

ECCO v4 development notes

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

June 8, 2015

abstract

These notes pertain to the ECCO v4 state estimate, model setup, and associated codes (Forget et al., 2015). Section 1 points to the other elements of documentation that are available online, and associated download procedures. Section 2 provides guidance to ECCO v4 users interested in operating the ECCO v4 model set-up and/or reproducing the ECCO v4 solution. Section 3 documents the revamped and augmented estimation modules of MITgcm. Some of the included material is expected to move to to the MITgcm [manual](#).

Contents

| | | |
|----------|--|-----------|
| 1 | downloads | 3 |
| 1.1 | MITgcm | 3 |
| 1.2 | ECCO version 4 setup | 3 |
| 1.3 | solution | 4 |
| 1.4 | analysis tools | 4 |
| 2 | MITgcm runs | 6 |
| 2.1 | regression tests | 6 |
| 2.2 | full ECCO v4 runs | 6 |
| 3 | re-implemented ecco and ctrl packages | 11 |
| 3.1 | pkg/ecco run-time parameters | 12 |
| 3.2 | pkg/ctrl run-time parameters | 14 |
| 3.3 | MITgcm compiling options | 14 |

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 Discussions*, **8** (5), 3653–3743, doi:10.5194/gmdd-8-3653-2015, URL <http://www.geosci-model-dev-discuss.net/8/3653/2015/>.

1 downloads

Here I document locations and directions to download the MITgcm, the ECCO v4 model setup, the ECCO v4 state estimate output, and related diagnostic matlab tools.

1.1 MITgcm

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

- The MITgcm web-page is mitgcm.org
- Install MITgcm using cvs as explained @ [cvs](#)
- Run MITgcm using testreport (for one experiment) as explained @ [manual](#), [howto](#)

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

- cvs @1.11.23_1 (active)
- wget @1.14.5+ssl (active)
- gcc48 @4.8.2_0 (active)
- mpich-default @3.0.4_9+gcc48 (active)
- mpich-gcc48 @3.0.4_9+fortran (active)
- netcdf @4.3.0_2+dap+netcdf4 (active)
- netcdf-fortran @4.2.10+gcc48 (active)

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

- sudo port select –set gcc mp-gcc48
- sudo port select –set mpich mpich-gcc48-fortran

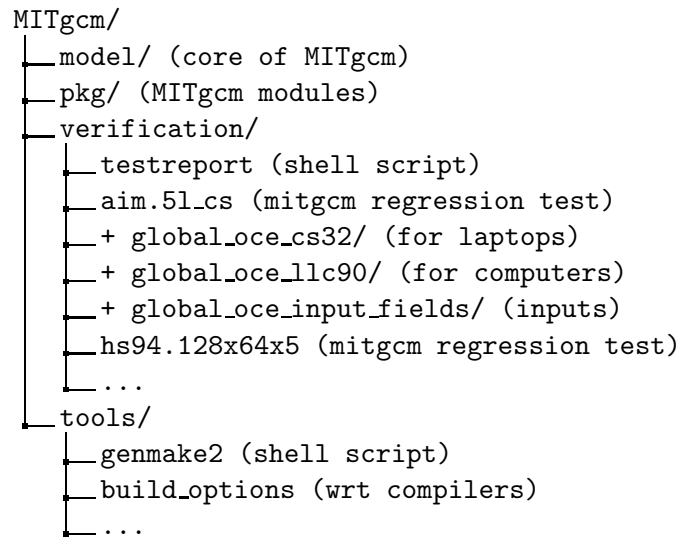
Side note – using mpi and netcdf within MITgcm requires two environment variables :

- export MPI_INC_DIR=/opt/local/include
- export NETCDF_ROOT=/opt/local

1.2 ECCO version 4 setup

Pre-requisites are MITgcm (see above) and mpi (except for small setup). User can then install the ECCO v4 setups, as explained @ [README](#), using the [setup_these_exps.csh](#) shell script. This script downloads [global_oce_cs32/](#) (small setup), [global_oce_llc90/](#) (bigger setup) and [global_oce_input_fields.tar.gz](#) binary model inputs to [global_oce_tmp_download/](#) (local sub-directory). User can then move these directories to [MITgcm/verification/](#) to allow for automated execution by [testreport](#) using [genmake2](#) (Fig.1).

Figure 1: MITgcm directory structure downloaded using `cv`s. The ECCO v4 directories indicated with "+" were downloaded separately using `setup_these_exps.csh` script and moved to `MITgcm/verification/`.



30 1.3 solution

31 The release 1 solution directory linked to ecco-group.org contains :

- 32 • 20 year solution output : [readme](#), [fields](#), [profiles](#), [grid](#)
- 33 • additional input files required to run the full 20 year solution (coming soon...).

34 1.4 analysis tools

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

- 36 • download, set-up [gcmfaces](#) + [MITprof](#) using [shell script](#) or manually (see [getting_started.m](#))
- 37 • download [MITgcm/utils](#) using `cv`s (basic functionalities only).

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

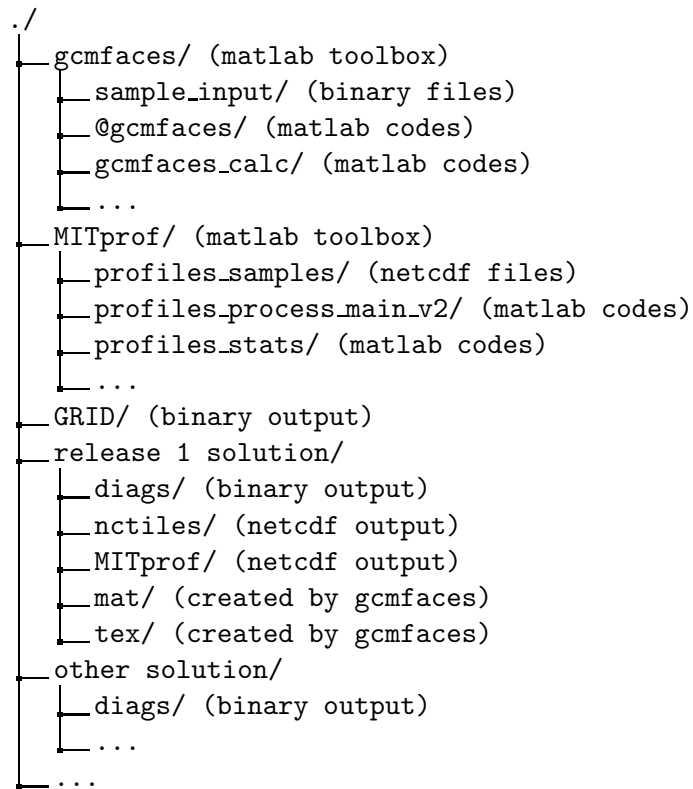
```

40 diags_driver('release1/', 'release1/mat/', 1992:2011);
41 diags_driver_tex('release1/mat/', {}, 'release1/tex/standardAnalysis');

```

42 assuming the conventional directory structure shown in Fig.2.

Figure 2: 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).



2 MITgcm runs

Here I document a few procedures, commands and submission scripts that may be relevant to run the ECCO v4 MITgcm setup – either in short regression tests or for multi-decadal simulations such as the full 20 year state estimate. Downloading MITgcm and the ECCO v4 setups is a pre-requisite (section 1.2).

2.1 regression tests

MITgcm and ECCO v4 regression tests are run using testreport utility (see Fig.2; [howto](#)). Serial regression tests can always be executed simply with, e.g.

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

```
./testreport -skipdir global_oce_llc90
```

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

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

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

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

2. Before submitting the Fig.3 script, add `-q` to the `'opt'` variable to skip compilation.

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

2.2 full ECCO v4 runs

Note to self ... ¹

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

- compilation and run are done without testreport and with compiler optimization.

¹Mention memory and disk space requirements, and additional input downloads.

- 77 • additional forcing, control vectors and/or observational inputs is necessary.
- 78 • additional memory and/or disk space is often necessary.

79 The typical compilation sequence is shown in Fig.4. The `tamc.h_itXX` and `profiles.h_itXX`
80 headers allow for additional time steps, and additional in situ data input, respectively. Also note
81 that compiling the adjoint requires a TAF license. Once that is done, user creates and enters a
82 run directory, links everything into place (see Fig.5), and finally submits a job to the queueing
83 system (see Fig.6).

84 A mechanism, analogous to `testreport` but for long runs, has been introduced recently
85 (Forget et al., 2015) that is `testreport_ecco.m` run within Matlab, which requires the downloaded
86 `'MITgcm/verification/global_oce_llc90/results_itXX/'` to be in the Matlab path. By itself it com-
87 pares cost functions and global mean time series to the reference state estimate values. These
88 can be extended to meridional transports, which requires `gcmfaces`. The typical call sequence is
89 indicated in the help of `testreport_ecco.m` and Fig.7.

Figure 3: Example script to run mpi testreport via a queuing system (machine dependent).

```
#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 'mpiexec -np TR_NPROC ./mitgcmuv' -t ${nameExp}
else if ( ${fwdORad} == 2 && ${excludeMpi} == 0 ) then
  ./testreport ${opt} -MPI \
  ${NBproc} -command 'mpiexec -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 4: Compilation directives, outside testreport, for intensive model runs. On a different machine (computer) another build option file such as linux_amd64_gfortran or linux_amd64_ifort11 should be used. To compile the adjoint, users need a TAF license and to replace ‘make -j 4’ with ‘make adall -j 4’. Note : the ‘-mods=../code’ specification can be omitted if the build directory contains the ‘genmake_local’ file).

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

make depend

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

make -j 4
```

Figure 5: 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 6: Example script to run the 20 year ECCO v4 state estimate on 96 processors (machine dependent).

```
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
```

Figure 7: Calling sequence to be executed from within matlab to verify that their re-run of the 20 year ECO v4 state estimate is acceptably close to the released state estimate.

```
addpath ../results_itXX;%necessary .m and .mat files
mytest0=testreport_ecco([], 'release1'); mytest0.info.interactive=0;%initialization
mytest=testreport_ecco(mytest0, 'release1', [-1:4], './', 1);%compute the tests
testreport_ecco(mytest, 'release1');%display the results
%testreport_write(mytest, 'myRun');%save the results to a mat file
```

90 3 re-implemented ecco and ctrl packages

91 State estimation consists in minimizing a least squares distance, $J(\mathbf{u})$, that is defined as

$$J(\mathbf{u}) = \sum_i \alpha_i \times (\mathbf{d}_i^T \mathbf{R}_i^{-1} \mathbf{d}_i) + \sum_j \beta_j \times (\mathbf{u}_j^T \mathbf{u}_j) \quad (1)$$

$$\mathbf{d}_i = \mathcal{P}(\mathbf{m}_i - \mathbf{o}_i) \quad (2)$$

$$\mathbf{m}_i = \mathcal{SDM}(\mathbf{v}) \quad (3)$$

$$\mathbf{v} = \mathcal{Q}(\mathbf{u}) \quad (4)$$

$$\mathbf{u} = \mathcal{R}(\mathbf{u}') \quad (5)$$

92 where \mathbf{d}_i denotes a set of model-data differences, α_i the corresponding multiplier, \mathbf{R}_i^{-1} the
 93 corresponding weights, \mathbf{u}_j a set of non-dimensional controls, β_j the corresponding multiplier,
 94 and additional symbols appearing in Eqs. 2-5 are defined below. The implementation of Eqs.1-
 95 5 and the adjoint interface within the MITgcm is charted in Fig. 8. A general presentation
 96 of Eqs.1-5 and Fig. 8 can readily be found in Forget et al. (2015). The focus here is on the
 97 underlying recent code development in the ‘pkg/ecco’ and ‘pkg/ctrl’ packages of MITgcm. These
 98 features are now tested daily via globaloce.cs32/ (adjoint experiment) that will also serve here
 99 for illustration in this document.

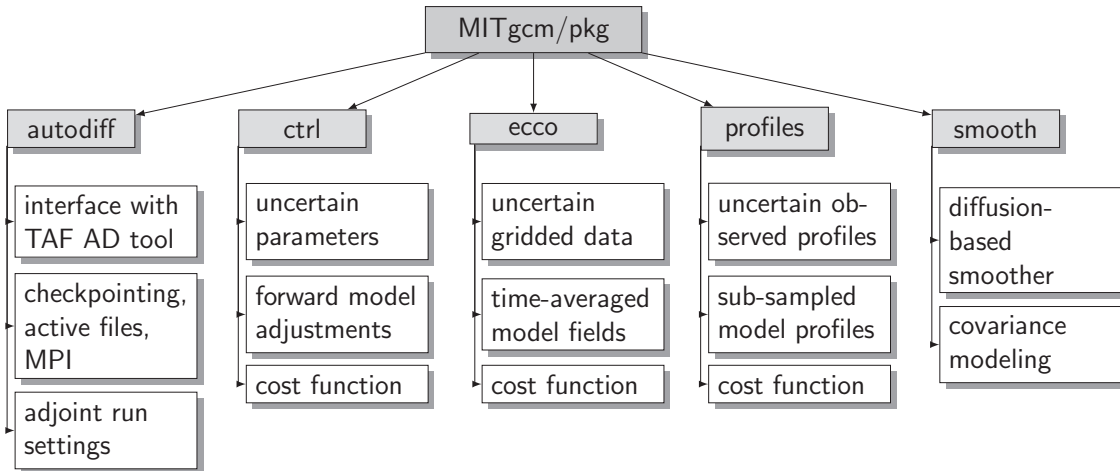


Figure 8: Chart of the organization and roles of MITgcm estimation modules. Additional details are reported in the MITgcm [manual](#), Forget et al. (2015), and section 3.

100 3.1 pkg/ecco run-time parameters

101 Note to self ... ²

102 Model counterparts (m_i) to observational data (o_i) derive from control parameters (\mathbf{v})
103 through the model dynamics (\mathcal{M}), diagnostic computations (\mathcal{D}), and averaging (or subsampling
104 in ‘pkg/profiles’) in space and time (\mathcal{S}). The physical variable in m_i is specified at run time via
105 the first characters in ‘gencost_barfile’ (to match the observed variable in o_i) as illustrated in
106 [this data.ecco](#) and [that data.ecco](#). The list of implemented variables as of the MITgcm check-
107 point c65l consists of ‘eta’, ‘sst’, ‘sss’, ‘bp’, ‘tauZon’, ‘tauMer’, ‘theta’, ‘salt’ (list obtained by:
108 `grep gencost_barfile pkg/ecco/cost_gencost_customize.F`). In the case of three dimensional vari-
109 ables (e.g. ‘theta’ or ‘salt’) the ‘gencost_is3d’ run-time option must be set to `.TRUE`. (it `.FALSE`.
110 by default). The file name for the observational fields (o_i) and the model-data uncertainty field
111 ($\sqrt{\mathbf{R}_i}$) are specified at run time via ‘gencost_datafile’ and ‘gencost_errfile’ respectively. The cost
112 function multiplier (α_i) further needs to be specified by ‘mult_gencost’ (it is 0. by default).

113 Both \mathcal{D} and \mathcal{S} in Eq.3 are mainly carried out as the forward model steps through time,
114 respectively by [ecco_phys.F](#) and [cost_averagesgeneric.F](#), and m_i is written to file periodically.
115 m_i and o_i normally are time series of daily or monthly averages, as specified at run time via
116 ‘gencost_avgperiod’. However dense time series of model time steps can also be employed for
117 testing purposes as illustrated in [this data.ecco](#). Furthermore climatologies of m_i can be formed
118 from its time series by [cost_genread.F](#) to allow for comparison with observational o_i climatologies.
119 This part of the m_i processing is carried out after the full time series has been written to file.
120 It is activated via the ‘gencost_preproc’ option as illustrated in [this data.ecco](#).

121 Model-data misfits are computed (Eq. 2) upon completion of the forward model simulation
122 by [cost_generic.F](#) that relies on [ecco_toolbox.F](#) for elementary operations and on [cost_genread.F](#)
123 for re-reading m_i from file. Plain model-data misfits ($m_i - o_i$) can be penalized directly (i.e.
124 used in Eq. 1 in place of d_i). More generally though misfits to be penalized (d_i in Eq. 1) derive
125 from $m_i - o_i$ through a generic post-processor (\mathcal{P} in Eq. 2). They can thus be smoothed in
126 space at run time via ‘gencost_posproc’ for example (see [this data.ecco](#)). The overall sequence
127 of operations for one cost function term is charted in Fig.9.

²Mention summary in stdout and cost function printouts

Algorithm 1 Generic cost function algorithm.

| | |
|--|---|
| 1: function COST_GENERIC(...) | ▷ Argument list defines the cost function |
| 2: call ecco_zero | ▷ Initialize local array to 0 |
| 3: call ecco_cprsr1 | ▷ Copy mask to local array |
| 4: for irec = 1, nrecloop do | ▷ Loop over time steps, days or months |
| 5: call cost_gencal | ▷ Get file names, pointers |
| 6: Begin cost_genread | ▷ Read, process model field |
| 7: if no preproc then | |
| 8: call ecco_readbar | ▷ Read one record |
| 9: else if preproc=clim then | |
| 10: call ecco_readbar within loop | ▷ Average records |
| 11: end if | |
| 12: End cost_genread | |
| 13: call mdsreadfield | ▷ Read observational field |
| 14: call ecco_diffmsk | ▷ Compute masked model-data misfit |
| 15: if posproc=smooth then | |
| 16: call smooth_hetero2d | ▷ Smooth masked misfit |
| 17: end if | |
| 18: call ecco_addcost | ▷ Add to cost function |
| 19: end for | |
| 20: end function | |

Figure 9: Chart of the generic cost function routine in pkg/ecco.

128 3.2 pkg/ctrl run-time parameters

129 Note to self ... ³

130 The control problem is non-dimensional, as reflected by the omission of weights in control
131 penalties ($u_j^T u_j$, Eq.1). Non-dimensional controls are scaled to physical units through multipli-
132 cation by their respective uncertainty fields, as part of the generic pre-processor Q (Eq.4) that
133 can also include the spatial correlation model and/or a mapping in time such as the cyclic rep-
134 etition of mean seasonal controls for example. Pre-conditioner \mathcal{R} (Eq.5) does not appear in the
135 estimation problem itself (Eq.1), as it only serves to push an optimization process preferentially
136 towards certain directions of the control space.

137 Key pkg/ctrl generic routines :

- 138 • [ctrl_map_ini_gen.F](#) computes dimensional control vector adjustments (Eq.4).
- 139 • [ctrl_map_ini_gentim2d.F](#) computes dimensional control vector adjustments (Eq.4).
- 140 • [ctrl_map_gentim2d.F](#) maps time varying controls to active model variables.
- 141 • [ctrl_map_ini_genarr.F](#) maps time invariant controls to active model variables.
- 142 • [ctrl_cost_gen.F](#) computes cost function penalties for all generic controls (in Eq.1).

143 3.3 MITgcm compiling options

144 Note to self ... ⁴

145 Much of the legacy code that has been distributed as part of ‘pkg/ecco’ and ‘pkg/ctrl’ in
146 the past is now deprecated – it is superseded by the generic cost functions and controls codes
147 presented above. Most of the deprecated codes had not been tested or maintained for many
148 years, and consist of variations of the same operations duplicated many times. Another issue
149 was the lack of organization amongst the deprecated codes (unlike in Fig.8). The consensus was
150 that there was no point in keeping them around much longer.

151 For the time being the deprecated codes still exist but they are not compiled anymore unless
152 the ‘ECCO_CTRL_DEPRECATED’ compile option is added in e.g. ‘ECCO_CPPOPTIONS.h’
153 (see below for details). To further facilitate the transition from old to new setup, the ctrlUseGen
154 run-time parameter was added that switches between the old and new (generic) treatment of
155 control vectors (assuming that ‘ECCO_CTRL_DEPRECATED’ was defined at compile time).
156 As a side note: there is one non-generic feature that ISN’T deprecated since it has not been
157 re-implemented in generic fashion, which is the control of open boundary conditions.

158 The deprecation of the legacy codes leads to a vast reduction in the volume of estimation
159 codes (30% of the code treated by automatic differentiation, which includes the entire phys-
160 ical model, was removed in the process), a vast addition of capabilities (new or pre-existing
161 functionalities are now available for any gridded data set), and a greatly improved flexibility
162 (virtually all options can now be switched on/off at run time). Furthermore, the ecco, ctrl
163 and autodiff packages were made independent of each other, and to follow general principles of
164 MITgcm packages. Thus they can now be switched on/off at run time, independently (by virtue
165 of useECCO, useCTRL, useAUTODIFF).

³Mention [this data.ctrl](#) ... [that data.ctrl](#) ... [this CTRL_OPTIONS.h](#) ... eccodevel email

⁴Mention optim and packing

166 Compiling options are typically found in the ‘code/’ directory of any given setup of MIT-
167 gcm (when customized) or in the corresponding MITgcm package (when using defaults). The
168 most obvious difference between [the new setup](#) and [an old setup](#) is that [CPP_OPTIONS.h](#) now
169 disregards [ECCO_CPPOPTIONS.h](#) and uses the following instead :

- 170 • [AUTODIFF_OPTIONS.h](#) contains the few compile directives of pkg/autodiff. The maxi-
171 mum numbers of time steps are set in tamc.h
- 172 • [ECCO_OPTIONS.h](#) contains compile directives of pkg/ecco. Very few remain necessary,
173 since all generic cost function settings can now be chosen at run time. The maximum
174 numbers of cost terms are set in ecco.h
- 175 • [CTRL_OPTIONS.h](#) contains compile directives of pkg/ctrl. Very few remain necessary,
176 since all generic control settings can now be chosen at run time. The maximum numbers
177 of controls are set in CTRL_SIZE.h
- 178 • along with [MOM_COMMON_OPTIONS.h](#), [GMREDL_OPTIONS.h](#), [GGL90_OPTIONS.h](#),
179 [PROFILES_OPTIONS.h](#), [EXF_OPTIONS.h](#), [SEAICE_OPTIONS.h](#), [DIAG_OPTIONS.h](#)