

In this exercise, you will use CESM to compute the surface mass balance of the Greenland ice sheet. You will make a simple code modification to perform a crude global warming or cooling experiment.

Create a case

You will create a case that consists of an active land component (CLM, the Community Land Model), an active ice-sheet component (CISM, the Community Ice Sheet Model), an active river component (MOSART, the Model for Scale Adaptive River Transport), a data atmosphere component, and stub ocean, sea ice and wave components. The data atmosphere component is from GSWP3 (the Global Soil Wetness Project Phase 3). The land model will (among many other things) compute a surface mass balance that is used to force the ice sheet model.

```
> cd /glade/p/cesm/tutorial/cesm2_0_alpha07c/cime/scripts
> ./create_newcase --case ~/cases/i.day5.1 --mach cheyenne --res f19_g17 --
compset I1850Clm50SpG --run-unsupported
```

I1850Clm50SpG is the short name for
1850_DATM%GSWP3v1_CLM50%SP_SICE_SOCN_MOSART_CISM2%EVOLVE_SWAV.

The resolution f19_g17 implicitly sets the CISM grid to its default of a 4-km Greenland grid (if you'd like, you can use the equivalent alias, f19_g17_gl4, in order to be more explicit).

You need the '--run-unsupported' option here because CESM doesn't have any automated tests of this exact compset and resolution combination. If you'd like, try the above command without '--run-unsupported' and see what happens. In general, you should be wary of running configurations where you need to specify '--run-unsupported', because you're more likely to run into problems (obvious or subtle), but in this case I'll vouch for this being an okay configuration.

Some relevant information printed by create_newcase is:

- LND component is clm5.0 Satellite phenology:
 - o CLM5.0 gives us the latest suite of CLM parameterizations.
 - o Most simulations use active land biogeochemistry rather than satellite phenology (e.g., using the compset I1850Clm50BgcCropG rather than I1850Clm50SpG), but here we'll use satellite phenology so the model runs faster.
- GLC component is cism2:cism ice evolution turned on:
 - o CISM2 gives us the latest suite of CISM dynamics and parameterizations.
 - o In general, compsets with a "G" at the end (or near the end) of their name have an evolving Greenland ice sheet (e.g., B1850G, T1850G, and this I1850Clm50SpG); others use a static Greenland ice sheet that merely diagnoses down-scaled surface mass balance

Configure and build the model

```
> cd ~/cases/i.day5.1
```

Now setup and build the case. You will be making source code modifications, so will have to rebuild the code later. But let's start a build using the unmodified code, so that this can proceed while you work on the modifications. The rebuild with the modified code will then proceed much faster.

```
> ./case.setup
```

Remember to do this tutorial-specific thing:

```
> ~/compile_node.csh
```

And then run:

```
> ./case.build
```

Modify source code to perform a global warming / cooling experiment

Your IG case is configured to run with atmospheric data for the years 1901 – 1920. What if the surface air temperature were higher or lower than these prescribed data?

To find out, you can change a single line of code in CLM. One way to change a file is to make a copy of the file in one of your SourceMods directories and change it there, leaving the original version unchanged. (Here we need to use SourceMods because we cannot change code in the shared tutorial location. In general, using SourceMods is best for changes that just apply to one or two cases, whereas it's better to change code in-place – leveraging the version control system – when you're doing long-term developments.)

Since the initial build is taking up your terminal window, you will need to open a new terminal window to do the following.

The file you will modify is called atm2lndMod.F90, and it is part of the CLM source code. Type the following to copy it to the CLM SourceMods directory:

```
> cd ~/cases/i.day5.1
> cp /glade/p/cesm/tutorial/cesm2_0_alpha07c/components/clm/src/main/atm2lndMod.F90
SourceMods/src.clm/
```

Now go to the CLM SourceMods directory:

```
> cd SourceMods/src.clm
> ls
```

There will be a single file, "atm2lndMod.F90". When the code is built, any files that are in this directory will automatically replace files of the same name in the CLM source code directory.

You will edit the version of "atm2lndMod.F90" in the CLM SourceMods directory. Find this line of code:

```
tbot_c = tbot_g-lapse_rate*(hsurf_c-hsurf_g) ! sfc temp for column
```

This part of the code sets the surface air temperature, "tbot_c", for all columns in CLM. Recall that in each grid cell containing glacial ice, the glaciated area is divided into several columns, each with a different elevation. The surface air temperature at the mean gridcell elevation is "tbot_g". For each column, "tbot_g" is adjusted by adding a term proportional to the elevation difference between the mean atmospheric gridcell elevation "hsurf_g" and the local column elevation "hsurf_c".

You will change this line to something like this:

```
tbot_c = tbot_g-lapse_rate*(hsurf_c-hsurf_g) + 5.0 ! sfc temp for column,  
plus five degrees
```

In this way you will have a crude version of a global warming (or cooling) simulation. Later we will see how this temperature change affects the surface mass balance of the Greenland ice sheet.

Once you've made your code changes in SourceMods/src.clm, go back to the case directory and rebuild the model. **You will need to wait for the initial build to complete before doing this step; if the initial build is still underway, you could take a few minutes to explore the ice sheet code directory structure, as discussed below in the appendix.**

From the login window where you performed the first build, rerun:

```
> ./case.build
```

This build should go much faster than the initial build, because only the file you changed and a few other files that depend on it need to be rebuilt.

If the build fails, you will get an error message pointing you to a log file in another directory. If you look at the bottom of that log file, you will usually see what went wrong.

If the code builds successfully, you will see on your screen a message like this:

```
Building cesm with output to  
/glade/scratch/sacks/i.day5.1/bld/cesm.bldlog.170817-201321  
Time spent not building: 6.756262 sec  
Time spent building: 410.564775 sec
```

Add annual-average output of CLM's surface mass balance

Two useful diagnostic fields in CLM are QICE_FORC and QICE. QICE_FORC gives the surface mass balance that CLM sends to CISM, separated by elevation class. QICE is the same, but

without a separate elevation class dimension, so it can give you gridcell averages. Let's add an annual-average history file from CLM with both of these fields.

In `user_n1_clm`, add the following lines:

```
hist_nhtfrq = 0,-8760
hist_mfilt = 1,1
hist_fincl2 = 'QICE_FORC','QICE'
```

The `hist_nhtfrq` setting says to use monthly frequency for the first (default) CLM history file (0 is a special value meaning monthly frequency), and every 8760 hours for the second history file (negative values specify a frequency in hours; note that 8760 hours is 365 days - i.e., 1 year). The `hist_mfilt` setting says to put a single time slice in each history file. The `hist_fincl2` setting says to add these two fields to the second history file. (If we didn't make these changes, the default would be to get a single history file with `hist_nhtfrq = 0` and `hist_mfilt = 1`. So the initial 0 and 1 values are simply avoiding changing the defaults.)

To make sure there were no typos, run

```
> ./preview_namelists
```

and make sure that command completes successfully.

Submit the run

It's often a good idea to first do a short test run, to make sure you haven't introduced any errors that prevent the model from running. So, we'll first do the default 5-day run.

Shorten the wall-clock limit so that the run can get through the queue more quickly:

```
> ./xmlchange JOB_WALLCLOCK_TIME=0:10
```

Remember to do this tutorial-specific thing:

```
> ./xmlchange --subgroup case.run JOB_QUEUE=R1578619
> ./xmlchange --subgroup case.st_archive JOB_QUEUE=R1578619
```

Then submit the run:

```
> ./case.submit
> qstat -u $USER
```

When the run completes, confirm that it was successful:

```
> tail CaseStatus
```

You should see lines noting “case.run success” and “st_archive success” near the bottom of this file. If you instead see lines indicating failure, you will have to look through your log files to determine the cause of the failure, fix it, then resubmit this short test run.

You can now do a longer, 3-year run:

```
> ./xmlchange STOP_OPTION=nyears,STOP_N=3
```

Also, lengthen the wall-clock limit to allow enough time for this longer run:

```
> ./xmlchange JOB_WALLCLOCK_TIME=0:50
```

Then submit the run:

```
> ./case.submit
```

```
> qstat -u $USER
```

Wait for the run to finish

The run will take 20-30 min. to complete. While you’re waiting, you can take some time to look through these three things:

(1) The various namelist settings that can be used to control CISM’s evolution:

http://www.cesm.ucar.edu/models/cesm2.0/namelist/cism_nml.html

This is a bit out of date and will change somewhat for the final CESM2 release, but it gives a general sense of the possible options.

The settings under the headings “cism_config_options”, “cism_config_parameters” and “cism_config_ho_options” are particularly useful.

Note that the options under headings beginning with “cism_config” will appear in the cism.config file in your run directory; options in the other sections (which do not mention “cism_config”) will appear in cism_in. However, both sets of options are modified in the same way: via the *user_nl_cism* file in your case directory.

(2) The glc log file from your initial 5-day run. From your case directory, do the following:

```
> cd logs
```

```
> ls
```

```
> less glc.log.*
```

Scroll through the log file. You will see information about the CISM input parameters, followed by some model diagnostics that are written out during the run. If the run finished successfully, you will see this near the end:

```
> GLC: end of main integration loop
```

There are similar log files for other components (e.g., *lnd* and *atm*). The logfile with the *cesm* prefix combines diagnostics from each component.

(3) The CISM source code, as described in the Appendix.

View the results

Now we will look at some output from your IG run. When a run completes successfully, output is written to several archive directories. Let's first look at some CLM output:

```
> cd /glade/scratch/$USER/archive/i.day5.1/lnd/hist
> ls
```

You should see one file for each month of the run. Each of the monthly files contains a number of monthly average fields. This command will tell you about the contents of the first monthly file:

```
> module load netcdf
> ncdump -h i.day5.1.clm2.h0.0001-01.nc | less
```

In addition, recall that you set up the CLM namelist to also output supplementary history files with annual-average values of QICE and QICE_FORC. This is the surface mass balance of glaciated grid cells in units of mm/s. (To convert to the more useful units of m/yr, you would multiply by 3.16e4.) Values are positive where the ice is growing and negative where ice is melting.

Look at the annual average file for year 3:

```
> module load ncview
> ncview i.day5.1.clm2.h1.0004-01-01-00000.nc
```

Click on the box labeled 'QICE'. You should see a global plot of QICE on the global land grid. We recommend changing the color scheme: First, click on the button labeled "3gauss"; continue to click until it reads "default". Then click on the button labeled "range"; change the range to -2e-5 to 2e-5, then click "OK".

You should be able to see QICE in both Greenland and Antarctica. Note that the annual surface mass balance is positive for almost all of Antarctica. The mass balance is positive for most of Greenland, except for a patch in the southwest part of the ice sheet.

Recall that the surface mass balance is computed in CLM for multiple elevation classes and then downscaled to the ice sheet grid in CISM. To see the surface mass balance in each elevation class, click on the box labeled 'QICE'. Then click on the box in the row labeled "elevclass" to cycle through the elevation classes. Elevation class 0 is the "bare land" elevation class, which is used to trigger glacial inception. Greenland has valid data for all elevation classes (since this is needed to force CISM), whereas Antarctica only has valid data for elevation classes with non-zero areas in their gridcell (to reduce computational expense).

Let's see what the mass balance looks like after downscaling. First go to the directory with GLC history files:

```
> cd /glade/scratch/$USER/archive/i.day5.1/glc/hist
> ls
```

First, combine the available years, so that you can easily scroll through these years in ncview:

```
> module load nco
> ncr_cat i.day5.1.cism.h.0002-01-01-00000.nc i.day5.1.cism.h.0003-01-01-
00000.nc i.day5.1.cism.h.0004-01-01-00000.nc cism.h.allyears.nc
```

(You can ignore the warning about *scale_factor*.)

View the resulting file:

```
> ncview cism.h.allyears.nc
```

You will see several 2D fields. First, view the ice thickness, "thk". The thickness of the ice sheet does not change much during this short run.

Next, look at the surface mass balance, which in CISM is called "acab" (for accumulation/ablation). The units are m/yr. Change the color scheme: First, click on the button labeled "3gauss"; continue to click until it reads "default". Then click on the button labeled "range"; change the range to -2 to 2, then click "OK". Near the bottom of the panel, click on the current time box to advance by one year at a time. Watch how the surface mass balance changes.

This run uses initial conditions that are mostly spun-up, so there aren't huge changes over the first few years. For unaltered forcing, the pattern is as follows:

- In most of the ice sheet, the mass balance is small and positive (green shading).
- In the southeast, the mass balance is large and positive (red shading).

- In the southwest and in some coastal regions, the mass balance is negative (blue and purple shading).

These features are fairly realistic. The pattern would be improved if the surface mass balance were computed in CLM at higher (e.g., 1 degree) resolution.

Remember that you were running a climate change experiment. Depending on the size of the temperature change, you may or may not see similar features in your CISM history file.

Finally, create a difference map showing the changes in your run relative to a control run that was done without any changes to the source code:

```
> ncdiff cism.h.allyears.nc  
/glade/scratch/sacks/archive/i.day5.1.CONTROL/glc/hist/cism.h.allyears.nc  
diffs.nc
```

```
> ncview diffs.nc
```

Click on acab, change the color scheme as desired (for a 5-degree warming experiment, the "default" scheme with range -1 to 1 works well), and scroll through to year 3.

Positive values indicate more accumulation (or less ablation) in your experiment; negative values indicate less accumulation. If you performed a warming experiment, you should see mostly negative values around the margins of Greenland.

Appendix: Ice sheet model directory structure

Go back to the CESM code:

```
> cd /glade/p/cesm/tutorial/cesm2_0_alpha07c/components
> ls
```

The components directory contains subdirectories corresponding to each possible active CESM component.

Let's take a quick tour of "cism":

```
> cd cism
> ls
```

The source code resides in several subdirectories:

- glimmer-cism = The core CISM code; especially see the following subdirectories:
 - o libglissade: the glissade dynamical core, which is the default when running with CISM2
 - o libglide: the serial, shallow ice approximation-based glide dynamical core; this also contains some modules shared between glide and glissade
 - o libglimmer: lower-level, shared utilities
- source_glc = code needed to link CISM to the CESM coupler
- drivers = more code associated with the coupler

Feel free to explore these directories while waiting for the model to build or run.