

ECCO v4 development notes

Gaël Forget

Department of Earth, Atmospheric and Planetary Sciences

Massachusetts Institute of Technology

November 21, 2014

abstract

These notes pertain to the revamped and augmented estimation modules of MITgcm; to the global setups and files used in regression tests; and to the analysis of the state estimate solution, including its re-runs. The intention is to compile and condense practical information for users and developers of ECCO v4 setups. Some of the included material will find its way in Forget and co authors (in preparation), and some should be added to the MITgcm [manual](#).

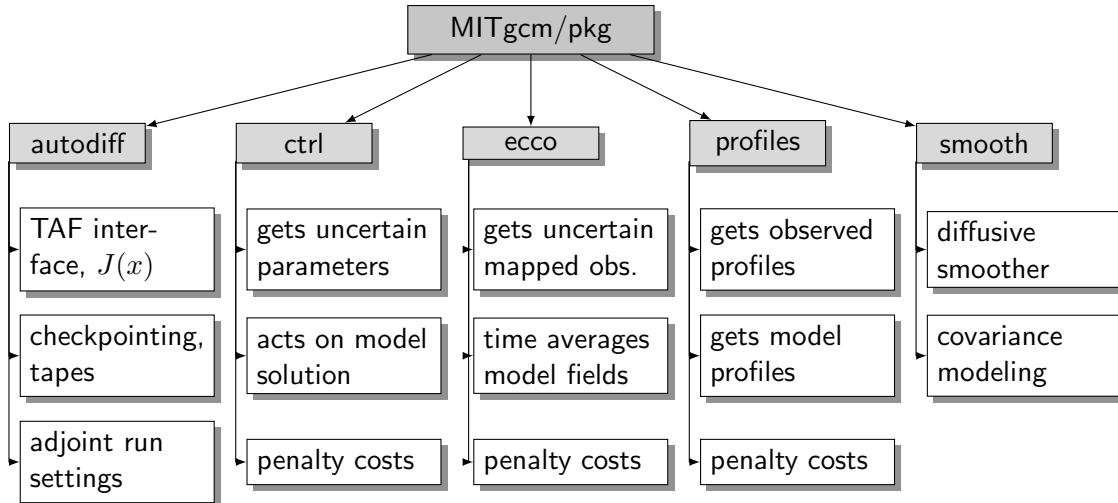


Figure 1: Organization and roles of MITgcm estimation modules.

Contents

1	downloads	3
1.1	MITgcm	3
1.2	setup	3
1.3	solution	4
1.4	analysis	4
2	MITgcm runs	6
2.1	regression tests	6
2.2	full ECCO v4 runs	6
	References	11

1 **1 downloads**

2 **1.1 MITgcm**

3 Pre-requisites are cvs, gcc, gfortran (or alternatives), and mpi (only for parallel runs). Then :

- 4 • The MITgcm web-page is [mitgcm.org](#)
5 • Install MITgcm using cvs as explained @ [cvs](#)
6 • Run MITgcm using `testreport` as explained @ [manual](#), [howto](#)

7 For example, my laptop setup, including mpi and netcdf, involved the following mac ports :

- 8 • cvs @1.11.23_1 (active)
9 • wget @1.14_5+ssl (active)
10 • gcc48 @4.8.2_0 (active)
11 • mpich-default @3.0.4_9+gcc48 (active)
12 • mpich-gcc48 @3.0.4_9+fortran (active)
13 • netcdf @4.3.0_2+dap+netcdf4 (active)
14 • netcdf-fortran @4.2_10+gcc48 (active)

15 Side note – overriding the default mac gcc and mpich with the above, further requires

- 16 • sudo port select –set gcc mp-gcc48
17 • sudo port select –set mpich mpich-gcc48-fortran

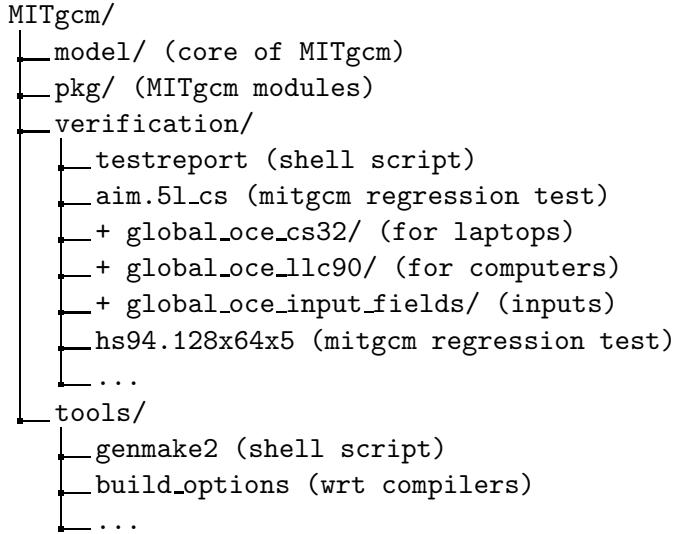
18 Side note – using mpi and netcdf within MITgcm requires two environment variables :

- 19 • export MPI_INC_DIR=/opt/local/include
20 • export NETCDF_ROOT=/opt/local

21 **1.2 setup**

22 Pre-requisites are MITgcm (see above) and mpi (except for small setup). User can then in-
23 stall the ECCO v4 setups, as explained @ [README](#), using the `setup_these_exps.csh` shell
24 script. This script downloads [global_oce_cs32/](#) (small setup), [global_oce_llc90/](#) (bigger setup)
25 and [global_oce_input_fields.tar.gz](#) binary model inputs to `global_oce_tmp_download/` (local sub-
26 directory). User can then move these directories to `MITgcm/verification/` to allow for automated
27 execution by `testreport` using `genmake2` (Fig.2).

Figure 2: MITgcm directory structure downloaded using [cvs](#). The ECCO v4 directories indicated with "+" were downloaded separately using [setup_these_exps.csh](#) script and moved to MITgcm/verification/.



28 1.3 solution

29 The release 1 solution directory linked to [ecco-group.org](#) contains :

- 30 • 20 year solution output : [readme](#), [fields](#), [profiles](#), [grid](#)
 31 • additional input files required to run the full 20 year solution (coming soon...).

32 1.4 analysis

33 Tools (e.g. matlab scripts) available to analyze the release1 solution, and others, include :

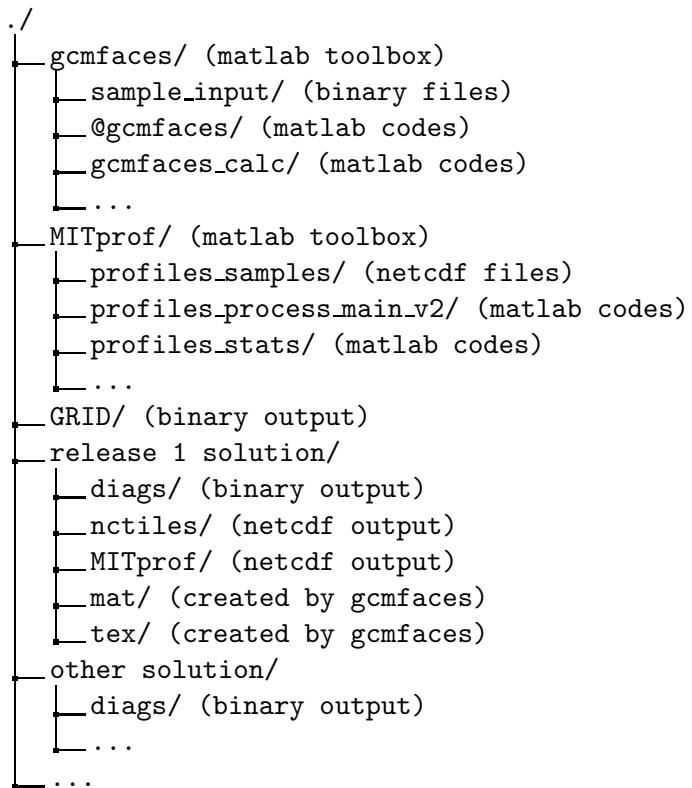
- 34 • download, set-up [gcmfaces](#) + [MITprof](#) using [shell script](#) or manually (see [getting_started.m](#))
 35 • download [MITgcm/utils](#) using [cvs](#) (basic functionalities only).

36 The so-called [standard analysis.pdf](#) is generated in matlab by means of [diags_driver.m](#) and
 37 [diags_driver_tex.m](#) in the following sequence :

38 `diags_driver('release1/','release1/mat/',1992:2011);`
 39 `diags_driver_tex('release1/mat/',[{}],'release1/tex/standardAnalysis');`

40 assuming the conventional directory structure shown in Fig.3.

Figure 3: Directory structure as expected by gcmfaces and MITprof toolboxes. The toolboxes themselves can be relocated anywhere as long as their locations are included in the matlab path. Advanced analysis using diags_driver.m and diags_driver.tex.m will respectively generate the mat/ directory (for intermediate computational results) and the tex/ directory (for [standard analysis](#)). This diagnostic process relies on the depicted organization of GRID/ and solution/ for automation (user will otherwise be prompted to enter directory names) and depends on downloaded copies of [fields](#) to nctiles/ (local subdirectory).



41 **2 MITgcm runs**

42 Here I document a few procedures, commands and submission scripts that may be relevant to run
43 the ECCO v4 MITgcm setup – either in short regression tests or for multi-decadal simulations
44 such as the full 20 year state estimate.

45 **2.1 regression tests**

46 MITgcm and ECCO v4 regression tests are run using testreport utility (see Fig.3; [howto](#)). Serial
47 regression tests can always be executed simply with, e.g.

48 `./testreport -t global_oce_cs32`
49 `./testreport -skipdir global_oce_llc90`

50 If something goes wrong and/or interrupts the process it is often safer to clean up experiment
51 directories (e.g., by executing `./testreport -clean -t global_oce_*`) and start over. For example,
52 the `global_oce_llc90` experiments require 12 processors in forward (96 in adjoint), and may crash
53 your laptop if you attempted to run them in serial mode.

54 Often in massive computing environments, however, mpi jobs can only be run within a
55 queuing system. The, machine specific, submission script in Fig.4 provides an example. It
56 contains 3 hard-coded switches : `fwdORad = 1` (2 for adjoint); `numExp = 1` (2 for llc90);
57 `excludeMPI = 0` (1 for serial). This script is located and submitted from MITgcm/verification.
58 If compute nodes cannot access the remote adjoint compiler (TAF), then proceed in two steps :

59 1. compile outside of the queuing system using e.g.

60 `./testreport -of ./tools/build_options/linux_amd64_ifort+mpi_ice_nas \`
61 `-j 4 -MPI 96 -command 'mpieexec -np TR_NPROC ./mitgcmuv_ad' -ad \`
62 `-t global_oce_llc90 -norun`

63 2. Before submiting the Fig.4 script, add `-q` to the `'opt'` variable to skip compilation.

64 Adjoint test require access to the TAF compiler. Then the call to testreport only needs to
65 be altered by appending the `'-ad'` option. The ECCO v4 regression tests do not include the
66 common, adjoint specific `'code_ad/'` directory, which is generally un-necessary. Since testreport
67 relies on the existence of `'code_ad/'` for its adjoint option though, it is necessary to soft link
68 `'code/'` to `'code_ad/'` in both [global_oce_cs32/](#) and [global_oce_llc90/](#) to run `'testreport -ad'` on
69 those experiments.

70 **2.2 full ECCO v4 runs**

71 There are three main differences between regression tests and full model runs (see [howto](#)) :

- 72
 - compilation and run are done without testreport and with compiler optimization.
 - additional forcing, control vectors and/or observational inputs are often necessary.
 - additional memory and disk space are often necessary.

75 The typical compilation sequence, taking ECCO v4 and pleiades as an example, is shown
76 in Fig.5. On a different machine, another opt file (see build_options) should be used. The
77 tamc.h_itXX and profiles.h_itXX headers allow for additional time steps, and additional in situ
78 data input, respectively. Also note that compiling the adjoint requires a TAF license. Once
79 that is done, user creates and enters a run directory, links everything into place (see Fig.6), and
80 finally submits a job to the queueing system (see Fig.7).

81 A mechanism, analogous to testreport but for long runs, has been introduced recently
82 (Forget and co authors, in preparation) that is `testreport_ecco.m`. Within matlab, its use re-
83 quires '`MITgcm/verification/global_oce_llc90/results_itXX/`' to be in the path. By itself it com-
84 pares cost functions and global mean time series to the state estimate values. These can be
85 extended to meridional transports, which requires `gcmfaces`

86 **material for separate paragraph :** memory and disk space requirements (provide specs,
87 timings); results of testreport; results of testreport_ecco.m

Figure 4: Example script to run mpi testreport via a queueing system.

```
#PBS -S /bin/csh
#PBS -l select=1:ncpus=16:model=ivy+4:ncpus=20:model=ivy
#PBS -l walltime=02:00:00
#PBS -q devel
#PBS -m n

#environment variables and libraries
#-----
limit stacksize unlimited
module purge
module load modules comp-intel/2013.1.117 mpi-sgi/mpt.2.10r6 netcdf/4.0
#
setenv LD_LIBRARY_PATH ${LD_LIBRARY_PATH}:${HOME}/lib
setenv MPI_IB_TIMEOUT 20
setenv MPI_IB_RAILS 2
setenv MPI_IB_FAILOVER 1
setenv MPI_CONNECTIONS_THRESHOLD 2049

#local variables and commands
#-----
set fwdORad = 1
set numExp = 1
set excludeMpi = 0
#
if ( ${numExp} == 1 ) then
    set nameExp = global_oce_cs32
    set NBproc = 6
else
    set nameExp = global_oce_llc90
    set NBproc = 96
endif
#
if ( ${excludeMpi} == 1 ) then
    set opt = '-of ../tools/build_options/linux_amd64_ifort -j 4'
else
    set opt = '-of ../tools/build_options/linux_amd64_ifort+mpi_ice_nas -j 4'
endif
#
if ( ${fwdORad} == 1 && ${excludeMpi} == 0 ) then
    ./testreport ${opt} -MPI \
    ${NBproc} -command 'mpieexec -np TR_NPROC ./mitgcmuv' -t ${nameExp}
else if ( ${fwdORad} == 2 && ${excludeMpi} == 0 ) then
    ./testreport ${opt} -MPI \
    ${NBproc} -command 'mpieexec -np TR_NPROC ./mitgcmuv_ad' -ad -t ${nameExp}
else if ( ${fwdORad} == 1 && ${excludeMpi} == 1 ) then
    ./testreport ${opt} -t ${nameExp}
else if ( ${fwdORad} == 2 && ${excludeMpi} == 1 ) then
    ./testreport ${opt} -ad -t ${nameExp}
endif

exit
```

Figure 5: Compilation directives, outside testreport, for intensive model runs.

```
cd verification/global_oce_llc90/build
../../../../tools/genmake2 -optfile=\\
../../../../tools/build_options/linux_amd64_ifort+mpi_ice_nas -mpi

make depend

\rm tamc.h profiles.h
cp ..../code/tamc.h_itXX tamc.h
cp ..../code/profiles.h_itXX profiles.h

make -j 4

(or for the adjoint :
make adtaf
make adall -j 4
)
```

Figure 6: Example script to setup the 20 year ECCO v4 state estimate. It is implied that user has filled directories /bla, /blaa, /blaaa and /blaaaa with appropriate forcing, observational, control vector, and pickup files.

```
#!/bin/csh -f

set forcingDir = ~/bla
set obsDir = ~/blaa
set ctrlDir = ~/blaaa
set pickDir = ~/blaaaa

source ..../input_itXX/prepare_run
cp ..../build/mitgcmuv .
\rm pick*ta EIG*
ln -s ${forcingDir}/EIG* .
ln -s ${obsDir}/* .
ln -s ${ctrlDir}/xx* .
ln -s ${pickDir}/pick* .

exit
```

Figure 7: Example script to run the 20 year ECCO v4 state estimate on 96 processors.

```
PBS -S /bin/csh
#PBS -l select=1:ncpus=16:model=ivy+4:ncpus=20:model=ivy
#PBS -l walltime=12:00:00
#PBS -q long

#environment variables and libraries
#-----
limit stacksize unlimited
module purge
module load modules comp-intel/2013.1.117 mpi-sgi/mpt.2.10r6 netcdf/4.0
#
setenv LD_LIBRARY_PATH ${LD_LIBRARY_PATH}:${HOME}/lib
setenv MPI_IB_TIMEOUT 20
setenv MPI_IB_RAILS 2
setenv MPI_IB_FAILOVER 1
setenv MPI_CONNECTIONS_THRESHOLD 2049

#run MITgcm
#-----
mpiexec -np 96 dplace -s1 ./mitgcmuv
exit
```

⁸⁸ **References**

- ⁸⁹ Forget, G. and co authors, in preparation: ECCO version 4: A global ocean modeling and state
⁹⁰ estimation framework. *Geoscientific Model Development*.
- ⁹¹ Forget, G. and R. Ponte, submitted: The partition of regional sea level variability. *Progress in*
⁹² *Oceanography*.
- ⁹³ Hansen, P. C., 1992: Analysis of discrete ill-posed problems by means of the l-curve. *SIAM*
⁹⁴ *review*, **34** (4), 561–580.
- ⁹⁵ Weaver, A. and P. Courtier, 2001: Correlation modelling on the sphere using a generalized
⁹⁶ diffusion equation. *Quarterly Journal of the Royal Meteorological Society*, **127** (575), 1815–
⁹⁷ 1846.