

A Hitchhikers Guide to the Black Arts (of Earth system modelling)

Lesson #1: snowball Earth

Stuff to keep in mind:

- Models ARE NOT the 'real World' (it is going to be pretty obvious this is the case here).
- Don't believe what you read in Nature or Science.

Relevant reading:

Snowball Earth and ice-albedo feedback & hysteresis

Hoffman and Schrag [2002] (*Terra Nova* **14**, 129-155)

→ Snowball review.

Hyde et al. [2000] (*Nature* **405**, 425-429)

→ Model analysis of the inception of a snowball Earth and ice-albedo thresholds.

0. Readme

0.1 You will need to download a new *restart* file prior to embarking on the snowball Earth experiments. To fetch this: change to the `cgenie_output` directory, and type (or copy and paste carefully from the PDF ...):

```
$ wget http://www.seao2.info/cgenie/labs/Bristol.2016/LAB_1.SPIN.tar.gz
```

This downloads an archived/compressed copy of the experiment `LAB_1.SPIN` – effectively, just an experiment (*spin-up*) that has been run for 10,000 years for you. Extract the contents of this archive by typing:

```
$ tar xfv LAB_1.SPIN.tar.gz
```

A new experiment results directly will then appear as if you had just run the entire 10,000 year experiment yourself, and you could in fact have done so (remember to refresh the WinSCP window).

You'll then need to change directory back to `genie-main` to run the model.

1. Brrrrrrrrrrrr – it's chilly on ... snowball Earth!

1.0 To illustrate how 'easy' it can be to configure an Earth system model such as cGENIE and explore the behavior of the Earth system and its response to perturbation – you are going to induce an extreme cooling of climate and see what happens. Solar output was weaker during the late Neoproterozoic, a time when the Earth experienced a series (2 ish) of extreme glaciations. Thus, having a mild climate state to start with must have been dependent on sufficient CO₂ and/or CH₄ in the atmosphere and hence presumably highly elevated compared to the modern World, so sort of the opposite of the problem we have today ...

1.1 You are going to be running experiments similar to before and using the *restart*:

```
$ ./runmuffin.sh cgenie.eb_go_gs_ac_bg.p0650e.NONE LABS LAB_1.EXAMPLE 100  
LAB_1.SPIN
```

but rather than use the provided experiment configuration file `LAB_1.EXAMPLE`, why not get into the habit of creating new and uniquely named *user-config* files (no harder than copying it and renaming it!). If you keep using the same experiment name, the results will be over-written each time. Also, be especially careful not to have 2 (or more) experiments running simultaneously with exactly the same name as they try and over-write each other in a somewhat entertaining way.

Overall: your task in this exercise will be to determine the radiative forcing (or pCO₂ equivalent) threshold required to drive the climate system into a full ice-covered ocean (snowball Earth) state. (Read the *Hyde et al.* [2000] paper.)

Useful 2-D (netCDF—Panoply) variables to view are surface air temperature and sea-ice extent (and/or thickness). Ocean surface temperature and salinity can be viewed in the 3-D NetCDF results file (apologies for ocean temperature being in units of K ...). You can also save the data seasonally if you like – see Section 5.2.3 in the User Manual (your configuration has 48 time-steps per year for the BIOGEM module).

Time-series (ASCII `.res` files) are useful for providing simple mean indicators of global climate such as global ocean fractional sea-ice covered.

Note that the model configuration of an idealized super-continent, positioned symmetrically about the Equator, is pretty unrealistic. But the further you go back in time, the more uncertain it becomes as to exactly where and in what orientation the continents were. Sometimes modelers have to resort to somewhat idealized experiments if the uncertainties are too great. In addition, one can conduct sensitivity experiments to test whether the continental configuration is important to the results. For instance, *Hoffman and Schrag* [2002] discuss the potential importance of continental configuration, while the entire hypothesis of *Donnadieu et al.* [2004] rests on specific details of the continental configuration being realistic.

For this configuration, the solar constant is set weaker than modern to reflect the fact that the sun's output has increased with time and during the Neoproterozoic the solar constant would have been ca. 5% weaker. This is set by the parameter (hidden in the *base-config* file):

```
ma_genie_solar_constant= 1285.92
```

(compared to a modern value of 1368 W m⁻²)

Other questions to think about with regards to numerical modeling (and this experiment) are:

- *(Is the model configuration and experimental design 'realistic' ... ?)*
- *What is 'missing' in the model and what might the implications for your predictions and conclusions be? For example, there is no land-surface scheme (and hence no concept of 'snow') in this particular configuration.*
- *Are the simulations being run for sufficiently long? Why not if not (i.e., justify your choices of parameter values and experimental assumptions)? How might the results and conclusions be biased (if at all)?*
- *How would you test model predictions and your overall conclusions?*
- *How could the experimental design be improved?*
- *Etc.*

- 1.2 To search for the atmospheric CO₂ concentration (or rather, radiative forcing equivalent) that would lead to a ‘snowball Earth’ state in the Neoproterozoic and answer the question:

‘How low does CO₂ have to be to trigger a ‘snowball’?’

you are going to edit the file that controls the specific details of the experiment. This is the *user-config* file. From the `genie-userconfigs/LABS` directory, open one of the snowball experiments in the SciTE text editor. At the top of the file you should see something like:

```
#
#
# --- CLIMATE -----
#
...
# scaling for atmospheric CO2 radiative forcing, relative to 278 ppm
ea_radfor_scl_co2=20.0
```

Each line that is not commented out (i.e., no #) contains a parameter of the format:

PARAMETER=VALUE

The value of each parameter can be edited to form a new experiment. (Additional parameter value specifications can also be added, or existing ones deleted.) In this example, the line:

```
ea_radfor_scl_co2=20.0
```

specifies a radiative forcing of climate by CO₂ equivalent to ×20 modern (20×278 = 2560 ppm).

Note that CO₂ is not being explicitly modeled in this experiment, but the long-wave radiative forcing associated with a specified concentration of CO₂ (as a ratio to modern concentrations) is being set instead.

Edit the value of `ea_radfor_scl_co2` (lower or higher) and save the file. Re-run the experiment to see whether sea-ice extent is approaching a new steady state. You may want to try even longer simulations if it becomes clear that the model is still far from steady-state. You can judge how close to equilibrium things have got by following (and/or plotting) the evolution of e.g., global surface air temperature or sea-ice extent (both time-series files). Note that you *might* want to run the experiment longer than 100 years ...

HINT: By submitting the experiments to the cluster will allow you to run all these experiments simultaneously.

- 1.3 For each experiment you want to be assessing how far towards the Equator the sea-ice limit encroaches through some of the time-series and time-slice files or even the on-screen summary lines (assuming running interactively rather than via a job submission to the cluster queue). Informative time-series variables include (but not necessarily be limited to: atmospheric temperature and sea-ice cover. (Sea-ice thickness, on account of the simple physics in the model, low resolution and long time-step, can fluctuate a little in area and volume at times.)

For the time-slice data: atmospheric and ocean surface temperature and sea-ice extent (2-D biogem NetCDF file) may be informative.

HINT: Be careful with the default ‘auto-scaling’ feature in Panoply. At near complete sea-ice cover, you may find Panoply scaling min and max sea-ice between 99.1 and 99.9% or something – it always tries to maximize color contrasts by default and will be at all helpful in this case.

In answering the question think about what an appropriate degree of accuracy might be for your experiments. Just because computer models generally calculate to around 16 significant places of precision, does not mean you have 16 significant figures of accuracy (or realism). For instance – how many significant figures is the solar constant quoted to and what do you think is the uncertainty in this? Harder to judge is how the assumed (incorrect) continental configuration creates additional uncertainty, or the physics assumed in the ocean or sea-ice, or lack of snow on land ...

- 1.4 Once you are happy about the controls on the snowball threshold try and answer:

How high does the (CO₂) radiative forcing have to be in order to escape from a snowball?

If you run the model with an appropriate radiative forcing to create a snowball, you can use that experiment as a *restart* and be able to carry out a series of experiments with increasing radiative forcing, all starting from the snowball state you have created. Defining the radiative forcing / climate path going out of a snowball would complete the hysteresis loop of *Hyde et al.* [2000]. Note that a good restart is one for which the experiment did not sit long in the snowball state before finishing. You can fine-tune the number of years the experiment is run for to achieve this.

Overall: think critically about the model configuration, the experimental design, and the nature of the scientific question (based on your background reading of snowball Earth). Some of the exploration/testing suggestions (above) may not necessarily give substantially different results. Such a finding would be as valid and interesting as determining an important dependence of a certain assumption, and would for instance indicate that the associated paleo uncertainties are not critical to model assessment of the question.

Always be prepared to justify all your choices for experimental design and model settings, e.g., range of radiative forcing assessed, continental configuration(s), solar forcing, use of *re-starts* (if any), run duration, etc. etc. etc. etc.

2. Further ideas

2.0 Feedback loop analysis following *Hyde et al. [2000]*

In order to quantify the snowball Earth hysteresis loop in cGENIE as per Figure 2 in *Hyde et al. [2000]* you will need to extract from the model ‘meaningful’ measures of climate (e.g., global surface air temperature, fractional sea-ice coverage) as a function of CO₂ multiples, CO₂ concentration, or (better) radiative forcing. For the latter, in cGENIE, the radiative forcing for a doubling of CO₂ is set at: 5.77 W m⁻². See: *Myhre et al. [1998]* (*Geophys. Res. Lett.* **25**, 2715–2718) and/or *IPCC [2007]* for more on what radiative forcing is and how it is related to a relative change in CO₂ concentration. Also, for making a comparison with *Hyde et al. [2000]* for going into the snowball, note that they plot the change in radiative with a ‘cooling’ as positive (a bit daft). Their baseline radiative forcing state (an anomaly of 0 W m⁻²) you might assume is equivalent to 278 ppm and hence ~130 ppm is an approximately halving of CO₂ and hence creates ~5 W m⁻² of cooling. (You might prefer to plot the radiative forcing change as warming being positive, which makes rather more sense ...)

For coming out of a snowball, because the CO₂ and hence radiative forcing threshold is so high compared to going in, you may want to be creative in the plotting (assuming attempting to combine both thresholds into a single plot) and, for instance, one might break the scale between the low radiative forcing interval spanning going in and the high one spanning coming out.

Another example is as per Figure 3,4 in *Stone and Yao [2004]* (*Clim. Dyn.* **22**, 815–822) (although here it is the solar constant rather than long-wave radiation forcing that is being varied).

2.1 Continental configuration

It was mentioned earlier that the position of the continents is an area of modelling uncertainty and might be important. You can test for this. Four alternative *base-configs* are provided which each define a different continental configuration:

1. `cgenie.eb_go_gs_ac_bg.wopol1.NONE` – **a single polar super-continent**, with an ocean resolution of 36×36 with 8 vertical levels. (Note potential ‘l’ and 1’1 confusion in ‘wopol1’.)
2. `cgenie.eb_go_gs_ac_bg.wopol2.NONE` – **one continent at each pole**, with an ocean resolution of 36×36 with 8 vertical levels.
3. `cgenie.eb_go_gs_ac_bg.woreq1.NONE` – **a single Equatorially-centred super-continent**, with an ocean resolution of 36×36 with 8 vertical levels.
4. `cgenie.eb_go_gs_ac_bg.woreq2.NONE` – **two continents straddling the Equator**, with an ocean resolution of 36×36 with 8 vertical levels.

You can use the same *user-config* file (`LAB_1.EXAMPLE`) as before as an experiment template and any of the alternative configurations can be run very similarly to as per before, i.e.:

```
$ ./runmuffin.sh cgenie.eb_go_gs_ac_bg.xxxxx.NONE LABS LAB_1.EXAMPLE 100
```

Note that you are using a different *base-config* file name: `cgenie.eb_go_gs_ac_bg.xxxxx.NONE` here compared to Lab #1, where `xxxxx` is one of: `wopol1`, `wopol2`, `woreq1`, or `woreq2`.

Also note that no *restarts* are provided for these configurations. You may (or may not) want to create some (you will need to judge for yourselves how long to run the *restart* experiments for to achieve as close to steady-state as you think is ‘sufficient’). Remember that *restarts* are just ‘normal’ experiments that have already been run.

Be careful that when changing from one *base-config* to another, the model re-compiles. Simply running the new configuration briefly is sufficient to ensure this. Experiments can then be safely submitted to a cluster queue. i.e. do not try and submit an experiment using a different *base-config* straight to the cluster queue without having run it (or a short version of the experiment you want) first (to ensure the model is re-compiled). This is also good practice – checking that a new sort of experiment and/or model configuration works as you intend and without hiccups.

2.2 Geothermal heat input

Finally, cGENIE will fairly happily build up sea-ice, apparently without limit (with the remaining wet ocean becoming progressively colder and more saline). In the real world, one might expect some

sort of limit to the maximum thickness achieved as the heat diffusion across a progressively greater thickness of sea-ice approaches the heat input at the bottom of the ocean from geothermal energy. Different modes of ocean circulation are also possible if one considers heating from the bottom as well as cooling (and brine rejection) from the top and which might affect the entry into or exit from a snowball state.

In the experimental setup you have been given, a geothermal heat input is specified in the ocean circulation module via the following parameters

```
bg_ctrl_force_GOLDSTEInTS=.TRUE.  
bg_par_Fgeothermal=100.0E-3
```

The first enables the temperature and/or salinity of the ocean to be modified by processes other than redistribution via ocean circulation and exchange with sea-ice and atmosphere (i.e. to impose external forcings). The second sets the geothermal flux in units of $W m^{-2}$. (Note that in the Neoproterozoic, the geothermal heat flux could have been somewhat higher than modern. How much? A question for Google...)

An appropriate research question might be to determine in radiative forcing vs. geothermal space (and requiring a 2D grid of parameter combinations to be created and submitted to the cluster): the equilibrium sea-ice thickness and region in which a snowball solution is not possible. However, more simply and suitable to a short workshop: **How much of a difference to the estimated entry and exit thresholds of radiative forcing, does the inclusion of a geothermal input make?** e.g what happens if you set it to zero? What about 10 times modern (or more, although *extreme* seafloor heating can cause numerical instability)?