

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

PART #0001: Ocean circulation and Atlantic overturning stability

Stuff to keep in mind:

- Nothing at all – keep your mind completely empty and let the wonderful truths of GENIE permeate your entire being.



Relevant reading (and references therein):

Atlantic circulation and stability in GENIE:

Hargreaves et al. [2004] (Climate Dynamics, Volume 23, Issue 7 - 8, Dec 2004, Pages 745 – 760)

→ Simple assessment of the likelihood of AMOC collapse.

Marsh et al. [2004] (Climate Dynamics, Volume 23, Issue 7 - 8, Dec 2004, Pages 761 – 777)

→ Characterization of thresholds of AMOC collapse.

Singaraye et al. [2008] (GRL 35, doi:10.1029/2008GL034074)

→ Role of changing ocean circulation in atmospheric radiocarbon variability during the Younger Dryas.

Chikamoto et al. [2008] (JGR 113, doi:10.1029/2007JG000666)

→ Deep-sea CaCO₃ sediment and atmospheric CO₂ response to AMOC shutdown.

Miscellaneous (model) Atlantic circulation and stability:

Rahmstorf et al. [2006] (In: Encyclopedia of Quaternary Sciences, Edited by S. A. Elias. Elsevier, Amsterdam)

→ Provides the background to the Atlantic Meridional Overturning Circulation and hypothesized hysteresis.

IPCC [2007] (e.g., Section 10.3.4)

→ Future predictions of AMOC strength.

Schmittner [2005] (Nature 434, 628– 633)

→ Impacts on marine ecosystems and carbon cycling.

Obata [2007] (J. Clim. 20, 5962–5976)

→ Climate-carbon cycle model response to freshwater discharge.

0. Readme

- 0.1 You will need to download a new *restart* file prior to embarking on the experiments with modern ocean circulation.

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/UoB.2013/EXAMPLE.worjh2.PO4Fe.SPIN.tar.gz
```

This downloads an archived/compressed copy of the experiment `EXAMPLE.worjh2.PO4Fe.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 EXAMPLE.worjh2.PO4Fe.SPIN.tar.gz
```

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

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

1. Geochemical forcing of cGENIE (“poking the climate beast” [Wally Broecker])

- 1.0 The ocean biogeochemistry module (*BIOGEM*) in cGENIE provides a framework for applying time- and spatially-variable ‘*forcings*’ of the Earth system – fluxes or restored-to boundary conditions that can be prescribed for any gas, dissolved substance (including temperature and salinity), or particulate matter. Examples include freshwater input (== a negative salinity flux *forcing*) of the North Atlantic to alter ocean circulation, fossil fuel CO₂ emissions to the atmosphere (== a CO₂ gas flux *forcing*), or aeolian iron supply to the surface ocean (a 2-D dust flux *forcing*). For example: view the *user-config* file: LAB.0001.colorinjection – you will see the following lines (under the heading: ‘# --- FORCINGS ---’)

```
bg_par_forcing_name="pyyyyz_Fred"  
bg_par_force_point_i=22  
bg_par_force_point_j=33  
bg_par_force_point_k=8  
bg_par_ocn_force_scale_val_48=0.0
```

The first line points cGENIE to a directory located in `cgenie.muffin/genie-forcings` that contains a set of files that define what geochemical property is going to be altered plus information about how the magnitude of the forcing changes with time.

There are then three lines (`bg_par_force_point_i=20, ..`) that specify the location in the ocean of the geochemical forcing is going to be applied. The point sources are specified in (*i, j, k*) coordinates, which in this case is (22, 33, 08). For the ocean model resolution we are using, the grid is 36×36×16, longitude (*i*) is counted from left-to-right (1 to 36); latitude (*j*) is counted from bottom-to-top (1 to 36); level depth (*k*) is counted from downwards top-to-bottom (16 down to 1). Thus, (22, 33, 08) is a release of tracer in the North Atlantic, a little south of Greenland, and intermediate depth (level = 8 out of 16). Refer to the Figure overleaf for how the horizontal and vertical grid is specified. Note that Panoply appears to ‘count’ the ocean layers in the *opposite* direction to the way in which the ocean model is actually counting them – the correct definition is with ‘1’ being very deepest level possible.

Finally, there is a scaling parameter (`bg_par_ocn_force_scale_val_48`) which modifies the magnitude of the flux to be applied (in units of mol yr⁻¹).

(Section 4 in the *User Manual* describes the original and most flexible provision for applying time-dependent *forcings*.)

- 1.2 You are going to run a brief experiment in which you will be injecting a conservative ‘dye’ tracer in the ocean. BIOGEM has two tracers defined for this purpose – ‘blue’ and ‘red’. Open the *user-config* file: LAB.0001.colorinjection and edit the parameter controlling the flux of red dye to read:

```
bg_par_ocn_force_scale_val_48=1.0E12
```

which specifies a flux of 1.0×10¹² (mol yr⁻¹) rather than zero as given.

The *base-config* you will be using is different from previously:

```
cgenie.eb_go_gs_ac_bg.worjh2.rb
```

 – this specifies a 16 vertical levels ocean and also includes seasonality of solar insolation.

- 1.3 Run the model for ... whatever, 20 years will do. Use the *restart* experiment you have just downloaded to start from:

```
$ ./runmuffin.sh cgenie.eb_go_gs_ac_bg.worjh2.rb LABS  
LAB.0001.colorinjection 20 EXAMPLE.worjh2.PO4Fe.SPIN
```

View the results – how the Red tracer distribution evolves – in the *time-slice* files (full ocean (/atmosphere) properties saved in the netCDF format (.nc) files). You can follow the progress of the dye (and hence diagnose the properties of ocean circulation in the model) by plotting vertical and/or horizontal slices that go through (or near) the cell location in which you inject the dye tracer in the 3D netCDF file. You can also view the tracer distributions in terms of a water-column integrated tracer inventory (netCDF variable name: `ocn_int_colr`; long name: `colr water-column integrated tracer inventory`) in the 2D netCDF output. (See: *Sabine et al. [2004]* for the use of water column integrals in the context of the distribution of anthropogenic CO₂ uptake and storage.) Changes in tracer inventory with time can be tracked in the time-series file `biogem_series_ocn_colr.res`.

You can also plot the overturning circulation from the 2D netCDF file – variable `phys_opsi == global overturning streamfunction`, `phys_opsia == overturning in the Atlantic` to provide a visualization of the large-scale ocean circulation that drives tracer movement.

Spend a little while altering the flux (`bg_par_ocn_force_scale_val_48`) and/or location (`bg_par_force_point_i`, `bg_par_force_point_j`, `bg_par_force_point_k`) of tracer input. Note how you can use ‘tracers’ to help diagnose (and better understand) the circulation of the ocean.

- 1.4 An interesting (honest!) and illustrative exercise is to use the dye tracer to pick out the path taken by Mediterranean Intermediate Water. Despite the low resolution of the cGENIE ocean circulation model component and the highly restricted representation of the Mediterranean, the model does predict a salty Mediterranean as a consequence of P-E in this basin (and its catchments) being negative and this water makes its way out in the subsurface into the Atlantic.

Simply specify a dye injection somewhere in the Mediterranean (be careful with the restricted depth of the Mediterranean – if you inject too deeply (into the crust!) then you will not see anything (refer to the figure for the depth level (`k`) number of the maximum depth of the water column in each location), and it is better to inject it relatively close to the opening of the gateway (try some different locations and see which ones produce a reasonably instructive tracing of Mediterranean outflow). Run for e.g. 20 or 50 years (from the provided *spin-up*). Then:

(1) View the dye-tagged plume of Mediterranean Intermediate Water by plotting a `lat-lon` slice (from the 3D netCDF file). This will give you the depth of the plume. How does this compare with salinity observations (salinity observations and appropriate global datasets can be found on the web with a little patience)? You can also view the water-column integrated distribution (2D netCDF).

(2) Try viewing the plume via a `lat-depth` slice. Refer to the figure to determine the ‘`i`’ value up the Atlantic that will just graze the edge of what passes for Spain at this low model resolution. Which direction does it head after exiting the Mediterranean? Is this ‘realistic’?

- 1.5 Instead of adding a dye tracer, you could add fresh water to assess the sensitivity of the Atlantic Meridional Overturning Circulation (AMOC) to collapse in a classic ‘hosing’ experiment.

The *user-config* file for this is called: `LAB.0001.hosing`. The default (`i, j`) location of the flux input is the same (as the dye tracer), but now the injection at the surface (level: `k=16`). Note that the forcing of the salinity tracer is negative (freshwater = negative salinity compared to sea-water)! To orientate you in freshwater forcing space: `bg_par_ocn_force_scale_val_2=-2.0E17` should be sufficient to make ‘stuff happen’ and quickly. **BUT**, this is a pretty extreme flux (see overleaf for a rough conversion between salinity forcing units (`mol yr-1`) and fresh water flux (in `m3 s-1` or Sv). Much more than this and the model may crash or at the very least, you’ll be left with a large freshwater pond in the North Atlantic ... (See later for some exciting discussion on units!)

To run the model for e.g. 20 years using the same *restart*:

```
$ ./runmuffin.sh cgenie.eb_go_gs_ac_bg.worjh2.rb LABS  
LAB.0001.hosing 20 EXAMPLE.worjh2.PO4Fe.SPIN
```

20 years should be long enough to see a collapse start to occur, but you might want to run the model for longer (and it can be submitted as a job, of course). Running for longer will also allow you to have a smaller, less extreme (and maybe more realistic) freshwater input flux.

The most obvious property of the Earth system to follow is the Atlantic overturning strength (`biogem_series_misc_opsi.res`). The AMOC stream-function (in `fields_biogem_2d.nc` 2-D time-slice netCDF results file, field: `phys_opsia`) is also illustrative. You can also try and identify the salinity **anomaly** (see below) due to freshwater input in the 3D salinity tracer field.

There are also important impacts on surface air temperatures and maybe sea-ice extent (in `fields_biogem_2d.nc`). Note the importance of the AMOC in transporting heat to the N Atlantic region (the film *the Day After Tomorrow* was not entirely inaccurate in this particular respect). Be aware of the possibility of climate impacts far from the location of fresh water forcing. Look out for any significant-looking impacts on sea-ice extent, etc.

Note that as the model is running rather *s l o w e r* than in the snowball configuration, you might want to think carefully of making use of cluster queuing possibilities (i.e. running multiple experiments at once in the background).

THOUGHTS FOR FURTHER (INDEPENDENT) INVESTIGATION:

What is the largest freshwater flux that can be sustained without ‘collapsing’ the AMOC? Is there a ‘threshold’ (‘tipping point’) of freshwater input, beyond which the AMOC rapidly decreases in strength?

Is the precise location of the freshwater input important (i.e., try tipping it in somewhere else)? What would you expect to see in the paleo (e.g., ice core) record of both hemispheres if such a shutdown occurred in the past?

Are any other major regions of deep water formation (where are they) sensitive to freshwater perturbation and what are the consequences (could it happen in the future)?

1.6a To more easily assess some of these impacts (and for other sorts of analysis) it is possible to create an **anomaly** (difference) map in Panoply:

1. First open a dataset, e.g., `atm_temp` (surface air temperature) in the **2D** netCDF file. You can either double-click the variable name, or, with the variable name highlighted, click the ‘Create Plot’ icon.

2. Now, with the `atm_temp` still selected (and the first plot window still open), click on the ‘Combine Plot’ icon. A dialogue box will appear and ask you to select a plot to combine the new one with. Make sure the name of your first plot window is selected/highlighted. Click ‘Combine’.

3. You now have a plot window that by default it is showing you the difference between two identical (in time) slices. The two different slices are labeled `Array 1` (LH side) and `Array 2` (RH side).

Keep one array (`Array 1`) fixed to the initial (year 1 (centered on 0.5)) and vary the year in the second array (`Array 2`). Note that you can select in Panoply whether `Array 1 - Array 2` is plotted, or `Array 2 - Array 1`, or various proportional or relative differences.

Note that you can switch off the auto-scaling feature (`Always fit to data`) and center the scale so that no change is white, with positive deviations = red and negative = blue by clicking on `Center on 0` (an often used convention in climate field plotting).

1.6b You can also plot ocean current fields which is sort-of fun and maybe even informative(!):

1. In the **3D** netCDF file, the three components of ocean velocity are represented by the variables: `ocean velocity - u` (Eastwards), `ocean velocity - v` (Northwards), and `ocean velocity - w` (upwards). 2. Open up `velocity - u`. Chose ‘lon-lat’.

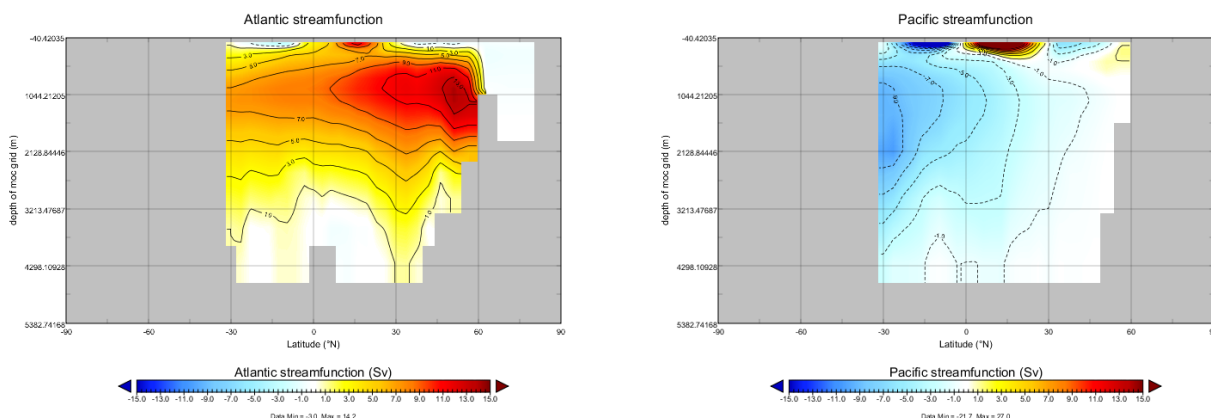
2. Select/highlight `velocity - v`. and click on the ‘Combine Plot’ icon (as per before).

3. Rather than a difference map, which is what you get by default, i.e. ‘Array 1 - Array 2’ – from the drop-down menu (next to the ‘Interpolate’ button) select ‘Vector Magnitude’.

4. You should have a color contoured (or not if you prefer plotting without contouring on) map of ocean current speed, with velocity vectors (direction and magnitude) overlain. You’ll need to re-scale the velocity vectors to properly see them – from the ‘Contours and Vectors’ tab – change the ‘Scale Length’ to e.g. 0.1. When fresh-water hosing – look out for impacts on the N. Atlantic current system associated with the AMOC.

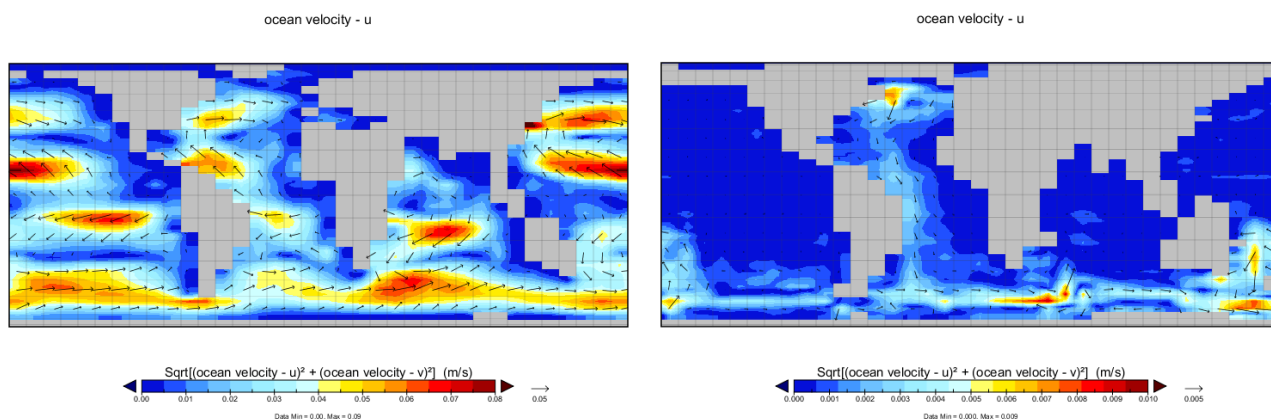
5. You can repeat this for deeper depth levels in the ocean – e.g. between about 1500 and 2000 m is a good place to go looking for the Western boundary current (and AMOC return) in the model (such as it exists at this low resolution) but you’ll need to re-scale the velocity vectors again (e.g. to 0.01 to less).

1.6c Example plots of the overturning streamfunction (2D netCDF file). (e.g. for Atlantic: netCDF parameter name: `phys_opsia`, long-name: Atlantic streamfunction).



Note that autoscaling has been turned off and the min and max plotting limits set manually. By convention, streamfunctions are plotted with their scale symmetrical around zero, giving red and 'warm' colors for positive value and clockwise overturning, and blues and 'cold' colors for negative values and anti-clockwise overturning. Also shown is the same for the Pacific. (The plots have been tart-ed up by overlaying solid contours plus contour labels.)

1.6d Example plots of ocean current fields (3D netCDF file). Again scaling has been set manually to create an easy-to-interpret axis scale. On the left is the surface field, and on the right an intermediate depth (illustrating what approximates the Deep Western Boundary current in the model in the Atlantic).



1.6e Finally, a brief note on units ... the freshwater forcing is implemented as negative salinity, just to really screw with your mind. The generic internal cGENIE model units for the forcing end up as $\text{PSU kg}^{-1} \text{ yr}^{-1}$. Which sort of does not make much sense ...

Start, by thinking of a value of `bg_par_ocn_force_scale_val_2` of -34.9 as equivalent to taking all the salt out of 1 kg of freshwater (since the mean global salinity is 34.9 PSU). Or equivalently, since the ocean volume is fixed, an applied forcing value of -34.9 is equivalent to adding 1 kg of freshwater to a (surface) box. So, a value of `bg_par_ocn_force_scale_val_2` of -3.49×10^4 ($-3.49\text{E}04$) would be a flux of $1 \text{ m}^3 \text{ yr}^{-1}$ (1000 kg m^{-3}) of freshwater.

So, in the example earlier (`bg_par_ocn_force_scale_val_2 = -1.0\text{E}18`), the freshwater flux is $1.0 \times 10^{18} / 3.49 \times 10^4 = 2.8653 \times 10^{13} \text{ m}^3 \text{ yr}^{-1}$.

The literature invariably gives freshwater fluxes in units of Sv ($10^6 \text{ m}^3 \text{ s}^{-1}$). So in the example, the freshwater flux is: $9.0797 \times 10^5 \text{ m}^3 \text{ s}^{-1}$ ($365.25 \times 24 \times 3600 = 31557600 \text{ s yr}^{-1}$). Or 0.9 Sv. Read the literature ... but generally, fluxes of ca. 0.05 Sv and larger (and to quite specific places) are applied in models in order to induce a collapse of the AMOC.