dr Grenville MS Lister (NCAS), Dr Christopher C. Symonds (University of Leeds), Dr David H Case (University of Reading), Michael Bareford (EPCC)
Project Start Date:	01/06/2021
Project End Date:	30/06/2022
Number of Project Months Funded:	8

Technical Report

Abstract

The Community Earth System Model (CESM2) is a fully-coupled global climate model developed at the US National Center for Atmospheric Research (NCAR) [1] that allows state-of-the-art simulations of the Earth's past, present, and future climate system. The availability of this model will provide scientists with tools to study these topics, and also to make comparisons with other models, which is an essential part of climate science. Beyond the scientific advantages, the model can be more easily configured compared to other models currently available on ARCHER2, such as the Unified Model [2], as all of the infrastructure can be maintained and run in a self-contained manner on one system.

The most current scientifically validated CESM2 version was released in June 2020 (CESM2.1.3), and this report details its porting, optimisation, testing and validation to ARCHER2. The key choices made during the model porting, specifically with regards to the software and hardware available on the system, are described. The model has been run in several configurations, which are the typical set-ups used by the climate research community, and the performance is characterised. These configurations have also been scientifically tested, and results are shown to satisfy validation requirements.

1. Introduction

CESM2 is an Earth System Model which can be configured and run in many different ways and resolutions, from global atmosphere-only to fully coupled. The model comprises different components, including CAM (atmosphere), CICE (sea ice), CISM (ice sheets), CTSM (land), MOSART (adaptive river transport), POP2 (ocean), RTM (river transport) and WW3 (waves), all of which are connected using the Common Infrastructure for Modelling the Earth (CIME). The model is a key tool to study climate changes and its causes, and due to the complexity of the code, it needs to be run on ARCHER2. CESM2 has been ported to many platforms, some of which are similar to ARCHER2.

The technical specification of ARCHER2 is: AMD EPYC2 processors (2x64 cores per node) with the Cray specific Linux, built using the Gnu compiler collection (GCC), and makes use of pre-built software libraries: `cray-hdf5-parallel`, `cray-netcdf-hdf5parallel`, `cray-parallel-netcdf`, `cray-libsci`, `cray-hdf5`, and `cray-netcdf`. Important work was done to inform/train/configure CESM2 of the batch scheduler (SLURM) queues specific to this service.

2. Background

UK Researchers interested in atmospheric science and earth system modelling use simulations to support their work. One such tool is CESM2. For a researcher new to the software it can be a daunting task to build, install, validate and then run their own simulations on a new machine. Typically prior to this eCSE, for a user to run CESM2 on ARCHER2 they would have to work through the NCAR guidance on "Porting" [3] to edit the

machine configuration files themselves before either carrying out extensive testing (pre-alpha, small ensemble bitwise comparisons and script regression tests). In some cases the work involved in this task would be quite onerous for a scientist given the large amount of technical details to get familiar with, including possibly large amounts of manual debugging, at the expense of time spent on the scientific problem to be investigated. This project allows UK researchers using ARCHER2 to confidently and reliably use CESM2 with minimum efforts.

3. Work packages

The work involved in the eCSE can be split into three work packages (WPs), these being porting, validation and dissemination. WP1 involves creating and testing the ARCHER2-specific configuration of CESM2, and running the recommended test suites. This task also included the identification and mitigation of code errors that prevented these test suites from properly running in some cases, and informing the code maintainers at NCAR about this. WP2 involves configuring and running a set of four key model set-ups, which are the typical configurations that scientists are likely to use on ARCHER2, and then validating the results from these defined test case runs to ensure the simulations are scientifically accurate. WP3 involves dissemination of the results from the first two WPs, including informing the UK CESM community of the availability of CESM2 on ARCHER2 and the ways the model can be accessed and used.

3.1 WP1: Porting and Testing

3.1.1 ARCHER2 specific configurations

In CESM2 the machine specific configurations that specify the programming environment, batch scheduler and compiler options are contained within XML files. Previously a user wanting to use CESM2 on ARCHER2 would create versions of these files containing the configuration for ARCHER2 and put them in the directory ``${HOME}/.cime``. These user specific configurations would generally be based on those used by colleagues or collaborators or, if that option is not available, then a user would use a template and spend effort to write their own configuration files. As one would expect, this approach is very prone to errors.

For machines known by the CESM2 software, these configurations are instead contained in central configuration XML files, located in the directory for the linking component of CESM2, CIME (Common Infrastructure for Modelling Earth). Machines in these configuration files are split into two categories - 'supported' machines which are checked for validity by NCAR staff and 'unsupported' machines which are included by collaborators from outside NCAR and are the responsibility of those collaborators to check for validity. ARCHER2 is now an 'unsupported' machine, as we have validated the output ourselves, but the machine details are included in the repository for CESM version 2.1.3.

The Intel compiler is used as standard on the NCAR machines, and as such it is the Intel compiler that is most often recommended for use with CESM. The Intel compiler is not available on ARCHER2 and is unlikely to be in the future, with the compilers available being

CCE by Cray, AOCC by AMD, and GCC from GNU. GCC has been tested previously on NCAR machines, [4], and CCE is known not to work with CESM to the extent that NCAR had advised not to use the Cray CCE. Later versions of CESM may address the issues of building with Cray CCE, however this is currently not a priority for NCAR and it was decided that the effort required to change the build system for an entirely new compiler is outside the scope of this eCSE. GCC has however been used on various other supported machines to run CESM and so it was decided that the porting and building of CESM2 on ARCHER2 would use GCC also. Due to the available compilers however, this is the only compiler for which an ARCHER2 configuration has been prepared.

As a result of the porting process, ARCHER2 is now included as an unsupported machine in the maintenance branch of CIME used by CESM2. As a result of these configurations, CESM2 is now set up to make use of the standard, short and serial queues, use the correct versions of the GNU compiler and pre-installed netcdf/hdf5 libraries, and build efficiently using compiler flags that have been verified as producing a stable run with a throughput rate similar to that obtained on NCAR machines. Preparing and validating these configurations was an iterative process which made use of the in-built testing suites to ensure that no unexpected errors were included in the configuration. This was important as, through testing, some errors in commonly used ARCHER2 configuration options were discovered which a user following the previous workflow may have overlooked.

3.1.2 Port verification through running the test suites

The changes necessary for a correct porting of CESM to a new machine can only be committed back to the NCAR repository when several blocks of testing have completed and results have been approved by NCAR. As stated earlier, CESM is made of interconnected individual components, and in CESM2 the component which links all of the modelling components, CIME, contains not only the machine specific configuration settings as detailed above but also the porting tests. There are three test suites contained within CIME: scripts regression, pre-alpha, and ensemble comparison.

The first test suite, named the scripts regression test, is a series of over 100 unit tests which increase in complexity, concluding with a full system test. It is very important to complete this test suite before any porting changes are brought into the NCAR repo. By running these tests a greater understanding of the configurations needed was achieved, including but not limited to configurations of the SLURM setup used on ARCHER2 which has differences to the setup used on other machines (including known issues with the way that environment variables are imported to a slurm environment), and also showing the importance of using the correct compiler version and compiler optimisation level. It was also discovered that some tests could only pass successfully with the gnu compiler if the branch of CIME used by CESM was updated, and as such a change of branch from the tagged version 5.6.32 to the maintenance branch `maint_5.6` is included as a step in the setup instructions. After the necessary changes were made all but one of the regression tests passed, with the exception being a test that failed during an intercomparison step. This failure however was deemed acceptable after consultation with staff at NCAR.

The second category of tests, named the 'cheyenne pre-alpha tests' is a series of 70 small test runs which test different aspects of the CESM model, looking at various components.

This suite is designed as the minimum number of tests that should be successfully completed by new versions of the code when run on the UCAR machine `cheyenne` prior to receiving an alpha tag. The running of these tests requires downloading around 2TB of input data. Initially over 60 of the 70 tests failed, and so a large amount of work was carried out to solve these problems. The errors encountered had many causes, and the investigation of these failed tests gave clarity on multiple underlying issues. Four specific issues are the importance of the NetCDF comparison program `cprnc`; the grid remapping software `ESMF` (Earth System Modelling Framework); issues caused by the configuration and importing of the netcdf libraries; and the need in some cases to change the decomposition and processor spacing to allow for higher memory availability per core. Another issue that seemed to occur with a surprising regularity was the corruption of downloaded input files which would not be caught by the md5 checksum process. As a result of finding and solving this issue for the tests, some advisory notes have been included in the user guide on the ARCHER2 docs pages [5] to advise users to check their input files for consistency, along with information about the types of errors that can indicate a corruption in the downloaded input files.

Another set of test failures seemed to be caused by coding errors in three components: atmosphere (CAM), land (CTSM) and waves (WW3). In all three components, there is a point where the conditionals use short-circuit logic which, when using the gnu compiler at optimisation levels below -O1 (as is done during testing), had the potential to create divide-by-zero errors. This behaviour has been previously reported in a github issue for the WW3-CESM component and a fix applied, however the fix was applied after the tagged release for CESM2. The file in which the error appears in CAM does not exist in the 2.2 release branch, only the 2.1 release branch, and the code in question could not be found in the files for 2.2. In the CTSM repo the problem code has been fixed as well in the master branch, but not in any of the current tagged releases (up to release 5.0.35). As such, these errors will not be a problem for future scientifically verified versions of CESM (such as ≥ 2.3). It should also be noted that as this behaviour only occurs when `DEBUG=TRUE` and running with the gnu compiler, it will not affect any production runs, and so where the tests ultimately passed once the necessary code changes were applied to the test version of the CESM code, these tests were classed as successful.

After all changes were made, a small number (~10) of failures remained. Some of these failures were due to the configuration of the tests themselves, such as tests which used a hardcoded file path to some data stored on cheyenne, or tests that made certain assumptions about the underlying architecture of the host machine. Other tests were confirmed by NCAR staff as being safe to disregard.

The final set of tests, named the CESM Ensemble Comparison Tests (CESM-ECT), evaluate the bitwise compatibility of runs by comparing a run against a known good output file. There are two ECTs, first the UF-ECT (ultra-fast ECT) which looks at a nine timestep run using just the atmospheric component, and secondly the POP-ECT, a twelve month run to look at issues in the ocean (POP) and sea ice (CICE) components. The process for running these tests requires the output files from each test to be uploaded to an online comparison tool. Both of these ensemble tests have passed completely, verifying bitwise compatibility of runs.

3.2 WP2: Configuring, Testing and Verifying Defined Test Cases

3.2.1 Configuration and testing of Test Cases

Four configurations, or compsets, of CESM2 have been run, chosen as typical model configurations used by the climate community. F2000 is an atmospheric-only model configuration, ETEST includes an atmosphere coupled to a simple (slab) ocean model, B1850 is a fully coupled model (without chemistry), and FXHIST has a well-resolved atmosphere up to the stratosphere which includes electrochemistry and requires the ESMF library [5]. In the F2000, ETEST and FXHIST compsets, the atmospheric model dominated the cost, whilst the ocean model also made a significant contribution to the cost of the coupled (B1850) configuration.

Both the atmospheric and ocean models scale well with cores, and with OpenMP threads. Figure 2 displays the simulated years per day for the B1850 case at a 1 degree (f09) resolution. The ocean is faster than the atmosphere, but not enormously so. Increasing the number of cores through MPI or OpenMP increases throughput, and whilst at a low PE count it is more efficient to spend the resources in MPI, OpenMP points overlap on the figure (to the eye) and allow larger totals to be run. The cost of the ocean component is reasonably significant to the total in this case, and can be run on separate nodes to concomitantly run both components: instructions for this are given in the notes for the ARCHER2 website [6].

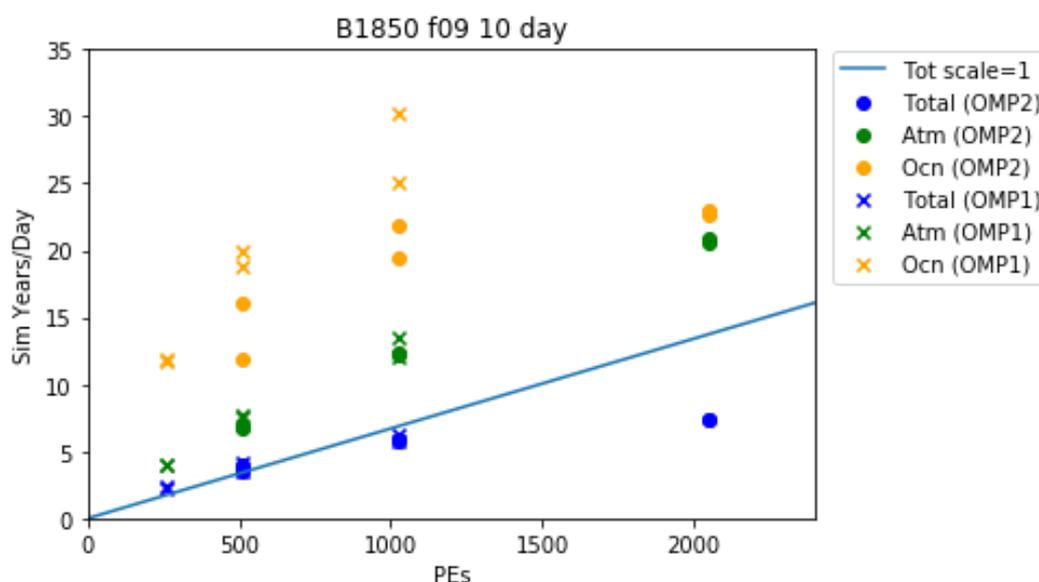


Figure 2: The scaling performance, in simulated years per day, of a 10 day B1850 simulation with a one degree grid, is plotted vs PEs. Total performance is shown in blue, with the atmospheric and ocean components in green and yellow respectively. The blue line shows ideal scaling of total performance. The crosses represent pure MPI configurations and there is some linearity for atmosphere only. Filled circles are configurations that have 2 OpenMP threads. The total performance assumes that model components are run consecutively, although slightly better performance can be obtained by running concomitantly (see text). Two data points have been calculated for each measurement, the difference between these gives a rough estimate of errors.

3.2.2 Validation of the defined test cases

Four specific simulation configurations were run for extended duration to confirm that the porting has been successful for their science.

The above-mentioned cases have been run at 1 degree atmospheric resolution for 30 years in the cases of F2000, ETEST and B1850, and at 2 degree resolution for 15 years for FXHIST. The archives are stored on JASMIN (partially so for FXHIST) and have been analysed for validity. This has involved a comparison with observations as well as with the output of similar simulations conducted with CESM at NCAR for many variables and multiple metrics, such as annual mean values, vertical profiles, zonal mean, annual cycle. An example of the output available from NCAR is in [6]. To provide some examples of the comparative analysis, Figure 3 displays the climatological (30-year) annual mean precipitation (mm/day) from GPCP observations (upper panel), CESM F-compset (middle panel), and CESM B-compset (bottom panel). The model shows a very good agreement with observations, and the performance is consistent with that of other CESM versions [6].

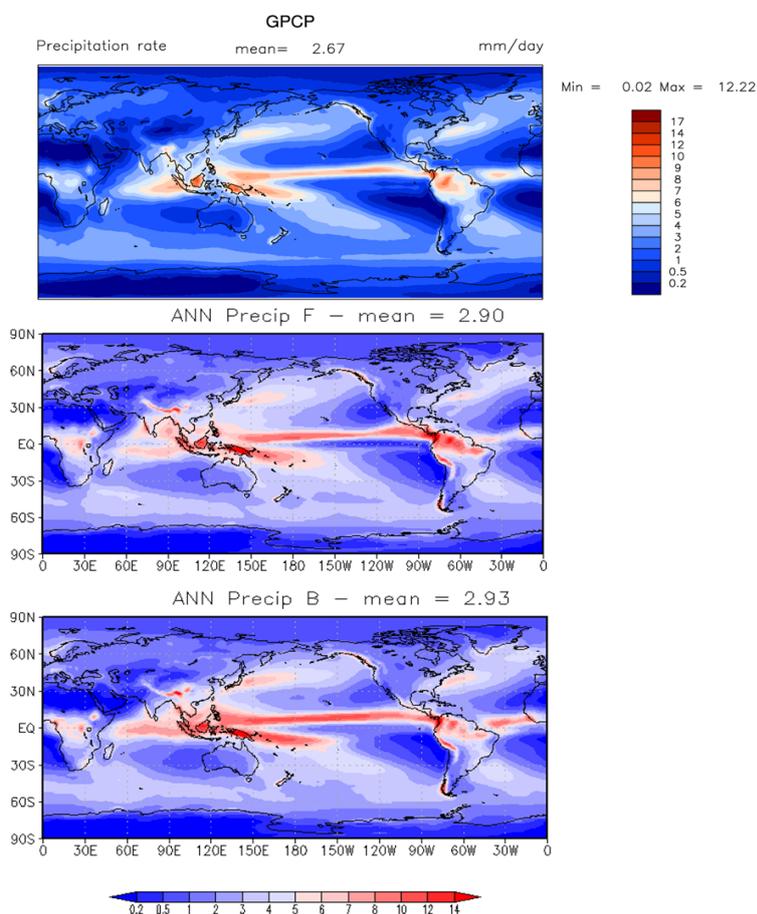


Figure 3: Climatological (30-year) annual mean precipitation (mm/day) from GPCP observations (upper panel), CESM F-compset (middle panel), and CESM B-compset (bottom panel).

Similarly, Figure 4 shows the annual cycle of the climatological (30-year) 200-hPa zonal-mean zonal wind (m/s) from MERRA Reanalysis (upper panel), CESM F-compset (middle panel), and CESM B-compset (bottom panel). This indicates that the model captures not only the spatial pattern of annual mean quantities, but also their variations within the seasonal cycle.

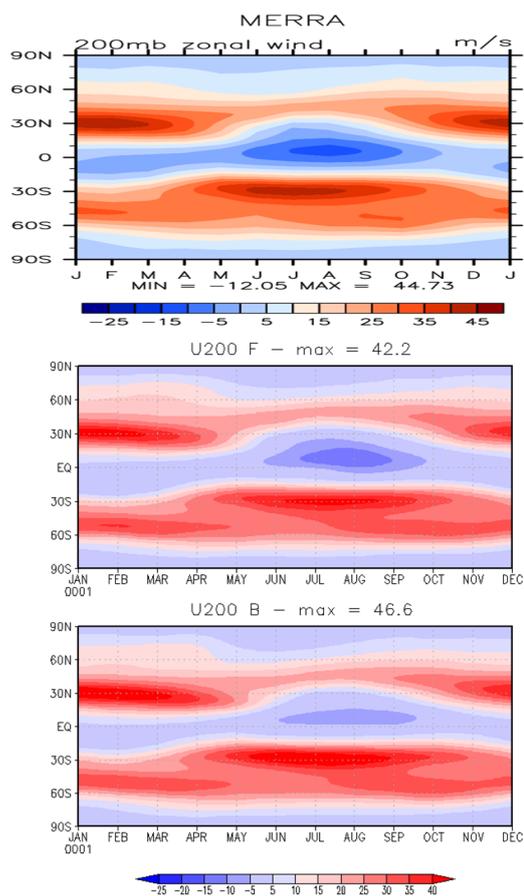


Figure 4: Annual cycle of the climatological (30-year) 200-hPa zonal-mean zonal wind (m/s) from MERRA Reanalysis (upper panel), CESM F-compset (middle panel), and CESM B-compset (bottom panel).

Furthermore, short runs (5 days) of the four cases have been performed by a different user (M Bareford). This check establishes that a new user is correctly running the simulations with the set-up scripts provided and following the guidelines given. The simulations run as anticipated.

3.3 WP3: Dissemination

The instructions to set up an installation of CESM2 and run a simple case have been incorporated into the ARCHER2 documentation website. Now when a user needs to build and install CESM2 on ARCHER2 they can follow the setup guidance provided here (

https://docs.archer2.ac.uk/research-software/cesm213_setup/), which will guide the user through the initial setup process including cloning the CESMs repository, building tools used by CESM and creating a personalised configuration file which will sit in the user's \$HOME directory. The user can then use the quickstart guide <https://docs.archer2.ac.uk/research-software/cesm> which will guide them through the steps required to set up and run a short run for a typical fully coupled simulation.

Further to this, the machine specific configuration of CESM2 for ARCHER2 has been included in the NCAR github repo for CIME, the interconnection component of CESM2, on the maintenance branch used by CESM2 (<https://github.com/ESMCI/cime/tree/maint-5.6>).

The project outcomes will be communicated to the CESM2 UK user mailing list as well as to the broader UK climate research community to raise awareness also to potential new users. Finally, the NCAS Training Committee is currently discussing organising a 3-day CESM2 workshop (NCAS-NCAR) tentatively for the spring of 2023. This will entail the joint participation of UK and US scientists.

4. Summary and concluding remarks

The project has achieved all the objectives and CESM2 is now a model running on ARCHER2.

There are four specific simulations that were used to satisfy UK scientists of the validity of the software port, and an independent user, who did not set up the port (MB), has successfully followed the online documentation to run these test cases for themselves.

The test suites were exercised and results were reported back to NCAR. The work was undertaken from a github “forked repo” and so it was straightforward to issue a “Pull Request (PR)” and the software developers at NCAR could merge the changes into a maintenance branch.

NCAS supported the proposal and provided personnel for the long simulations. CEMAC provided personnel to undertake the porting exercise.

The tagged maintenance branch of CESM2 was forked in github and changes were made to enable it to build, install and run on ARCHER2. Documentation has been created within the ARCHER2 github repository to explain to users how to build, install and run the software, along with instructions on setting up the configuration of the four test cases detailed above. They are processed and published along with all Archer2 documentation and appear in [7]

This success of this project is of great benefit for the UK climate research community, with a foreseeable user workshop in Spring 2023.

This work was funded under the embedded CSE programme of the ARCHER2 UK National Supercomputing Service (<http://www.archer2.ac.uk>).

- [1] <https://www.cesm.ucar.edu/>
- [2] <https://www.metoffice.gov.uk/research/approach/modelling-systems/unified-model>
- [3] http://esmci.github.io/cime/versions/maint-5.6/html/users_guide/porting-cime.html

- [4] <https://bb.cgd.ucar.edu/cesm/threads/porting-cesm2-13-to-archer2-memory-allocate-issue.5883/>
- [5] <https://www.cesm.ucar.edu/models/cesm2/config/2.1.3/compsets.html>

- [6] https://webext.cgd.ucar.edu/FAMIP/gpci_cam5.1_cosp_1d_001/atm/gpci_cam5.1_cosp_1d_001-obs/

- [7] <https://docs.archer2.ac.uk/research-software/cesm/>