

Using ECCO v4

Gaël Forget
Department of Earth, Atmospheric and Planetary Sciences
Massachusetts Institute of Technology

October 30, 2015

abstract

These notes pertain to the ECCO v4 state estimate, model setup, and associated codes (Forget et al., 2015). Section 1 summarizes download procedures and links to additional documentation¹. Section 2 explains how ECCO v4 solutions, or corresponding short regression tests, can be re-run.

Contents

1	Downloading ECCO v4	2
1.1	Released ECCO v4 Solution	2
1.2	Diagnostic Tools	2
1.3	ECCO v4 setup	2
2	Running ECCO v4	3
2.1	Baseline ECCO v4 solutions	3
2.2	Short Regression Tests	5
2.3	Iterative Optimization Test Case	6

References

Forget, G., J.-M. Campin, P. Heimbach, C. N. Hill, R. M. Ponte, and C. Wunsch, 2015: Ecco version 4: an integrated framework for non-linear inverse modeling and global ocean state estimation. *Geoscientific Model Development*, **8** (10), 3071–3104, doi:10.5194/gmd-8-3071-2015, URL <http://www.geosci-model-dev.net/8/3071/2015/>.

¹Throughout this document links are indicated by blue colored font.

1 Downloading ECCO v4

This section first provides direction to download the ECCO v4–release 1 state estimate output (section 1.1) and associated matlab analysis tools (section 1.2). It then explains download procedures for the ECCO v4 model setup and MITgcm (section 1.3). Sections 1.2 and 1.3 rely on [the MITgcm cvs server](#) (see section 1.3). As an alternative, frozen versions are available [here](#).

1.1 Released ECCO v4 Solution

The model output for the ECCO v4–release 1 state estimate is available via [this.opendap server](#) and [this ftp server](#) from [ecco-group.org](#). The servers provide the [grid files](#) and [monthly output fields](#) in ‘nctiles’ format, as well as [collocated in situ and state estimate profiles](#) in ‘MITprof’ format. The ‘nctiles’ and ‘MITprof’ format are described in Forget et al. (2015). The files can be downloaded at the command line, e.g. within a linux environment, by typing

```
wget --recursive ftp://mit.ecco-group.org/ecco_for_las/version_4/release1/nctiles_grid
wget --recursive ftp://mit.ecco-group.org/ecco_for_las/version_4/release1/nctiles
wget --recursive ftp://mit.ecco-group.org/ecco_for_las/version_4/release1/MITprof_release1
```

1.2 Diagnostic Tools

To analyze model output from section 1.1 or section 2.1, two sets of Matlab tools are available:

- The [gcmfaces+MITprof](#) framework (see Forget et al., 2015) normally gets installed using [this shell script](#) (via [cvs](#)). Alternatively, frozen versions are available [here](#).
- Basic MITgcm [tools](#) can also be downloaded via [cvs](#).

For example, one may readily regenerate [this set of diagnostics](#) for the ECCO v4–release 1 state estimate using ‘diags_driver.m’ as explained in [the gcmfaces.pdf documentation](#).

1.3 ECCO v4 setup

First, install the MITgcm using [cvs](#) as explained [here](#). Second, install the ECCO v4 model setup on the LLC90 and CS32 grids (see Forget et al., 2015) also via [cvs](#) according to:

```
cd MITgcm/verification
cvs co -P -d global_oce_llc90 MITgcm_contrib/gael/verification/global_oce_llc90
cvs co -P -d global_oce_cs32 MITgcm_contrib/gael/verification/global_oce_cs32
```

Alternatively, frozen versions are available [here](#). [global_oce_cs32/](#) (614M) is a small setup used only for testing, whereas [global_oce_llc90/](#) (595M) is the production setup that typically runs on 96 processors. Running and verifying the section 2.1 solutions furthermore requires downloading the three-hourly forcing fields (127G) and in-situ data inputs (7G) from:

```
cd MITgcm/verification
wget --recursive ftp://ecco2.jpl.nasa.gov/data1/ecco/version4r1/StandAloneForward/era_xx/
wget --recursive ftp://mit.ecco-group.org/ecco_for_las/version_4/release1/MITprof_inputs/
```

Running the section 2.2 regression tests instead requires ‘[global_oce_input_fields/](#)’ (1.6G):

```

36 cd MITgcm/verification
37 wget http://mitgcm.org/~gforget/global_oce_input_fields.tar.gz
38 gunzip global_oce_input_fields.tar.gz
39 tar xf global_oce_input_fields.tar
40 \rm -f global_oce_input_fields.tar

```

41 2 Running ECCO v4

42 This section explains how the ECCO v4 setup is used to rerun 20-year state estimates (section
43 2.1) or short regression tests (section 2.2). As a pre-requisite, one must have downloaded the
44 MITgcm as well as the ECCO v4 model setup and inputs (section 1.3). Based upon the section
45 1.3 directions, the various downloaded directories should be organized as illustrated in Fig.1
46 within the MITgcm directory. Running MITgcm furthermore requires the following software:
47 gcc, gfortran (or alternatives), mpi (only for parallel runs) and netcdf (only if ‘pkg/profiles’ is
48 used). Additional information can be found in [the MITgcm howto](#) and in [the MITgcm manual](#).

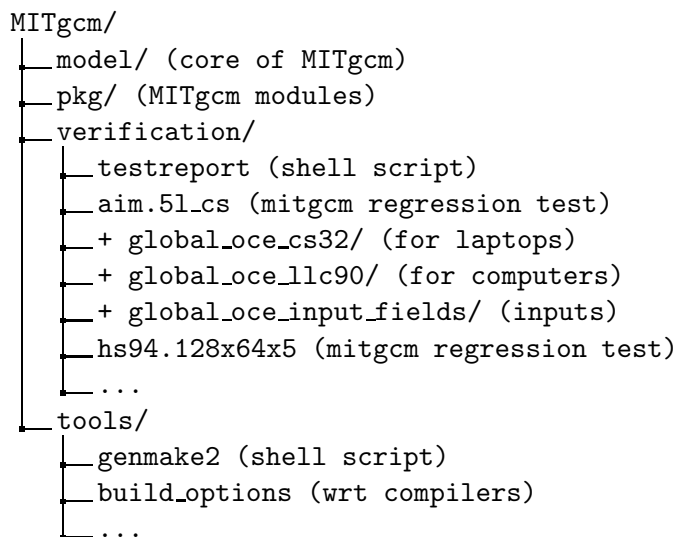


Figure 1: MITgcm directory structure including the ECCO v4 directories (indicated with “+”) downloaded according to the section 1.3 directions.

49 2.1 Baseline ECCO v4 solutions

50 The ‘baseline1’ and ‘baseline2’ versions of the 1992-2011 ECCO v4 state estimate are minor up-
51 dates of ‘release1’ (Forget et al., 2015) that are easiest for outside users to reproduce. ‘baseline1’
52 most closely matches the ‘release1’ output (section 1.1). ‘baseline2’ further benefits from: 1)
53 inclusion of geothermal heating at the sea floor; 2) increased adjoint stability due to additional
54 bottom viscosity. The small impact of these additional features is documented by Forget et al.
55 (2015). The standard analysis for ‘baseline1’ and ‘baseline2’ are furthermore available [here](#)².

²coming soon...

56 To re-run ‘baseline2’ one proceeds according to Fig. 2. A 20-year ECCO v4 model run
 57 typically takes between 6 to 12 hours on 96 cores (depending on the computing environment).
 58 To verify the re-run results one proceeds according to Fig. 3. The expected level of accuracy
 59 for 20-year re-runs (with an up to date MITgcm; on any given computer) is reached when the
 60 displayed values are all ≤ -3 (see Forget et al., 2015, for details).

61 To re-run ‘baseline1’ instead of ‘baseline2’ one needs a couple modifications to the setup:
 62 (a) define ‘ALLOW_KAPGM_CONTROL_OLD’ and ‘ALLOW_KAPREDI_CONTROL_OLD’ in
 63 ‘global_oce_llc90/code/GMREDI_OPTIONS.h’ before compiling the model;
 64 (b) copy ‘global_oce_llc90/input.ecco_itXX/data’ over ‘global_oce_llc90/input.ecco_v4/data’ be-
 65 fore running the model. Users who may want to reproduce ‘release1’ even more precisely than
 66 ‘baseline1’ does should contact ecco-support@mit.edu to obtain additional model inputs.

```
#1) compile the model
cd MITgcm/verification/global_oce_llc90/build
../../../../tools/genmake2 -mods=../code -optfile \
    ../../../../tools/build_options/linux_amd64_gfortran -mpi
make depend
make -j 4

#2) link files into run directory
cd ../run
ln -s ../build/mitgcmuv .
ln -s ../input.ecco_v4/* .
ln -s ../input_fields/* .
ln -s ../input_profiles/* .
ln -s ../era_xx .

#3) run model
mpiexec -np 96 ./mitgcmuv
```

Figure 2: Procedure to re-run the ECCO v4 state estimate (‘baseline2’ version). Pre-requisites: (1) installation of gcc, gfortran (or alternatives), and mpi (only for parallel runs); (2) installation of the MITgcm and ECCO v4 setup installation according to section 1.3. The contents of ‘input.ecco_v4’ (short text files) and ‘input_fields’ (grid and other binary input) should match those found [here](#). The contents of ‘era_xx’ directory should match [this](#). The contents of ‘input_profiles’ should match [this](#). These files can be downloaded freely, e.g., as explained in section 1.3.

67 The number of cores (96 by default and in Fig. 2) can be reduced to, e.g., 24 by copying
 68 ‘global_oce_llc90/code/SIZE.h_24cores’ over ‘global_oce_llc90/code/SIZE.h’ before compiling the
 69 model and then running it with ‘mpiexec -np 24 ./mitgcmuv’. Different compiler options (alter-
 70 natives to ‘linux_amd64_gfortran’ in Fig. 2) are available in ‘MITgcm/tools/build_options’.

71 For users holding a TAF license: to compile the adjoint one replaces ‘make -j 4’ with ‘make
 72 adall -j 4’ in Fig. 2; to run the adjoint one then turns on ‘pkg/autodiff’ in input.ecco_v4/data.pkg
 73 (i.e. useAUTODIFF=.TRUE.) and replaces ‘mitgcmuv’ with ‘mitgcmuv_ad’ in Fig. 2.

Figure 3: Top: instructions to verify (using ‘testreport_ecco.m’ within Matlab) that a re-run of the baseline ECCO v4 solution (currently ‘baseline2’) is acceptably close to the reference result. Bottom: example output from testreport_ecco.m where the re-run agrees up to 6 digits with the reference result. To activate additional tests (of meridional transports) one needs to have installed [gcmfaces](#) (see section 1.2) and uncomment the ‘addpath’ and ‘gcmfaces_global’ commands below (where ‘ /Documents/MATLAB/gcmfaces’ is a user specific path).

```
cd MITgcm/verification/global_oce_llc90
matlab -nodesktop -nodisplay

%addpath ~/Documents/MATLAB/gcmfaces;
%gcmfaces_global;

addpath results_itXX;%necessary .m and .mat files
mytest=testreport_ecco('run/');%compute the tests and display result
```

```
-----
      & jT & jS &      ... & (reference is)
run/ & (-6) & (-6) &      ... & baseline2
-----
```

74 2.2 Short Regression Tests

75 To ensure continued compatibility with the up to date MITgcm, the ECCO v4 model setup
 76 is also tested on a daily basis using the ‘testreport’ command line utility (indicated in Fig.1)
 77 that compares re-runs with reference results over a few time steps. The reader is referred to
 78 ‘testreport -help’ and [the MITgcm howto](#) for details. The short regression test of the smaller
 79 ECCO v4 setup ([global_oce_cs32/](#)) is thus executed by typing:

```
80 ./testreport -t global_oce_cs32
```

81 If everything proceeds as expected then the regression test results are reported to screen as
 82 shown in Fig. 4. On other machines the degree of agreement (16 digits in the Fig. 4 example)
 83 may vary and ‘testreport’ may indicate ‘FAIL’. Despite the dramatic character of such message,
 84 it may not prevent reproducing 20-year solutions (see section 2.1) with acceptable accuracy. An-
 85 other side note: for short regression test purposes, one can copy ‘global_oce_llc90/code/tamc.h_short’
 86 and ‘global_oce_llc90/code/PROFILES_SIZE.h_short’ over ‘global_oce_llc90/code/tamc.h’ and
 87 ‘global_oce_llc90/code/PROFILES_SIZE.h’ to somewhat reduce the memory footprint and disk
 88 storage.

89 It should be stressed that the bigger ECCO v4 setup ([global_oce_llc90/](#)) requires at least 12
 90 cores in forward mode (96 in adjoint mode) and therefore should not be run using the above
 91 command (or on a laptop). Instead the [global_oce_llc90/](#) regression tests use mpi:

```

default 10  ----T-----  ----S-----
G D M      c          m s          m s
e p a R g m m e . m m e .
n n k u 2 i a a d i a a d
2 d e n d n x n . n x n .

Y Y Y Y>14<16 16 16 16 16 16 16 16 16 pass global_oce_cs32

```

Figure 4: Abbreviated output of testreport to screen.

```

92 ./testreport -of ../tools/build_options/linux_amd64_gfortran \
93 -j 4 -MPI 24 -command 'mpiexec -np TR_NPROC ./mitgcmuv' \
94 -t global_oce_llc90

```

95 with 24 processors and gfortran (these settings may differ on another machine). To prevent users
96 from running the llc90 tests in serial mode (e.g. via a plain './testreport' command) the results
97 are in 'global_oce_llc90/results/hidden/'. To activate the llc90 tests, the 'output*' files contained
98 in this directory must be linked to 'global_oce_llc90/results/'. Also 'testreport' assumes that a
99 separate 'code_ad' directory is needed for the adjoint, which does not apply to ECCO version
100 4. To activate the adjoint tests, users holding a [TAF](#) license therefore want to create a soft link
101 called 'code_ad' that points to 'code' in [global_oce_cs32/](#) and [global_oce_llc90/](#).

102 The ECCO v4 daily regression test results (forward & adjoint, [cs32](#) & [llc90](#), using gfortran
103 and 24 processors) from a machine called 'glacier' are reported [on this site](#).

104 2.3 Iterative Optimization Test Case

105 The [global_oce_cs32/input_OI](#) directory implements an iterative optimization test case. It boils
106 down to optimal interpolation solved by a variational method using the MITgcm adjoint (the
107 ocean model being replaced with a simple diffusion equation here). The pre-requisites are:

- 108 1. run the adjoint benchmark in [global_oce_cs32/](#) via testreport (see section 2.2).
- 109 2. Go to MITgcm/lsopt and compile (see section 3.18 of [manual](#)).
- 110 3. Go to MITgcm/optim, replace 'natl_box_adjoint' with 'global_oce_cs32' in [this Makefile](#),
111 and compile as explained in section 3.18 of [manual](#). An executable named 'optim.x' should
112 get created in MITgcm/optim. If otherwise, please contact ecco-support@mit.edu
- 113 4. go to MITgcm/verification/global_oce_cs32/input_OI and type 'source ./prepare_run'

114 Then the iterative optimization itself proceeds as follows

- 115 1. ./mitgcmuv_ad > output.txt
- 116 2. ./optim.x > op.txt
- 117 3. increment optimcycle by 1 in data.optim

- 118 4. go back to step #1, to run the next iteration
- 119 5. type 'grep fc costfunction000*' to display results (Fig. 5).

```
costfunction0000: fc = 4118.1987222194211      0.00000000
costfunction0001: fc = 1523.9310891186672      0.00000000
costfunction0002: fc = 1053.3611790049420      0.00000000
costfunction0003: fc = 790.10479375339185      0.00000000
```

Figure 5: Results of iterative optimization after 3 iterations carried out as explained in section 2.3.