Performance low resolution CESM1/CESM2 on Snellius

Michael Kliphuis

February 3, 2023

1 Introduction

We tested the performance of CESM1.0.5 and CESM2.1.3 on low degree grids on the dutch national supercomputer Snellius. This (informal) document shows the test results.

We tested two low degree grids:

0.9x1.25_gx1v6 (1° atm/ 1° ocn)
1.9x1.25_gx1v6 (2° atm/ 1° ocn)

In each test we used the so called "B1850" compset representing a fully coupled pre-industrial simulation in which all the components (see table below) are active.

version	**components**
CESM1.0.5	CPL LND ICE ATM (CAM4) OCN
CESM2.1.3	CPL LND ICE ATM (CAM6) OCN ROF GLC WAV ESP

CESM2.1.3 runs with the extra river (ROF), land-ice (GLC), ocean-wave (WAV) components as well as an external system processing (ESP) component that can be used for data assimilation (not in the tests). There is also a river component in CESM1.0.5 but this is embedded in the LND component.

In CESM2.1.3 it is also possible to switch on submodels of components like the Biochemistry submodel (MARBL) of the OCN component. We tested with MARBL on and off.

All tests were done with the Intel compiler and were optimized with flags -O2 and the Math Kernel Library mkl.

The test results will be used to determine which CESM version and resolution we will use for doing a 3000 year hosing experiment. Therefore the tables in the 'Results' section below also have a column with the corehour costs per 3000 modelyears.

2 Results

2.1 resolution: 0.9x1.25_gx1v6

The table below contain the performance test results for the resolution 0.9x1.25_gx1v6 (1° atm/ 1° ocn), for CESM1 as well as CESM2

resolution: 0.9x1.25_gx1v6 (1° **atm**/1° **ocn**)				
CESM1				
# cores	performance (modelyears/24h)	cost core-hours/modelyear	cost corehours/3000 modelyears	remarks
128	12.13	254	762.000	atmosphere CAM4 much more
768	38.39	481	1.443.000	expensive than ocean.
1280	45.61	674	2.022.000*	putting >1024 cores on atm
1792	49.99	861	2.583.000	gives communication overhead and performance loss
CESM2				
1024	6.24	3938	11.814.000	atmosphere CAM6 again much
1280	8.68	3539	10.617.000	more expensive than ocean.
1408	9.85	3432	10.296.000*	as well as CAM4! (see section
1792	9.41	4570	13.710.000	'conclusions'). putting >1152
2304	8.60	6432	19.296.000	cores on atm gives communication overhead and performance loss

Table 1: * Blue entry seems most optimal in terms of corehours and completion time

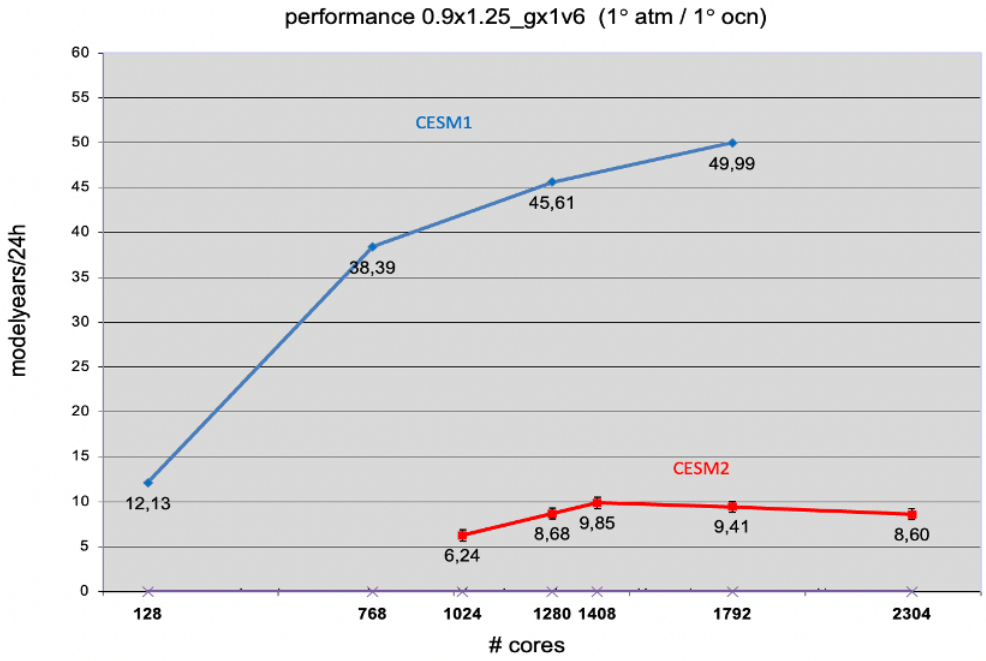

Figure 1: Performance 0.9x1.25_gx1v6 resolution of CESM on Snellius

2.2 resolution: 1.9x1.25_gx1v6

The table below contain the performance test results for the resolution 1.9x1.25_gx1v6 (2° atm/ 1° ocn), for CESM1 as well as CESM2

resolution: 1.9x1.25_gx1v6 (2° **atm**/1° **ocn**)				
CESM1				
# cores	performance (modelyears/24h)	cost core-hours/modelyear	cost corehours/3000 modelyears	remarks
768	57.56	321	963.000	atmosphere slightly less
1024	83.74	295	885.000	expensive than ocean but max
1536	93.81	393	1.179.000*	number of cores on atm with
1664	93.11	429	1.287.000	this resolution is 1024 (too few gridpoints)
CESM2				
768	19.42	949	2.847.000	atmosphere CAM6 much
1024	22.54	1091	3.273.000	more expensive than ocean.
1280	24.98	1230	3.690.000*	as well as CAM4! (see section 'conclusions')

Table 2: * Blue entry seems most optimal in terms of corehours and completion time

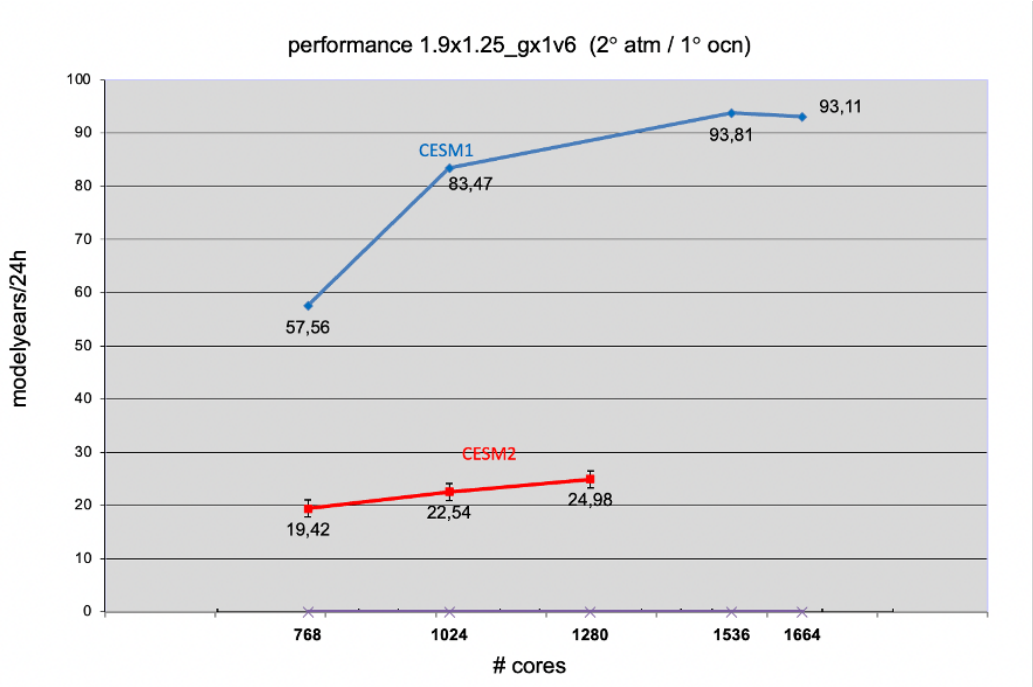

Figure 2: Performance 1.9x1.25_gx1v6 resolution of CESM on Snellius

NCAR published the performance of CESM (for both resolutions) on many different supercomputers in a so called timing table. The performance results for CESM1.0.5 can be found here. The ones for CESM2.1.3 can be found here. If we compare their results of a B1850 case then for CESM1.0.5 our performance is much better (but Snellius is a much newer machine) and we get a similar performance for CESM2.1.3.

3 Conclusions

CESM2 is roughly a factor 5 slower CESM1 for the f09_g16 resolution and about a factor 4 slower for the f19_g16 resolution.

The extra components ROF, GLC, WAV and ESP in CESM2.1.3 are not really causing the lower performance. They are on the same cores as the atmosphere is on and use relatively few computation time. Since the ocean model is not so expensive compared to the atmosphere (especially for the f09_g16 resolution) it is not even a problem to add MARBL. Since the atmosphere is the bottleneck and the ocean has to wait anyway we tested similar performances with MARBL on and off.

The lower CESM2 performance is mainly caused by the much more expensive atmosphere model CAM6 (instead of CAM4).

The figure below shows an overview of the evolution of CAM and the schemes that were used for 'Microphysics', 'Deep Convection' etc. The model became complexer and much more expensive in terms of cpu-time.

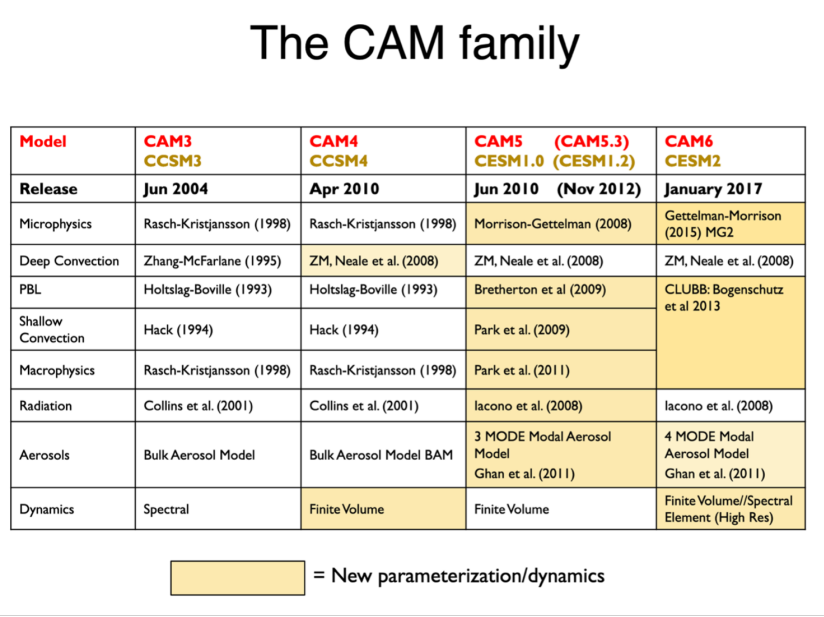

Figure 3: Evolution of CAM

If budget and time are no problem then we can use the 1° atm/ 1° ocn resolution of CESM. CESM2 will cost roughly 10 million cpu hours but it will take 300 wallclock days to complete the run. The CESM1 version of this resolution will cost about 2 million cpuhours and will take about 67 wallclock days to complete.

If we want to complete the run faster and spent less budget then it is better to use the 2° atm/ 1° ocn resolution of CESM. CESM2 will then cost roughly 3.6 million cpu hours and it will take 120 wallclock days to complete the run. The CESM1 version of this resolution will cost about 1.2 million cpuhours and will take about 33 wallclock days to complete.

Note that at IMAU we earlier ran the high resolution CESM1.0.4 with CAM5 instead of CAM4 (in order to run with a more sophisticated atmosphere). We could do this too with CESM1 but ofcourse this will decrease the performance.