

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

PART #0100 : Fossil fuel CO₂ release and 'ocean acidification'

Relevant reading (and references therein):

Kleypas, J.A., Feely, R.A., Fabry, V.J., Langdon, C., Sabine C.L. and Robbins, L.L. (2006). Impacts of ocean acidification on coral reefs and other marine calcifiers: a guide for future research. Report of a workshop held 18–20 April 2005, St Petersburg, FL, sponsored by NSF, NOAA, and the US Geological Survey, 1-88.

www.ucar.edu/communications/Final_acidification.pdf

Orr, J.C., K. Caldeira, V. Fabry, J.-P. Gattuso, P. Haugan, P. Lehodey, S. Pantoja, H.-O. Pörtner, U. Riebesell, T. Trull, M. Hood, E. Urban, and W. Broadgate (2009) Research Priorities for Ocean Acidification, report from the Second Symposium on the Ocean in a High-CO₂ World, Monaco, October 6-9, 2008, convened by SCOR, UNESCO-IOC, IAEA, and IGBP, 25 pp.


<http://ioc3.unesco.org/oanet/index.html>

Royal Society (2005) Ocean acidification due to increasing atmospheric carbon dioxide. Policy document 12/05 Royal Society, London.

<http://royalsociety.org/WorkArea/DownloadAsset.aspx?id=5709>

Turley, C., Findlay, H. S., Mangi, S., Ridgwell, A. and Schimdt, D. N., CO₂ and ocean acidification in Marine Climate Change Ecosystem Linkages Report Card 2009. (Eds. Baxter JM, Buckley PJ and Frost MT), Online science reviews, 25pp (2009).

<http://www.mccip.org.uk/elr/acidification/>



0. ReadMe

0.0 [none]

1. Exploring the consequences of fossil fuel CO₂ emissions

- 1.0 For the next experiment(s) you can chuck CO₂ into the atmosphere, just for the hell of it. As much as you want! Apparently, humans are actually doing this now. Imagine that!

The *user-config* for cGENIE, LAB.0100.CO2emissions is provided and configured with climate being responsive to CO₂ (i.e., it takes account of CO₂-climate feedbacks):

```
# set CO2-climate feedback  
ea_36=y
```

as well as having a rate of calcification by plankton at the surface ocean that is responsive to ocean acidification and saturation state (i.e., it takes into account CO₂-calcification feedbacks, which will additionally interact with climate – see *Ridgwell et al.* [2007b, 2009] from <http://www.seao2.info/pubs.html>). Anything could happen!!!

In the *user-config*, a release of CO₂ to the atmosphere is prescribed, which by default is set for 1000 PgC over an interval of a single year. (Releasing CO₂ just over a single year is obviously rather unrealistic, but represents a useful idealized experiment for assessing the time-scale(s) of fossil fuel CO₂ uptake by the ocean.) Additional netCDF output has also been prescribed (bg_par_data_save_level=10) so that more information relevant to assessing ocean acidification is saved.

- 1.1 Run the experiment for e.g., 20 (or more if you like) years, starting from the *re-start* used previously, i.e.:

```
$ ./runmuffin.sh cgenie.eb_go_gs_ac_bg.worjh2.BASEFe LABS  
LAB.0100.CO2emissions 20 EXAMPLE.worjh2.PO4Fe.SPIN
```

and view the run-time output, particularly atmospheric CO₂ (which decays away after the first year as fossil fuel CO₂ is progressively taken up by the ocean), ocean surface temperature (SST), sea-ice extent, and Atlantic Meridional Overturning ('AMO') strength (and/or other variables). Viewing the time-series results file `biogem_series_fexport_CaCO3.res` will show how global carbonate production responds to the ensuing ocean acidification, which itself is recorded in the file: `biogem_series_misc_surpH.res` for mean global surface pH, or `biogem_series_carb_sur_ohm_cal.res` for surface saturation with respect to calcite, with a variety of other carbonate parameters also outputted. The calcification response is encoded in the model and described in *Ridgwell et al.* [2007a,b] (see: <http://pubs.seao2.org>) and may or may not reflect the real World.

In the 3-D netCDF time-slice file, while ocean pH is a particularly (but not exciting) relevant field to consider together with calcite and aragonite saturation – note that ocean surface waters in which aragonite becomes under-saturated ($\Omega < 1.0$) is regarded as a critical threshold for organisms making aragonite shells and skeletons and spells TROUBLE for some poor calcifying marine organism somewhere. Temperature is also highly relevant to marine ecosystems under future global change as well as sea-ice extent.

Refer to the 'output supplement' for a summary of the key ocean acidification variables in the model.

- 1.2 Because 'accidents can happen' and the global environmental changes induced by the massive fossil fuel CO₂ release can obscure mistakes made in the experiment configuration (parameter values) and/or the *re-start* used, you are strongly advised to first (or in parallel, as a job submitted to the cluster) set a control experiment going:

```
$ ./runmuffin.sh cgenie.eb_go_gs_ac_bg.worjh2.BASEFe LABS  
LAB.0100.CONTROL 20 EXAMPLE.worjh2.PO4Fe.SPIN
```

If everything is OK, atmospheric pCO₂ (and climate) should be stable and there should be little (or no) drift in any of the output variables.

Here – the *user-config* defining the control experiment (LAB.0100.CONTROL) is identical to that for the actual experiment itself (LAB.0100.CO2emissions) with the exception of the scaling of the CO₂ emissions that have been set to zero, via the parameter described below (Section 1.3).

It is good practice (i.e. always do it!) to always run a control experiment for each different type of experiments – i.e. only one control experiment is needed for multiple different CO₂ emissions experiment, but a new control would be advisable for a different set of freshwater hosing experiments.

- 1.3 You can easily modify the experimental design to release more/less CO₂ very much as you did for the red dye tracer. In the *user-config* file, the lines:

```
bg_par_atm_force_scale_val_3=8.3333e+016
```

scales the CO₂ flux. The scaling values given to you produce a CO₂ release of 1000 PgC yr⁻¹ for just a single year (cGENIE uses units of mol yr⁻¹ for *forcings*, with 1 PgC being equal to 8.333×10¹³ mol C) compared to current emissions are about 8 PgC yr⁻¹. Altering the value assigned to *bg_par_atm_force_scale_val_3* in the *user-config* file gives you a quick and simple immediate control over emissions rate. (The parameter: *bg_par_atm_force_scale_val_4*=-27.0 specifies the carbon isotopic composition of fossil fuel carbon and can be ignored.)

You can also adjust the emissions to have a time-varying rate by editing the file:

```
biogem_force_flux_atm_pCO2_sig.dat
```

which can be found in the directory:

```
cgenie.muffin/genie_forcings/worjh2_FeMahowald2006_FpCO2_Fp13CO2
```

The format of this file is:

```
-START-OF-DATA-
    0.0  1.0
    1.0  1.0
    1.0  0.0
999999.9  0.0
-END-OF-DATA-
```

and defines an emission of 1 mol C per year over the first 1 year of the model experiment (between year 0.0 and 1.0), but which in the example *user-config* is then scaled by a value of 8.333×10¹⁶ (by the parameter *bg_par_atm_force_scale_val_3*) to give a total of 1000 PgC yr⁻¹. (Year 999999.9 has no special meaning and is simply just way in the future ...)

Pause ... and note briefly how the final CO₂ flux is arrived at. cGENIE calculates it by multiplying the value in the *forcing* file (1.0) by a modifying parameter in the *user config* file (8.3333e+016). The total flux is hence: 1.0 × 8.333×10¹⁶ = 8.333×10¹⁶ mol CO₂ yr⁻¹.

Equally, we could have had 8.3333e+013 in the *forcing* file, which is the units conversion from mol C to PgC, and then had a value of 1000.0 as the value of *bg_par_atm_force_scale_val_3*, i.e. a total flux of: 8.333×10¹³ × 1000.0 = 8.333×10¹⁶ mol CO₂. Doing it this alternative way around is convenient because it 'hides' the units conversion to mol C in the *forcing* file, allowing a much simple value in units of PgC yr⁻¹ to be entered in the *user-config* file. It is entirely up to you which way around to do it – they'll multiply up to the same value. (If you screw up and multiply 8.3333e+013 and 8.3333e+016 as the total flux, you'll soon know it as you cook the Earth ...)

- 1.4 You can control the shape of the emissions profile as well as its magnitude. Between the start and end 'tags' in the text file, the data is arranged into 2 columns: the first contains a series of tie-points for defining the timing of changes in emissions, and the 2nd column contains flux information (units of mol yr⁻¹). At each time-step of the model the CO₂ flux is interpolated between these time points.

The purpose of:

```
0.0  1.0
1.0  1.0
1.0  0.0
```

then specifies a uniform flux of 1.0 (scaled in the user-config to be 8.3333e+016) over the first full year of the model run, followed by a sharp turn-off to zero flux at the end of first year (and remaining zero thereafter). To extend the period of emissions – for example:

```
0.0  1.0
10.0 1.0
10.0  0.0
```

would result in a uniform flux lasting 10 years (i.e. once scaled by the parameter in the *user-config* – 1,000 PgC yr⁻¹ over 10 years – 10,000 PgC total emissions). In contrast;

```
0.0 0.0  
10.0 1.0
```

would result in a linear ramp, starting from zero at the start of year 0.0 to 1.0 mol yr⁻¹ (scaled to 1,000 PgC yr⁻¹ at year 10.0 (a total CO₂ emission of 1000×10×0.5 = 5,000 PgC over 10 years).

- 1.5 By editing (and saving) the flux (and/or timing information) you can control the CO₂ emissions trajectory and total fossil fuel burn. Explore different CO₂ release assumptions and note their impact on climate and ocean biogeochemistry. Much more realistic and appropriate to our current global experimenting is a lower rate (order of 10 or 20 PgC yr⁻¹) released over a longer interval (order 100 years). Because the experiments are getting longer to run in real time ... remember to make appropriate use of the cluster queuing facility – i.e. think about whether you want to sit around starting at the screen for 15 minutes waiting for a new line of numbers appear – if not: submit to the cluster queue.

Possible questions you might have (or be persuaded to have) and investigate by creating new model experiments (aka Earth system ‘play time’) are listed below (GREEN). But before playing, think about setting the experiment detailed in Section 1.6 going on the cluster queue (as it takes about an hour).

What is the maximum total CO₂ release that can be made without inducing aragonite under-saturation at the ocean surface anywhere? How important is the time-scale of emissions in determining this? For total emissions above this: where in the ocean does the surface first become under-saturated? How large would the emissions have to be in order to induce under-saturation at the surface in the tropics (home to socio-economically important reef systems). These are questions that are addressed with simple CO₂ release experiments in ocean carbon cycle models and everyone seems to get a GRL paper out of it each and every time!

How important are CO₂-climate feedbacks in amplifying or diminishing future climate and ocean carbonate chemistry changes and on the atmospheric pCO₂ value itself – i.e. is the same atmospheric pCO₂ value reached with and without climate feedback (and surface warming) – if not, why? You can investigate this by contrasting an experiment made including CO₂-climate feedback with one made without. The CO₂-climate feedback can be turned off by the setting: `ea_36=n`.

Also: How much CO₂ emissions does it take to significantly ‘collapse’ the AMOC and over what time-scale? (Or alternatively: what is the atmospheric pCO₂ threshold for collapse?) If the AMOC weakens or collapses ... why in the absence of any freshwater perturbation does this happen? (Plotting appropriate environmental anomalies between the CO₂ release experiment and the control might help.)

- 1.6 Historical and future (SRES) emissions scenarios can also be prescribed explicitly. A historical emissions *forcing* (technically: a prescribed concentration profile of pCO₂ and other anthropogenic gases) can be specified by adding/substituting the following lines to the *user-config*:

```
bg_par_forcing_name='worjh2_FeMahowald2006_historical2010'
```

Now, no additional scaling is needed because the *forcing* specification directly follows the observed change in atmospheric concentration with time (in units of atm CO₂) and the line containing the scaling parameter:

```
bg_par_atm_force_scale_time_3=1.0
```

should be deleted (or commented out with a # at the start of the line).

An additional line is needed in the *user-config* because the historical pCO₂ transient starts in the 1700s (for which a nominal date of 1765 is often used) rather than year zero. For example, to start from year 1765, the start year parameter must be set:

```
bg_par_misc_t_start=1765.0
```

Because the start year has changed, it is convenient to specify save points that are consistent with the historical period, e.g.:

```
bg_par_infile_slice_name='save_timeslice_historicalfuture.dat'  
bg_par_infile_sig_name='save_timeseries_historicalfuture.dat'
```

A *user-config* with these changes is provided for your convenience (or to double-check you were following it all) – LAB.0100.historical. A suitable experiment would then be one run for 245 years so that it reaches year 2010 (having started from year 1765):

```
$ ./runmuffin.sh cgenie.eb_go_gs_ac_bg.worjh2.BASEFe LABS  
LAB.0100.historical 245 EXAMPLE.worjh2.PO4Fe.SPIN
```

WARNING! Ignore the 'WARNING's at the start – these are simply telling you that more *tracer forcings* have been specified than you have selected tracers for in the *base-config* (cgenie.eb_go_gs_ac_bg.worjh2.BASEFe). (A different *base-config* with additional selected tracers could have been specified to make use of other historical changes in atmospheric composition, such as of radiocarbon (¹⁴C) and CFCs.) Also: from year 1765 onwards, changes in atmospheric CO₂ only rise very slowly initially. Don't expect to see anything happen in 10 seconds flat because relatively few people and countries in the 1800s could be bothered to burn much more than a little local coal. You could potentially start your experiment at year 1850, changing the value of `bg_par_misc_t_start` and specifying shorter experiment duration.

- 1.7 Given that there is observationally-based information on the distribution of anthropogenic CO₂ taken up by the ocean (e.g. Sabine et al. [2004]) and having run a historical transient experiment with the model driven by observed increases in atmospheric pCO₂, you are in a position to critically evaluate the model's ability (or lack of) to represent the future-critical process of oceanic fossil fuel CO₂ uptake and transport by large scale ocean circulation.

In the 2D netCDF output, there is a variable for the water column integrated inventory of DIC – equivalent to the Sabine map except you will need to subtract the preindustrial background of DIC first, i.e. to create a DIC anomaly map representing only the added fossil fuel CO₂ component of ocean DIC. The data in the Sabine paper clusters around 1994. A time-slice centered on this year (1994.5) has been configured in the model exactly for this purpose. Your baseline state can either be from prior to CO₂ emissions commencing at any significant rate (e.g. 1750.5) or (better), from a control experiment. Note that similar comparisons could be (and are regularly) made with other tracers such as CFCs, which provide additional insights into the patterns and time-scales of trace gas update and ocean circulation. (See: Cao et al. [2009])

Observational data, re-gridded to the cGENIE grid and in netCDF format can be downloaded from the 'usual place' (<http://www.seao2.org/mucgenie.html>) under 'Observations on cGENIE grid'. You could for instance, compare horizontal or vertical slices (3D netCDF) and create difference (anomaly) maps. Somewhat more representative of the entire ocean is to compare (or calculate difference maps) of zonal average profiles. Unfortunately, the observations are not in the form of water column integrals and hence you cannot create difference maps of model as per the Sabine paper ... unless you are MATLAB-friendly and you use the 3D BIOGEM MATLAB plotting scripts (genie-matlab) whose use is somewhat described in the cGENIE user-manual. Examples of MATLAB plotting of the model vs. observed anthropogenic anomaly are show in the Figure.

- 1.8 Finally, and the closest to being slightly interesting of all: rather than applying highly idealized pulses (or other artificial shapes) of CO₂ emissions, IPCC SRES emissions scenarios can be used to make future projections by running on from the end of the historical transient. An example forcing of this sort is provided and can be selected by changing the name of the forcing selection parameter (`bg_par_forcing_name`) to:

```
worjh2_FeMahowald2006_FpCO2_Fp13CO2_A2_02180PgC
```

which gives you the IPCC 'A2' scenario, which has been extended beyond year 2010 in this case to give a total cumulative fossil fuel burn of 2018 PgC (e.g. Ridgwell and Schmidt [2010]).

As this *forcing* has units of PgC yr⁻¹ in its time-series file, you will need to add a scaling parameter to the *user-config* file to turn into mol yr⁻¹, i.e.

```
bg_par_atm_force_scale_val_3=8.3333e+013
```

(remembering that the final flux *forcing* is obtained by multiplying the value in the forcing by any scaling value set in the *user-config*).

For complete 'realism' you will need to run this experiment starting from the end of the historical transient experiment (Section 1.5), e.g.

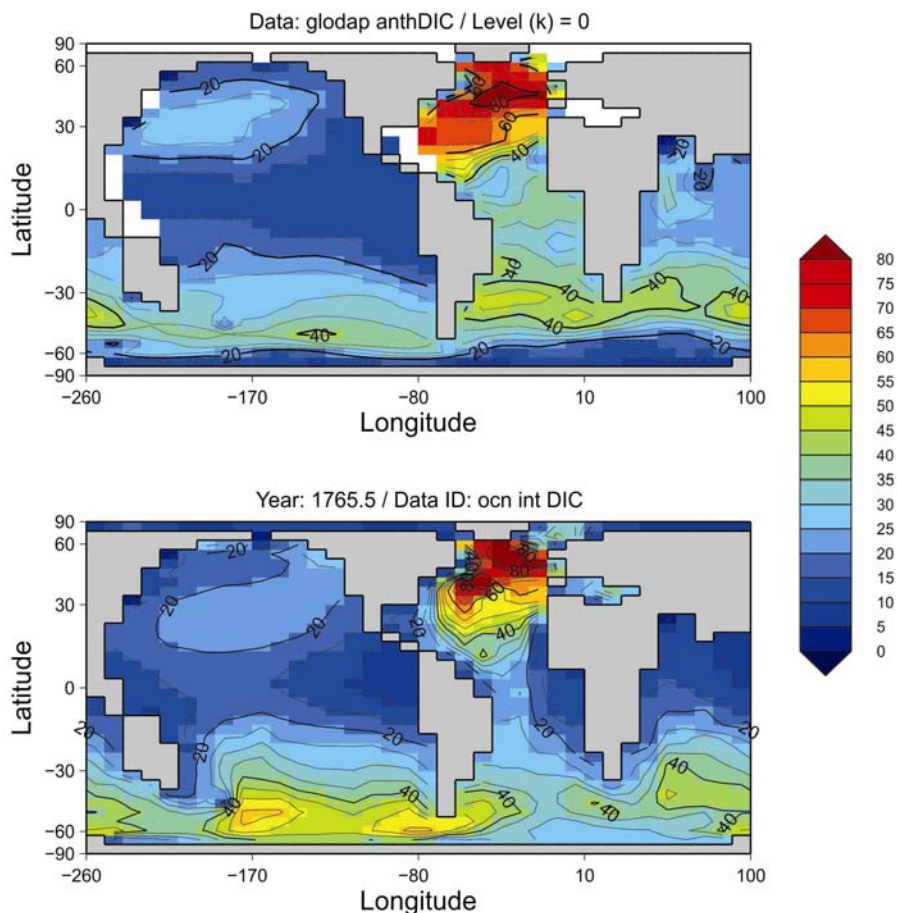
```
$ ./runmuffin.sh cgenie.eb_go_gs_ac_bg.worjh2.BASEFe LABS  
LAB.0100.future 90 LAB.0100.historical
```

and with the start from year now set to year 2010 (the end year of the historical transient):

```
bg_par_misc_t_start=2010.0
```

Note that the user-config LAB.0100.future is not provided for you this time – create this (or a file named however you like) by copying e.g. LAB.0100.historical and making the parameter changes described above (forcing specification parameter, scaling parameter, start year parameter).

You can also easily replace the details of the emissions with other SRES scenarios – simply find the year vs. emissions rate information from the interweb (e.g. http://sres.ciesin.columbia.edu/final_data.html) and edit (or copy-and-paste) the flux values for each decade into the file biogen_force_flux_atm_pCO2_sig.dat in the forcing directory. cGENIE will then automatically interpolate between the decadal tie-points to give a continuous change in emissions. Now you are able to make a rather more realistic/plausible assessment of when and where potential ecological impacts (via assumed ocean chemistry criteria) might occur.



Observed (top) vs. Model (bottom) anthropogenic CO₂ inventories.

Data and model water column integrals in units of mol CO₂ m⁻² and are nominally with respect to year 1994.

S1. Ocean acidification output supplement

S1.0 Variables relevant to ocean acidification are saved in the 'usual' formats – *time-series* (.res) files, 3D netCDF, and also 2D netCDF *time-slice* files. The most relevant ones of interest are as follows (although this is not an exhaustive list of what might be relevant to look at ...) Refer to the lecture for a refresher on ocean acidification ...

S1.1 *Time-series* files:

biogem_series_misc_surpH.res	- mean ocean surface pH
biogem_series_carb_sur_ohm_cal.res	- mean ocean surface carbonate saturation (calcite)
biogem_series_carb_sur_ohm_arg.res	- mean ocean surface carbonate saturation (aragonite)
biogem_series_fexport_CaCO3.res	- global CaCO ₃ export (from planktic calcifiers)
biogem_series_fexport_POC.res	- global organic matter export (from all plankton)

And of course, time-series of *temperature* (e.g. mean surface, ocean or air temperature), Atlantic meridional overturning strength (AMOC), and sea-ice cover.

S1.2 3D netCDF *time-slice* file (fields_biogem_3d.nc):

misc_pH	- ocean pH
carb_ohm_cal	- carbonate saturation (calcite)
carb_ohm_arg	- carbonate saturation (aragonite)
bio_fpart_CaCO3	- CaCO ₃ export flux (from planktic calcifiers)
bio_fpart_POC	- organic export flux (from all plankton)

And of course, *time-slices* of temperature and salinity, perhaps also dissolved oxygen (ocn_O2) that decreases in a warming ocean and also if the organic carbon flux increases.

S1.3 2D netCDF *time-slice* file (fields_biogem_2d.nc):

ocn_int_DIC	- is the water-column integrated inventory of dissolved carbon (refer to hand-out)
misc_sur_rCaCO3toPOC	- is the ratio between CaCO ₃ and POC export – i.e. comparatively, how much CaCO ₃ is exported

And of course, the overturning stream-function.