Geosci. Model Dev. Discuss., 8, 3653–3743, 2015
www.geosci-model-dev-discuss.net/8/3653/2015/
doi:10.5194/gmdd-8-3653-2015

Geoscientific Model Development Discussions

ECCO version 4: an integrated framework for non-linear inverse modeling and global ocean state estimation

G. Forget[1], J.-M. Campin[1], P. Heimbach[1,2,3], C. N. Hill[1], R. M Ponte[4], and C. Wunsch[5]

[1]Dept. of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge MA 02139, USA
[2]Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX 78712, USA
[3]Jackson School of Geosciences, The University of Texas at Austin, Austin, TX 78712, USA
[4]Atmospheric and Environmental Research, Inc., Lexington MA 02421, USA
[5]Dept. of Earth and Planetary Sciences, Harvard University, Cambridge, MA 02139, USA

Received: 25 February 2015 – Accepted: 8 April 2015 – Published: 5 May 2015

Correspondence to: G. Forget (gforget@mit.edu)

Published by Copernicus Publications on behalf of the European Geosciences Union.

Abstract

This paper presents the ECCO v4 non-linear inverse modeling framework and its baseline solution for the evolving ocean state over the period 1992–2011. Both components are publicly available and highly integrated with the MITgcm. They are both subjected to regular, automated regression tests. The modeling framework includes sets of global conformal grids, a global model setup, implementations of model-data constraints and adjustable control parameters, an interface to algorithmic differentiation, as well as a grid-independent, fully capable Matlab toolbox. The reference ECCO v4 solution is a dynamically consistent ocean state estimate (ECCO-Production, release 1) without un-identified sources of heat and buoyancy, which any interested user will be able to reproduce accurately. The solution is an acceptable fit to most data and has been found physically plausible in many respects, as documented here and in related publications. Users are being provided with capabilities to assess model-data misfits for themselves. The synergy between modeling and data synthesis is asserted through the joint presentation of the modeling framework and the state estimate. In particular, the inverse estimate of parameterized physics was instrumental in improving the fit to the observed hydrography, and becomes an integral part of the ocean model setup available for general use. More generally, a first assessment of the relative importance of external, parametric and structural model errors is presented. Parametric and external model uncertainties appear to be of comparable importance and dominate over structural model uncertainty. The results generally underline the importance of including turbulent transport parameters in the inverse problem.

1 Introduction

The history of inverse modeling in oceanography goes back at least four decades (see Wunsch, 2006, for a general presentation). The canonical oceanographic inverse problem as implemented by Wunsch (1977) consisted in estimating the time mean abso-

lute ocean circulation from synoptic, ship-based, hydrography transects. The physical model combined thermal-wind shear (diagnosed from observations) and a continuity equation. The model parameter to be estimated (i.e. the control vector) was the "reference level velocity". Least squares provide an adequate formulation to this inverse problem, and a practical method to avoid mis-interpreting geophysical noise (synoptic eddies, internal waves, etc.) for time mean ocean circulation features (Wunsch, 1977).

The original implementation has been extended substantially over subsequent decades, and some of the key technical developments are worth recalling (see also Wunsch and Heimbach, 2013a, for a review of the state of the art), as they provide the context for the present work. Non-linearities were introduced (Mercier, 1986) to incorporate optimal interpolation of hydrographic data in the inverse problem. While diapycnal and horizontal diffusion were also introduced early on (Schott and Zantopp, 1980; Olbers et al., 1985), the need for extending the inversion problem to parameterized advective eddy transports (Gent and Mcwilliams, 1990) was not fully appreciated until the study of Ferreira et al. (2005).

Time dependency and the use of Lagrange multipliers (i.e. the adjoint method) were first introduced in ocean inverse modeling by Thacker and Long (1988) and Holland and Malanotte-Rizzoli (1989) and applied to general circulation models by Tziperman and Thacker (1989) and Tziperman et al. (1992a, b). Later on algorithmic differentiation (AD) was introduced, making the use of Lagrange multipliers more practical (Griewank, 1992; Giering and Kaminski, 1998). Its application (Marotzke et al., 1999; Heimbach et al., 2002, 2005) to the Massachusetts Institute of Technology general circulation model (MITgcm; Marshall et al., 1997; Adcroft et al., 2004b) allowed for implementation of the time varying non-linear inverse problem, as envisioned by Wunsch and Minster (1982) and Wunsch (1984), to the case of actual observations (Stammer et al., 2002; Ferron and Marotzke, 2003).

The MITgcm AD capabilities remain exceptional amongst general circulation models. Over the last decade, in the context of the Estimating the Circulation and Climate of the Ocean (ECCO) project, the MITgcm non-linear inverse modeling framework (using

the adjoint method and algorithmic differentiation) has become a common tool for data synthesis, applied by many investigators to derive ocean state estimates (Stammer et al., 2004; Wunsch et al., 2007; Köhl et al., 2007; Köhl and Stammer, 2008; Forget et al., 2008b; Wunsch and Heimbach, 2009; Hoteit et al., 2009; Forget, 2010; Mazloff et al., 2010; Köhl et al., 2012; Speer and Forget, 2013; Köhl, 2014; Losch et al., 2014; Dail and Wunsch, 2014).

General circulation models implement the primitive equations, which extend far beyond the physics and numerics used in common inverse box models. On the one hand, they readily provide a versatile tool for dynamical interpolation of virtually all types of observations. On the other hand, numerical modeling has to be regarded as an integral part of non-linear inverse modeling, and as a primary responsibility of groups carrying ocean state estimation. Indeed, the quality of the model and the adequacy of its settings determine the physical consistency of ocean state estimates. Hence the state estimation group at MIT has become a main contributor of MITgcm code including, but not limited to the estimation framework implementation. Furthermore, the development of the new ECCO version 4 (ECCO v4) estimate described here started with an extensive revisit of MITgcm settings.

These considerations prompt the joint depiction of forward model setup and estimation framework developments as part of ECCO v4, and of the baseline solution of the non-linear inverse model. The overarching goal, which is essential to the oceanographic community, is the unification of the two pillars of science, namely observations (emphasis here is on data of global coverage) and theory (of which general circulation models are a vehicle). Thus, the synergy between data analysis and modeling is a guiding thread of this paper.

As a complement to this paper, and a number of associated publications, the setup and baseline solution of ECCO v4 are thoroughly documented by an extended suite of diagnostics (the "standard analysis" provided as Supplement) that users can readily download or reproduce. Daily and monthly regression tests are run for, respectively, a few time steps and 20 years. This will allow, for the foreseeable future, any user

to generate additional output that may be needed for extended data-model analyses. Thus, the authors aim to provide ECCO v4 as a fully-integrated non-linear inverse modeling framework including its baseline time-dependent solution that any interested user can readily start from, analyze and/or accurately re-run.

The foundation of the ECCO v4 model setup is a set of global grids of the earth surface (Sect. 2). The design, implementation and specification of the forward model setup and of the estimation framework are presented in Sects. 3 and 4 respectively. The baseline ECCO v4 solution (the ECCO-Production, release 1 state estimate) is the subject of Sect. 5, which is followed by conclusions and perspectives (Sect. 6).

2 Global grids

The most visible grid improvement, as compared with earlier ECCO configurations, is the extension of the gridded domain to the Arctic. This limitation of ECCO estimates produced until 2008 was due to the use of a latitude–longitude grid (LL; left panel of Fig. 1) that simply follows straight lines in spherical-polar projection, and requires an exponentially decreasing time step when approaching the North pole.

The cubed-sphere grid (CS; center panel in Fig. 1) has been successfully used in various MITgcm calculations (e.g., Menemenlis et al., 2005a; Marshall et al., 2007b) resolving the Arctic. The CS grid is a conformal mapping of a sphere to a cube surface, such that each face contains one sixth of the so gridded earth (Rančić et al., 1996; Purser and Rančić, 1998). At its 8 vertices (cube corners), CS grid lines converge exponentially as resolution increases, but more slowly than for LL at geographic poles.

However, a number of shortcomings of CS have been noted. First, loss of orthogonality near the cube corner is exacerbated when increasing horizontal resolution. Second, some of the vertices have to be placed on ocean-covered areas, and have an exceedingly high resolution, requiring unnecessarily small time steps. Third, such grids represent an obstacle for new users who were accustomed to latitude–longitude grids.

These considerations led to the design of the Lat-Lon-Cap grid (LLC; right panel in Fig. 1) such that

1. the grid reverts to a simple LL sector between 70° S and 57° N;
2. grid vertices are located over land;
3. grid heterogeneities remain acceptable at $\frac{1}{48}^{\circ}$ resolution.

At mid-latitudes, within the LL sector, the LLC grid is locally isotropic with grid spacing varying in $\cos(\varphi)$, where φ denotes latitude. At low latitudes LLC is refined in the meridional direction to better resolve the tropical system of zonal currents[1]. Grid scaling properties are shown in more detail in Figs. 2 and 3. The LL sector mesh derives from a simple analytic formulation based of geographic latitude and longitude. Users who may not be particularly invested in high latitude research may skip over the rest of this section, and simply extract the LL sector (70° S and 57° N) out of global LLC fields (see Sect. "Code availability").

Poleward of 57° N, LLC is topologically equivalent to CS minus one cube face (Fig. 1). The vertices of the *Arctic cap* are placed at a latitude of 67° N and in a specific orientation such that they all fall over land (Fig. 2, middle panel). For any given Arctic face dimension, LLC has the added advantage of an increased resolution in the Arctic as compared with CS, which has vertices at 45° N. The LLC grid is specifically designed for ocean simulations, whereas the CS grid is more suitable for atmospheric simulations. Between the LL sector and the Arctic cap, the grid makes a gradual, conformal transition that is evident in Fig. 3 between 57 and 67° N. To the South of 70° S, LLC is topologically equivalent to a *pillow case* that would have two vertices in each hemisphere (Fig. 2, right panel). The vertices are again placed over land at a latitude of 80° S and away from the Ross and Weddell ice shelves. Details of the grid generation method are reported in Appendix A.

[1] An alternative version of LLC that remains locally isotropic in the tropics is also available.

Looking beyond the immediate need for a truly global coarse resolution grid, we chose to generate a parent $\frac{1}{48}^\circ$ global grid[2]. The main advantages of this approach are that a full suite of lower resolution grids readily descend from the parent grid, and that the entire suite of commensurate grids share grid lines. Thus, in principle, one can easily nest back and forth between grids of different resolution. These global grids are being used in a number of model setups at 1° resolution (Danabasoglu et al., 2014; Marshall et al., 2014; Köhl, 2014) and at full resolution (D. Menemenlis, personal communication, 2014).

The parent $\frac{1}{48}^\circ$ global grid has 17 280 longitudinal grid cells. It is labelled "LLC4320" since the common face dimension is 4320. This grid size was chosen to maximize the number of integer factors available for coarsening the parent grid and for partitioning the computational domain in parallel computer environments[3]. It has 64 whole integer factors in total, so that the $\frac{1}{48}^\circ$ grid can readily, accurately be coarsened to e.g. $\frac{1}{24}^\circ$, $\frac{1}{16}^\circ$, $\frac{1}{12}^\circ$, $\frac{1}{8}^\circ$, $\frac{1}{6}^\circ$, $\frac{1}{4}^\circ$, $\frac{1}{3}^\circ$, $\frac{1}{2}^\circ$, 1°, 2° or 4°. This is a desirable property for a long-term project such as ECCO, in which spatio-temporal resolution is expected to increase in the future as computing capability and observational data base will keep increasing. A high degree of factorization also provides a convenient basis for down-scaled regional computations that employ boundary conditions from the state estimate (Sect. 5).

Advanced gridding has clear advantages from the standpoint of numerical ocean modeling. It can however put additional burdens on users of ocean model output who may find themselves coding the same diagnostics over and over again to accommodate different grids. One common approach is to distribute fields that were interpolated to a simpler grid (e.g. LL). This approach however tends to introduce sizable errors (e.g. in areal integrals and transports). A different and simple approach to the analysis of model output is chosen here that does not alter the results but alleviates the burden of grid specifics when analyzing model output – the "gcmfaces" analysis framework

[2]The resolution along the equator is quoted as $\frac{1}{48}^\circ$ (see Fig. 3 for details).

[3]Number 17 280 is known as the compositorial of 10, i.e the product of composite numbers less than or equal to 10 (see Wells, 2011).

that mimics the gridded earth decomposition of general circulation models in Matlab (Appendix C).

3 Model configuration

The model configuration presented below is the ECCO v4 setup used in state estimation (Sects. 4 and 5) and based on the LLC90 grid (Sect. 2). Variants of the ECCO v4 setup are also used in un-optimized model simulations (Danabasoglu et al., 2014; Marshall et al., 2014). The setup uses fully supported options of MITgcm software, is archived and regularly benchmarked (Appendix F), and is freely available along with the MITgcm itself (Sect. "Code availability").

The MITgcm, as configured in ECCO v4, solves the hydrostatic, Boussinesq equations (Marshall et al., 1997) using the z^* rescaled height vertical coordinate (Adcroft and Campin, 2004) and the vector-invariant form of the momentum equation (Adcroft et al., 2004a). This latter choice yields a discretized momentum equation without metric terms, which simplifies the handling of elaborate grids such as LLC90 (Sect. 2). This section summarizes the model equations, settings, and new MITgcm features (MITgcm Group, 2002; Adcroft et al., 2004b) used in ECCO v4. The novelty here largely resides in additions of forward model features to the body of adjointed codes (Sect. 4) and in their use in the state estimate (Sect. 5). Table 1 provides a list of basic model settings.

The relative importance of various model settings generally depends on the ocean state characteristic of interest. Here, a selection of ocean state characteristics is made amongst model-data differences (see Sect. 4), monthly time series of global mean quantities, and meridional transports (Table 2). These characteristics are used to gauge perturbations of 20 year solutions to various model settings (Table 3). They are also used to benchmark state estimate re-runs and model revisions (first 3 rows of Table 3; Appendix F) and to gauge the sensitivity of the state estimate to adjusted control parameters (Sect. 5).

3.1 Basic equations

For a water column that extends from the bottom at $z = -H$ to the free surface at $z = \eta$, the z^* vertical coordinate is defined as $z = \eta + s^* z^*$ with the scaling factor $s^* = 1 + \eta/H$. In this section, the notation ∇_{z^*} indicates the nabla operator at constant z^*, i.e., in a plane of constant z^* value. The z^* coordinate set of equations was introduced by Adcroft and Campin (2004) (their Eqs. 9–11 and 13). Written in vector-invariant formulation they read:

$$\frac{\partial \boldsymbol{v}}{\partial t} + (f + \zeta)\hat{\boldsymbol{k}} \times \boldsymbol{v} + \nabla_{z^*}\mathrm{KE} + w\frac{\partial \boldsymbol{v}}{\partial z} + g\nabla_{z^*}\eta + \nabla_h \Phi' = \boldsymbol{D}_{z^*,\boldsymbol{v}} + \boldsymbol{D}_{\perp,\boldsymbol{v}} + \mathcal{F}_{\boldsymbol{v}} \tag{1}$$

$$\frac{\partial \Phi'}{\partial z} = g\frac{\rho'}{\rho_c} \tag{2}$$

$$\frac{1}{H}\frac{\partial \eta}{\partial t} + \nabla_{z^*}(s^*\boldsymbol{v}) + \frac{\partial w}{\partial z^*} = s^*\mathcal{F} \tag{3}$$

$$\frac{\partial (s^*\theta)}{\partial t} + \nabla_{z^*}(s^*\theta \boldsymbol{v}_{\mathrm{res}}) + \frac{\partial (\theta w_{\mathrm{res}})}{\partial z^*} = s^*(\mathcal{F}_\theta + D_{\sigma,\theta} + D_{\perp,\theta}) \tag{4}$$

$$\frac{\partial (s^* S)}{\partial t} + \nabla_{z^*}(s^* S \boldsymbol{v}_{\mathrm{res}}) + \frac{\partial (S w_{\mathrm{res}})}{\partial z^*} = s^*(\mathcal{F}_S + D_{\sigma,S} + D_{\perp,S}) \tag{5}$$

where $\boldsymbol{v}$ is the horizontal velocity, $w^* = w/s^*$ is the vertical velocity in z^* coordinates[4], $\hat{\boldsymbol{k}}$ is the vertical unity vector, f and $\zeta = \nabla \times \boldsymbol{v}$ are the planetary and relative vorticity vertical component, KE is the horizontal kinetic energy, g is gravity, ρ' is the density anomaly relative to the constant Boussinesq density ρ_c ($\rho = \rho_c + \rho'$), Φ' is the pressure anomaly scaled by constant density ($\Phi' = p'/\rho_c$), θ and S are the potential temperature and salinity, $\boldsymbol{D}_{z^*}, D_\perp, D_\sigma$ are subgrid-scale (SGS) processes parameterized as mixing horizontally, vertically or along iso-neutral surfaces, and $\mathcal{F}_{\boldsymbol{v}}$, $\mathcal{F}$, $\mathcal{F}_\theta$, $\mathcal{F}_S$ are the forcing terms which are generally concentrated at the surface. Fields θ and S are advected (in Eqs. 4 and 5) by the residual-mean velocity field $(\boldsymbol{v}_{\mathrm{res}}, w_{\mathrm{res}}) = (\boldsymbol{v}, w) + (\boldsymbol{v}_{\mathrm{b}}, w_{\mathrm{b}})$ where $(\boldsymbol{v}_{\mathrm{b}}, w_{\mathrm{b}})$ is the bolus velocity parameterizing the effect of unresolved eddies (Gent and Mcwilliams, 1990, GM hereafter).

[4]For practical reason, the vertical velocity calculated by the model (w per Eq. 3) is neither the z^* vertical velocity w^* nor the true vertical velocity that would contain additional contributions (Adcroft and Campin, 2004, see their Eq. 4).

The z^* coordinate formulation introduces a major difference in the continuity equation (Eq. 3) as compared to the z coordinate. Thus, even in the abscence of fresh-water input ($\mathcal{F} = 0$), the divergence of the 3-D flow field (u, v, w) is no longer zero. Then, within the continuity equation, the rate of change of sea-surface elevation (the first term in Eq. 3) is uniformly distributed along the water column (as denoted by the $1/H$ scaling). The vertical velocity component w is obtained diagnostically from the continuity equation (Eq. 3). Furthermore, the horizontal momentum equation (Eq. 1) differs from the z coordinate case by the expression of horizontal pressure gradient (Adcroft and Campin, 2004, Eq. 15):

$$\nabla_h \Phi' = \nabla_{z^*}\Phi' + g\frac{\rho'}{\rho_c}\nabla_{z^*}\left(\eta\left(1 + \frac{z^*}{H}\right)\right) \tag{6}$$

where the second term represents the effect of gravity acting on the slope of constant z^* surface. The vertical momentum equation (Eq. 2) is reduced to the hydrostatic balance and sea-water density ρ is evaluated using the Jackett and McDougall (1995) equation of state in which pressure is assumed to be a function of only depth ($p = -\rho_c g z^*$) so that any compressible effect is completely removed.

Apart from the horizontal grid and the vertical coordinate z^*, the choice of time-stepping options used in ECCO v4 represents another major change compared to previous ECCO configurations. The time-discretized version of Eqs. (1)–(5) is reported in Appendix B, which is particularly important to understand budget and other diagnostics (Appendix C). In summary, the staggered time step approach is used, along with Adams–Bashforth 3 (AB-3) time stepping for momentum advection and the Coriolis term, 3rd order Direct Space Time tracer advection (DST-3; a multi-dimensional scheme), and 3rd order implicit tracer vertical advection (unconditionally stable). These

options improve the model stability, allowing for a longer time step. Thus, the time step restriction due to the Coriolis term in the Arctic is alleviated by the use of AB-3 (Δt = 1 h was unstable with AB-2 and $\epsilon_{AB} \sim 0.1$). Also, the chosen combination of staggered time-stepping and tracer advection schemes increases the stability limit related to internal-wave speed. With these choices a time step of Δt = 1 h is used (Table 1).

3.2 Volume and tracer conservation

ECCO v4 uses a non-linear free surface combined with *real freshwater flux* forcing and the z^* coordinate. This approach allows to include material exchanges through the free surface in a physically intuitive way (Campin et al., 2008) and to achieve exact tracer conservation, both locally and globally (Campin et al., 2004). To illustrate this point, it is useful to start from the vertical integral of Eqs. (3)–(5), which is

$$\frac{\partial \eta}{\partial t} + \nabla \cdot \int_{-H}^{\eta} \boldsymbol{v}\, \mathrm{d}z = (\mathrm{PmE})/\rho_c \tag{7}$$

$$\frac{\partial}{\partial t}((\eta + H)\overline{\theta}) + \nabla \cdot \int_{-H}^{\eta} \theta\, \boldsymbol{v}_{\mathrm{res}}\, \mathrm{d}z = Q_{\mathrm{net}}/(\rho_c C_p) + \int_{-H}^{\eta} D_{\sigma,\theta}\, \mathrm{d}z \tag{8}$$

$$\frac{\partial}{\partial t}((\eta + H)\overline{S}) + \nabla \cdot \int_{-H}^{\eta} S\, \boldsymbol{v}_{\mathrm{res}}\, \mathrm{d}z = S_{\mathrm{flux}}/\rho_c + \int_{-H}^{\eta} D_{\sigma,S}\, \mathrm{d}z \tag{9}$$

where the overbar denotes vertical averaging according to $\overline{\varphi} = \frac{1}{(\eta+H)} \int_{-H}^{\eta} \varphi\, \mathrm{d}z$.

The forcing terms $\mathcal{F}$, $\mathcal{F}_\theta$, $\mathcal{F}_S$ in Eqs. (3)–(5) are concentrated at or near the surface (unless geothermal heating at the bottom is active) and have been replaced by their integral form in Eqs. (7)–(9), namely the net fresh-water input at the surface (PmE, in $\mathrm{kg\,m^{-2}\,s^{-1}}$), the net heat flux into the water column (Q_{net}, in $\mathrm{W\,m^{-2}}$) and the salt-flux

at the surface (S_{flux}, in $\mathrm{g\,m^{-2}\,s^{-1}}$) which is zero in the absence of seaice and salinity relaxation (see Sect. 3.5).

With the non-linear free-surface, the water column thickness varies as the free-surface goes up and down (as apparent in Eqs. 7–9). With the z^* coordinate, this variation is distributed vertically over all grid-cells[5]. The fact that η enters the continuity equation (Eq. 7) also through $\int_{-H}^{\eta} \mathrm{d}z$ renders the free-surface non-linear; furthermore, time dependent grid-cell thickness introduces many more non-linearities which required code modifications to ensure efficient adjoint code generation via AD (Sect. 4.2).

Earlier ECCO configurations relied on the linear free-surface method (LFS), where column thickness and grid-cell thickness are fixed in time. The LFS version of Eqs. (7) and (9) is:

$$\frac{\partial \eta}{\partial t} + \nabla \cdot \int_{-H}^{0} \boldsymbol{v}\, \mathrm{d}z = \epsilon_{\mathrm{FW}}(\mathrm{PmE})/\rho_c \tag{10}$$

$$\frac{\partial}{\partial t}(H\overline{S}) + (\mathrm{PmE})/\rho_c \widetilde{S} + \nabla \cdot \int_{-H}^{0} S\, \boldsymbol{v}_{\mathrm{res}}\, \mathrm{d}z = S_{\mathrm{flux}}/\rho_c + \int_{-H}^{0} D_{\sigma,S}\, \mathrm{d}z \tag{11}$$

The Goldsbrough–Stommel circulation (Stommel, 1984) can be accounted for by setting $\epsilon_{\mathrm{FW}} = 1$ (*virtual fresh-water*) or ignored ($\epsilon_{\mathrm{FW}} = 0$). However, since grid-cell thickness is held fixed with LFS, the dilution effect due to surface freshwater flux needs to be represented explicitly as a *virtual salt flux* (2nd term in Eq. 11) using either the local surface salinity S or a constant S_o as $\widetilde{S}$, with drawbacks in both cases (see, e.g., Campin et al., 2008). By contrast, the non-linear free-surface formulation incorporates the dilution effect very naturally, within the time derivative of the water-column salt content (first term in Eq. 9).

[5]Each $z^* = -\alpha H$ level is a moving $z = z^* + (1-\alpha)\eta$ surface; $z = \eta$ at $z^* = 0$; $z = -H$ at $z^* = -H$.

The symmetry between continuity (Eq. 7) and tracer (Eqs. 8 and 9) equations allows for strict tracer conservation (Campin et al., 2004) when discretized consistently (Appendix B). In contrast, in the LFS case, this symmetry is lacking ($\partial\eta/\partial t$ in Eq. 10 has no counter part in Eq. 11) resulting in artificial tracer loss or gain (unless a global correction is added).

3.3 Tracer transports

Ocean tracers are advected by the residual mean velocity $\boldsymbol{v}^{\mathrm{res}}, w^{\mathrm{res}}$ (Eqs. 4 and 5). The present ECCO v4 uses the 3rd order DST scheme in the horizontal, and the implicit 3rd order upwind scheme in the vertical. Previous ECCO configurations used the explicit 3rd order upwind scheme in all directions. Flux limited advection schemes are also available in forward mode, although they are not used in the state estimate (Sect. 5), since they are not yet in the body of adjointed codes (Sect. 4). Choices of advection schemes are a concern to ocean state estimation, since their structural properties cannot generally be controlled by continuous parameters, and since numerical diffusion and advective overshoots could preclude an adequate fit to observations. Their importance can be gauged from Table 3. Thus, activating flux limiters has a sizable influence over 20 years, which is generally smaller than the impact of activating the C-D scheme, but exceeds the impact of activating geothermal heating for example (Table 3; see next section). Global mean times series often show an exceptionally high sensitivity to a variety of model settings, and to surface boundary layer settings in particular (Table 3).

Diffusion includes diapycnal and isopycnal components, the GGL mixed layer turbulence closure (Gaspar et al., 1990), and simple convective adjustment. The latter (GGL and convective adjustment) are used instead of the KPP vertical mixing scheme (Large et al., 1994) that was used in earlier ECCO configurations. The rationale for this choice and its impact on the 20 year solution is further discussed in Sect. 5. Time-invariant three dimensional fields of background diapycnal diffusivity ($\mathcal{K}_{\mathrm{d}}$), isopycnal diffusivity ($\mathcal{K}_{\sigma}$) and GM intensity ($\mathcal{K}_{\mathrm{gm}}$) are adjusted under observational constraints (Sect. 4) starting from constant first guess values (Table 1). The estimated parameter

maps for these highly uncertain coefficients become an integral part of the ECCO v4 model setup. In Sect. 5, the sensitivity to these parameters adjustments is evaluated, and compared with the results in Table 3. The geography of $\mathcal{K}_{\mathrm{gm}}$, $\mathcal{K}_{\sigma}$ and $\mathcal{K}_{\mathrm{d}}$, their impact on stratification, and their observability by means of Argo are further assessed in Forget (2015b).

3.4 Momentum discretization

Parameters of the momentum Eq. (1) currently used in ECCO v4 are provided in Table 1. Lateral eddy viscosity is harmonic and dependent on grid spacing, with coefficient given by $0.25 \times \mu L^2/\Delta t$, where $\mu = 2 \times 10^{-2}$ (viscAhGrid in Table 1) is a non-dimensional scaling number, L^2 is the spatially varying grid spacing squared (Fig. 2 shows L) and $\Delta t = 3600\,\mathrm{s}$ (deltaTmom in Table 1). The resulting viscosity varies from $\approx 10^3$ to $1.6 \times 10^4 \mathrm{m}^2\,\mathrm{s}^{-1}$, depending on location. The other dissipation contributions used in ECCO v4 are harmonic vertical viscosity and quadratic bottom drag (with parameters in Table 1) plus contributions from GGL.

Previous ECCO configurations used the C-D scheme (Adcroft et al., 1999) that interpolates the Coriolis term from the Arakawa C grid to a D grid and back. This scheme acts to reduce grid scale noise that is otherwise seen in the vertical velocity fields at all time scales, and particularly in the deep ocean (Fig. 4). Large vertical velocities have adverse effects on adjoint model stability, which ECCO originally resolved by means of the C-D scheme. The C-D scheme does however have a large impact on the large scale ocean circulation (Fig. 4) and hydrography (Table 3).

A comparable damping of the barotropic circulation could be obtained through a large increase in viscosity (not shown). Also vertical velocity noise is most intense near the ocean floor, which led us to the inference that adding viscosity more selectively near topography could suffice to damp the vertical velocity noise (Fig. 4) and, along with the use of vertical implicit advection, could stabilize the adjoint. Because the impact of this approach on the circulation and hydrography is more muted than that of the C-D scheme, the latter was abandoned in ECCO v4. Targeted viscosity increase near

topography remains needed to stabilize adjoint solutions (Sect. 4), but it can be omitted in forward solutions, as done in the state estimate (Sect. 5).

3.5 Surface boundary conditions

Upward buoyancy, radiative and mass fluxes (latent, sensible and radiative contributions to $\mathcal{F}_\theta$; evaporation as part of $\mathcal{F}$) through the free surface are computed using the bulk formulae of Large and Yeager (2004), and 6 hourly ERA-Interim re-analysis fields (Dee et al., 2011) for the near surface atmospheric state (temperature, humidity, downward radiation, precipitation). Downward shortwave radiation is allowed to penetrate, with exponential decay, to a depth of 200 m as part of $\mathcal{F}_\theta$ (Eq. 4). A seasonal climatology of runoff, from Fekete et al. (2002), is added as part of $\mathcal{F}$ (Eq. 3).

Earlier ECCO configurations using the *virtual salt flux* approach with $\epsilon_{\mathrm{FW}} = 0$ (Sect. 3.2) could only account for the dynamical impact of precipitation, evaporation and runoff as they affect buoyancy (see Ponte, 2006). Accordingly, they could only include seaice as a *levitating* layer without any direct effect on η (Campin et al., 2008). In contrast, ECCO v4 uses the *real freshwater flux* approach (Sect. 3.2) and thus further accounts for the dynamical effects of material exchanges through the free surface (either with the atmosphere, land or seaice) as shown in Campin et al. (2008).

Open ocean rain, evaporation and runoff simply carry (advect through the free surface) the local SST and zero salinity in the model. When seaice is present, buoyancy and mass fluxes[6] are recomputed based upon the thermodynamic balance of a fully interactive seaice model (Losch et al., 2010). In this model as configured in ECCO v4, seaice carries 0 °C and 4 g kg^{-1} salinity, while snow carries 0 °C and zero salinity.

The implementation of mass, buoyancy and momentum exchanges through the seaice–ocean interface in the rescaled z^* coordinate framework is presented in detail in Campin et al. (2008). A further correction was added in ECCO v4 to ensure conservation of heat for the combined ocean+seaice+snow system. While the ocean model

[6] $\mathcal{F}_\theta$, $\mathcal{F}_S$, $\mathcal{F}$ in Eqs. (3)–(5); PmE, Q_{net}, S_{flux} in Eqs. (7)–(9).

is configured to exchange freshwater at the local SST, the seaice model operates at constant internal heat, so it cannot freeze and melt at variable temperature. The added correction simply puts the heat differential back into the ocean. Ocean+seaice+snow budgets (as well as separate ocean, seaice, snow budgets) of mass, heat and salt are then closed to machine precision and readily diagnosed (Appendix C).

In centennial ocean model simulations, it is customary to add a Newtonian relaxation of surface salinity to an observation-based climatological map (e.g., Danabasoglu et al., 2014) as part of $\mathcal{F}_S$ (Eq. 5). While this method has no clear physical basis, it generally adds stability to centennial simulations. In contrast, the state estimate (Sect. 5) has no salinity relaxation term, so that $\mathcal{F}_S \neq 0$ only occurs when seaice (which salinity is set to 4 g kg^{-1}) melts or freezes. Salt rejected by seaice formation is distributed in the vertical using the parameterization of Duffy et al. (1999) and Nguyen et al. (2009) as part of $\mathcal{F}_S$.

Wind stress, also from ERA-Interim, is applied directly as part of $\mathcal{F}_{\boldsymbol{v}}$ (Eq. 1). A common alternative is to compute wind stress through bulk formulae. This approach typically requires backing out adequate drag coefficients – so that the results would approximately match surface stresses that, in atmospheric models, follow from a momentum balance rather than bulk formulae – or ad-hoc adjustments of atmospheric variables (see, e.g. Large and Yeager, 2004; Risien and Chelton, 2008). In the context of state estimation, the direct specification of wind stress further allows for a clear distinction between momentum controls on the one hand, and buoyancy and mass controls on the other hand (Sect. 4).

4 Estimation framework

The state estimation problem is defined here by a distance to observations (J) to be minimized under the constraint of a dynamical model. Section 4.1 formulates the state estimation problem in more detail. Section 4.2 provides an overview of the MITgcm adjoint, which is instrumental in solving the state estimation problem, and its recent devel-

opments in the context of ECCO v4. Aside from the dynamical model (Sects. 2 and 3) the defining ingredients of state estimation are model-data constraints and control parameters. Their ECCO v4 implementation and specifications are covered in Sects. 4.3 and 4.4.

The state estimate (Sect. 5) is a solution of the forward model (Sects. 2 and 3) at an approximate minimum of J. The process of finding such solutions, typically through an iterative optimization process and using the adjoint model, is not a focus of this paper. A number of well known optimization methods, with third party implementations freely available online, can be used to this end (see, e.g. Heimbach et al., 2005, and references therein).

Note that the existence of a unique global minimum of J is only rigorously established for linear least squares, when it can be solved for in matrix form to machine precision. In contrast, for non-linear inverse problems[7] one can only aim to find at least one approximate minimum of J that is an *acceptable* fit to the data (i.e. a fit within specified errors).

4.1 Problem formulation

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

$$J(\mathfrak{u}) = \sum_i \alpha_i \times \left(d_i^T \, \mathbf{R}_i^{-1} \, d_i \right) + \sum_j \beta_j \times \left(\mathfrak{u}_j^T \mathfrak{u}_j \right) \tag{12}$$

$$d_i = \mathcal{P}(m_i - o_i) \tag{13}$$

$$m_i = \mathcal{S}\mathcal{D}\mathcal{M}(\mathfrak{v}) \tag{14}$$

$$\mathfrak{v} = \mathcal{Q}(\mathfrak{u}) \tag{15}$$

$$\mathfrak{u} = \mathcal{R}(\mathfrak{u}') \tag{16}$$

[7] The degree of non-linearity may depend on the process of interest and increases substantially upon inclusion of meso-scale eddies.

where d_i denotes a set of model-data differences, α_i the corresponding multiplier, $\mathbf{R}_i^{-1}$ the corresponding weights, $\mathfrak{u}_j$ a set of non-dimensional controls, and β_j the corresponding multiplier. Additional symbols appearing in Eqs. (13)–(16) are defined below. The implementation of Eqs. (12)–(16) and the adjoint interface within the MITgcm is charted in Fig. 5.

Model counterparts (m_i) to observational data (o_i) derive from a set of adjustable model parameters ($\mathfrak{v}$) through the model dynamics ($\mathcal{M}$), diagnostic computations ($\mathcal{D}$), and subsampling or averaging in space and time ($\mathcal{S}$), performed as the forward model steps through time (Eq. 14). Model-data misfits are then computed, upon completion of the forward model simulation, in order to evaluate $J(\mathfrak{u})$ and provide the adjoint model forcing (Sect. 4.2). Raw model-data misfits ($m_i - o_i$) can be penalized directly (i.e. used in Eq. 12 in place of d_i). More generally though, as formulated in Eq. (13), misfits being penalized (d_i in Eq. 12) derive from $m_i - o_i$ through the generic post-processor $\mathcal{P}$ (Sect. 4.3).

The control problem, as implemented in ECCO v4, is non-dimensional, as reflected by the omission of weights in control penalties ($\mathfrak{u}_j^T \mathfrak{u}_j$, Eq. 12). Non-dimensional controls are scaled to physical units through multiplication by their respective uncertainty fields, as part of the generic pre-processor Q (Eq. 15; Sect. 4.4). Pre-conditioner $\mathcal{R}$ (Eq. 16) does not appear in the estimation problem itself (Eq. 12), as it only serves to push an optimization process preferentially towards certain directions of the control space.

The specification of (always approximate) error covariances (e.g. $\mathbf{R}_i$) is a key ingredient of ocean state estimation, and least squares in general. ECCO has contributed a large body of work in this respect (e.g. Forget and Wunsch, 2007; Ponte et al., 2007; Quinn and Ponte, 2008, 2010; Chaudhuri et al., 2013; Forget and Ponte, 2015). Although not a focus in this paper, the difficulty in providing accurate error covariances, and assessing their impact on the state estimate, requires careful analysis of misfit residuals after the fact. This process typically leads to another phase of state estimate production, and so-forth.

For problems as massive as ECCO v4 (see Tables 5–7), full error covariance matrices are impractical and will remain so for the foreseeable future. Matrix free approaches are of great practical value in this context. For example, the method of Weaver and Courtier (2001) is used in ECCO v4 to specify control parameter adjustment scales (Sect. 4.4) and penalize large-scale model-data misfits (Forget and Ponte, 2015).

Within pure linear least-squares theory, under the unrealistic assumption of perfect error covariance specifications, multipliers α_i, β_j should be omitted from Eq. (12). They are, however, adequate in practice as a means to partly compensate for approximations in error covariances, and the neglect of $\mathbf{R}_i$ non-diagonal terms in particular. They also provide a practical means to accelerate the optimization of data sets introduced in J during later stages of optimization. Furthermore, $\mathrm{u}_j^T \mathrm{u}_j$ (in Eq. 12) essentially are regularization terms included to limit control parameters adjustments, and the β_j multipliers provide the corresponding trade-off parameters (Hansen, 1992).

4.2 Adjoint modeling

The method of Lagrange multipliers (i.e. the adjoint method) and its application to numerical models being stepped forward in time is well documented elsewhere. In particular, the interested reader is referred to Thacker and Long (1988) for a succinct presentation, with application to the case of a simple wave equation. The fitting of model sea level variability to altimetry through forcing adjustments estimated by the adjoint method (see, e.g. Forget and Ponte, 2015) is analogous to the simple case treated in Thacker and Long (1988). A crucial advantage of this method, as used in ECCO, is that it avoids adding source/sink terms of unknown nature to the model equations[8]. Adjoint models have many useful applications in their own right, and we shall list a few that are particularly relevant to ECCO.

[8]Note that this desirable property does not hold in the case of incremental (or sequential) data assimilation schemes (whether or not using an adjoint model) but this is not a case of interest here. In particular, it does not hold in 4DVar as practiced in numerical weather prediction.

Integrating adjoint models over extended periods of time allows diagnosis of the sensitivity of model dynamics to various parameters. Two examples are provided in Fig. 6 pertaining to the tracer (left panels) and momentum (right panels) equations that were computed using the "autodiff" (Fig. 5; this section), "profiles", "ctrl" and "smooth" (Fig. 5; subsequent sections) MITgcm packages. Figure 6 illustrates that the sensitivity of model-data misfits (here they cover 2008–2010) extend far back in time (here to 1992). The ability to use information contained in observations backward in time is a powerful advantage of the adjoint method over sequential/filtering assimilation methods. Such adjoint sensitivities provide a practical means to reduce spurious model drifts and biases, through inversion of uncertain model parameters (see, e.g. Ferreira et al., 2005). In cases that are sufficiently linear, adjoint sensitivities to e.g. wind stress can further be convolved with forcing anomalies to reconstruct and attribute variability in the ocean circulation (see, e.g. Fukumori et al., 2015).

Unlike the simple case treated in Thacker and Long (1988), hand-coding the adjoint of the MITgcm would be a very tedious and daunting task. Algorithmic differentiation, through a source-to-source code transformation tool, is a powerful alternative (see Griewank and Walther, 2008). Computational aspects of algorithmic differentiation applied to the MITgcm are described in Heimbach et al. (2005). Since its origin, ECCO has relied on TAMC (Tangent Linear and Adjoint Model Compiler Giering and Kaminski, 1998) and its commercial successor TAF (Transformation of Algorithms in Fortran; Giering et al., 2005). Open source tools such as OpenAD (Utke et al., 2008) and Tapenade (Hascoët and Pascual, 2013) are on their way to provide alternatives for massive problems such as ECCO (Heimbach et al., 2011). The balancing of storage vs. recomputation via the checkpointing method is essential to computational efficiency (Griewank, 1992; Heimbach et al., 2005). This is particularly true for ECCO v4 since the non-linear free surface (see Sect. 3) expectedly increases storage requirements.

During the early development stages of ECCO v4, the adjoint handling of exchanges and storage was extended (partly hand-coded) to allow for elaborate grids such as

CS and LLC (Fig. 1). More generally, development of efficient adjoint code using TAF largely consists in accommodating non-linearities of added forward model features.

Overwhelmingly expensive recomputations of non-linear terms in the adjoint are treated by adding TAF storage directives[9]. These directives take the form of fortran comments (starting with "CADJ") embedded in the forward model code, which TAF transforms into sure code for storage operations (for details, see Heimbach et al., 2005). The ECCO v4 set-up involves 1458 such comments, which were all inserted manually in carefully chosen locations. Once all of the needed storage directives are in place, then "algorithmic differentiation" becomes the "automatic differentiation" that an ECCO v4 user holding a TAF license will experience.

The non-linear free surface, the Adams–Bashforth 3 time stepping scheme, and implicit vertical advection were thus added as adjoint capabilities as part of ECCO v4. Including the non-linear free surface, along with the real freshwater flux boundary condition, in the ocean state estimate is regarded as a major improvement in physical realism. The Adams–Bashforth 3 and implicit vertical advection schemes have a minor impact on the forward model solution but provide additional stability also in adjoint mode.

Exactness and completeness of the adjoint is the general goal of the MITgcm adjoint development. Exactness can be of particular importance to carry quantitative analyses of adjoint sensitivities (e.g. Verdy et al., 2014; Fukumori et al., 2015). For state estimation purposes, however, it is often advantageous, or simply convenient, to use an approximated adjoint (see, e.g. Jiang et al., 2002). The most basic approximation consists in switching off forward model features in the adjoint, which allows postponing the development of a stable adjoint.

In ECCO v4, the Gaspar et al. (1990), Nguyen et al. (2009), and Losch et al. (2010) components are thus omitted in the adjoint. Note that the approximated adjoint does take into account e.g. the diffusivities and viscosities computed by GGL (Gaspar et al.,

[9]TAF adopts a "recompute-all" strategy by default; OpenAD in contrast uses "store-all" by default.

1990). It is only the parametric dependency of these diffusivities and viscosities on the ocean state that is omitted. Until 2008 applications of the MITgcm adjoint were also omitting the Redi (1982) and Gent and Mcwilliams (1990) components, which precluded optimal control of their parameters. This situation was resolved by using a simple clipping scheme for large isopycnal slopes, and by omitting only the parametric dependency of isopycnal slopes on the ocean density field in the adjoint, following a reasoning similar to that of Jiang et al. (2002). Thus, the parametric dependency of turbulent transports on $\mathcal{K}_{\mathrm{gm}}$, $\mathcal{K}_{\sigma}$ and $\mathcal{K}_{\mathrm{d}}$ is retained in the adjoint, so that they can be optimally controlled.

Beyond the removal of unstable adjoint dependencies, other alterations of the adjoint are of practical value for optimization purposes. In particular, it is common practice to increase viscosity parameters to add stability to MITgcm adjoint simulations (Hoteit et al., 2005). Despite successful adjoint simulations with particular versions of the sea ice model (Heimbach et al., 2010; Fenty and Heimbach, 2013), the seaice adjoint is omitted in ECCO v4 due to persisting issues. A pseudo-seaice adjoint is introduced instead to account at least for the most basic effect of seaice – the shielding of sea water from the atmosphere. The adjoint pseudo-component is obtained by AD of a forward pseudo-component. The forward pseudo-component merely tapers air–sea fluxes to zero according to $(1 - a)$ where a is the seaice fraction computed by the actual forward seaice model. This gross, local approximation omits the thermodynamics and dynamics of seaice, and is never used in forward mode. In the adjoint, it masks out open ocean adjoint sensitivities that do not apply where ice cover is present. A fraction of open ocean sensitivity is preserved at the ice edge, which is physically reasonable and avoids a discontinuity in adjoint fields. The pseudo-seaice adjoint approach has been extended in the context of Arctic ice–ocean state estimation (A. Nguyen, personal communication, 2014).

4.3 Observational constraints

Ocean state estimation involves imposing data constraints upon ocean models. Model-data comparison (i.e. computing Eq. 12) becomes an integral part of numerical modeling. In forward mode, "ecco" and "profiles" are diagnostic packages that can be used in any MITgcm run to perform model-data comparisons and to compute Eqs. (12)–(14). In adjoint mode, they take the role of providing the adjoint model forcing (see Fig. 6).

In situ data constraints are handled by the "profiles" package. Model-data comparison is computed at the time-step and grid point nearest to each observed profile (see Appendix D). Aside from the primary goal of carrying out state estimation, the "profiles" output permits direct and rigorous assessments of modeled and observed statistics (and how they may differ) based upon a near identical and instantaneous sampling (e.g., see Forget et al., 2011). To this end, it alleviates the need to output global fields at full temporal resolution, which becomes overwhelming at high spatial resolution.

Gridded data constraints[10] are commonly based upon monthly or daily averaged fields and handled by the "ecco" package. Many features have been added to "ecco" over the course of the ECCO v4 development. In preparation for this paper, these features were generalized so they can immediately be applied, when adequate, to any gridded observational constraints. As of MITgcm's checkpoint65h, the generic "ecco" capabilities are those listed in Table 4.

In general, observable quantities (m_i in Eq. 13) are diagnosed from model state variables (via operator $\mathcal{D}$ in Eq. 14). For potential temperature and salinity ("theta" and "salt" in Table 4) the corresponding model state variables (θ and S in Eqs. 4 and 5) are simply time averaged, and $\mathcal{D}$ then simply denotes the identity operator. In contrast, sea surface height ("eta" in Table 4) is diagnosed as $\eta + \eta_{\text{ips}} + \overline{\eta}_{\text{nbs}}$ where η is the model free surface (see Sect. 3.1), η_{ips} is the weight of sea ice plus snow per unit area divided by ρ_{c} (see Campin et al., 2008), and $\overline{\eta}_{\text{nbs}}$ is a global steric sea level correction to the

[10] By "gridded" we mean either interpolated (e.g. for monthly sea surface temperature) or simply bin averaged (e.g. for along track altimetry).

Boussinesq model (see Griffies and Greatbatch, 2012). Furthermore, for comparison of sea surface height with altimetry, the time mean of $m_i - o_i$ computed at each grid point, and the time variable global mean of $m_i - o_i$, are further subtracted via post-processor $\mathcal{P}$ in Eq. (13) (see Forget and Ponte, 2015).

The basic steps in imposing e.g. a gridded data constraint using the "ecco" package are:

1. Mapping observational data (whether along satellite tracks, gridded, or interpolated) to the model grid, which is easily done e.g. in Matlab using gcmfaces (Appendix C).
2. Specifying the error covariances ($\mathbf{R}_i$ in Eq. 12) of model-data misfits (d_i in Eq. 12). To accommodate the great ocean heteroscedasticity (e.g. see Forget and Wunsch, 2007), spatially varying uncertainties are generally needed.
3. Carrying optimization until convergence to an approximate minimum of $J(\mathrm{u})$.

It should be stressed that all three steps are required to claim that an observational constraint has effectively been imposed on a state estimate, and that the specification of errors is the central scientific problem. This is also true for "profiles" although the first step is limited to a vertical interpolation to standard levels in this case. Table 6 provides the list of gridded observational constraints that, along with the in situ data constraints listed in Table 5, have been imposed on the ECCO-Production, release 1 state estimate (Sect. 5).

4.4 Control parameters

Within the MITgcm, the "ctrl" package (Fig. 5) handles adjustable control parameters (u in Eq. 15). In forward mode, "ctrl" is a package that influences the ocean state evolution (Eq. 14). Activating a new control parameter only requires a few lines of codes to map it to corresponding model parameters (Eq. 15). In adjoint mode, "ctrl" takes the diagnostic role of collecting adjoint variables and evaluating derivatives of Eq. (12) (see Fig. 6).

A penalty can further be added to $J(\mathrm{u})$ by setting $\beta_j > 0$ accordingly (Eq. 12), which will act as an adjoint forcing, to constrain the magnitude of control parameter adjustments.

Most features in "ctrl" were recently generalized so they can readily be applied, when adequate, to any set of controls. The generic pre-processor Q (Eq. 15) may thus include the Weaver and Courtier (2001) spatial correlation model (Appendix E), the cyclic application of climatological mean controls and/or a rotation of (zonal, meridional) vectors to the model C grid. Control parameters used in the state estimate are reported in Table 7.

Most generally, complete and accurate error covariance estimates are lacking for control parameters. For all controls used in the state estimate (Table 7) the error correlation scale was simply specified as 3 times the grid scale using the "smooth" package (as part of $\mathcal{Q}$; Appendix E). The estimation of an initial state that pre-dates Argo and of its uncertainty, given the sparsity of the ship-based ocean sampling, is a difficult problem in itself that is proposed for further, dedicated investigation (e.g., see Forget, 2010; Lyman and Johnson, 2014).

For atmospheric re-analyses fields, in the absence of formal error estimates, ad-hoc specifications of $\mathcal{Q}$ are based upon the spread of available atmospheric variable estimates Chaudhuri et al. (2013). Here the squared sum of time mean and seasonal differences between NCEP and ERA-Interim fields was computed, then capped to a maximum, and used as an ad-hoc estimate of error variances in atmospheric controls.

For $\mathcal{K}_{\mathrm{gm}}$, $\mathcal{K}_{\sigma}$ and $\mathcal{K}_{\mathrm{d}}$, the first guess values were 10^3, 10^3 and $10^{-5}\,\mathrm{m^2\,s^{-1}}$, respectively. The corresponding uncertainties were set to 500., 500. and $10^{-4}\,\mathrm{m^2\,s^{-1}}$. The adjusted parameters were further imposed to stay within $10^2 < \mathcal{K}_{\mathrm{gm}} < 10^4$, $10^2 < \mathcal{K}_{\sigma} < 10^4$, and $10^{-6} < \mathcal{K}_{\mathrm{d}} < 5\times10^{-4}\,\mathrm{m^2\,s^{-1}}$. The $\mathcal{K}_{\mathrm{gm}}$, $\mathcal{K}_{\sigma}$ and $\mathcal{K}_{\mathrm{d}}$ adjustments within the state estimate are assessed more specifically in Forget (2015b). In summary: the estimated $\mathcal{K}_{\mathrm{gm}}$, $\mathcal{K}_{\sigma}$ and $\mathcal{K}_{\mathrm{d}}$ adjustments have a strong impact on ocean stratification and mixed layer depth; these ocean characteristics are now well observed by the Argo program; their estimated and observed maps are in close agreement. These results are evidence

that regional turbulent transport parameter inversions have an observational basis in Argo data.

5 State estimate

The ECCO-Production, release 1 state estimate covers the period from 1992 to 2011 and is the baseline solution of the ECCO v4 forward model setup (Sects. 2 and 3), including the parameter adjustments derived from observational data constraints (Sect. 4). Its monthly mean output and model-data misfits are publicly available online. Several recent publications (Speer and Forget, 2013; Wunsch and Heimbach, 2013b, 2014; Buckley et al., 2014; Forget and Ponte, 2015; Balmaseda et al., 2015) analyze earlier iterations (see Appendix G). The solution fits altimetry (Forget and Ponte, 2015), SST (Buckley et al., 2014) and subsurface hydrography data (Sect. 5.2) at or close to the specified noise level. Many characteristics of the solution have been analyzed in some detail and found physically plausible, which warranted its public release. The reader is further referred to the extensive documentation (the "standard analysis" provided as Supplement) of model-data misfits and physical characteristics of the state estimate that is also publicly available online.

5.1 Select characteristics

The characteristics in Table 2 are a small subset that is representative of the multifaceted nature of ocean state estimation. To shed light on the observational and climate problems, this section assesses the sensitivity of these characteristics, and correlations amongst characteristics, as measured within Table 3 (model settings) and within Table 8 (adjusted controls). The different levels of sensitivity seen in Table 8 vs. Table 3 is addressed in Sect. 5.3, and pertains to controllability rather than observability.

The various observational constraints (the first seven characteristics) show contrasting levels of sensitivity to control parameter adjustments (Table 8). The same is true for

discrete model settings (Table 3). This behavior may reflect different levels of random errors amongst the different types of observations. In particular, the subsurface hydrography, as constrained by jT and jS, appears as the most sensitive model-data distance (Tables 3–8). Another noteworthy result is that global mean time series show more spread than do time averaged meridional transports (Tables 3–8 reporting normalized differences). The main exception to this behavior is for the meridional salt transport, whose time average is small in the state estimate (Fig. 7).

The correlation (or lack thereof) between columns of Tables 3–8 also yields salient conclusions (Fig. 8). High correlations between observational constraints (bottom panels) is suggestive of some redundancy between data sets (i.e. consistency amongst observations). High correlations between meridional transports and observational constraints (top and middle right panels) provides evidence that Argo and altimetry may efficiently constrain heat and freshwater transports (see also, e.g. Forget et al., 2008a, b).

In contrast, low correlations between global mean time series and distances to observations is striking (Fig. 8, top and middle left panels). Given that the time variable global mean model-data difference is omitted in computing jHa, the low correlation between mH and jHa indicates that a given global mean sea level time series could be associated with many regional solutions with equal uncertainty. The low correlation between jT and mT may further reflect that regional variations can be much larger than, and not necessarily related to, temporal changes in global mean properties.

Beyond the present study, the extent to which Argo and altimetry, amongst others, constrain temporal changes in global mean properties remains unclear. A related concern is that global mean properties are the most sensitive to discrete choices in model numerics and physics amongst the selected characteristics (see Table 3). It is tempting to attribute this behavior to the omission of atmospheric, continental, etc. hydrology modeling in ECCO v4, although this remains to be proven. Whether and how the large sensitivity of global mean properties seen in Fig. 7 translates into simpler models used to quantify climate change from observations (as in, e.g. Purkey and Johnson, 2010;

Llovel et al., 2014) emerges as a question of direct relevance to climate change monitoring.

5.2 Improved hydrography fit

In developing and producing the state estimate, a primary goal was to improve the fit to observed in situ profiles (of T and S) as compared with earlier ECCO estimates (see Forget, 2010). As already apparent in Table 3, the inclusion of parameterized physics as controls (i.e. adjustable model parameters) was instrumental in achieving that goal. The present section focuses on this defining characteristic of the baseline ECCO v4 solution.

The fit to in situ profiles of temperature and salinity is depicted in Fig. 9 as a function of time, for ECCO v4 and earlier MITgcm solutions. For ECCO v4 the model-data distance for in situ profiles (see Sect. 4.3) is $jT \approx 1.5$ for potential temperature and $jS \approx 1.5$ for salinity (on average over all depths, locations and times). Average values of 1 would be ideal if the error estimate (see Forget and Wunsch, 2007) was perfect and the state estimate was devoid of large scale errors (neither of which is true). It is suspected that jT and jS could be further reduced. Values of 1.5, however, are regarded as sufficiently low to justify analysis of the state estimate water masses (Speer and Forget, 2013) and stratification (Forget, 2015b). Furthermore, jT and jS are already much reduced (by a factor of 2 to 10) compared with earlier ECCO estimates (Fig. 9) throughout the period from 1992 to 2011. Amongst earlier ECCO estimates, the ECCO v3 solution comes the closest to the observed hydrography with a typical distance to observations of 3 (Fig. 9).

The contrasts in jT and jS amongst solutions (Fig. 9) reflect large scale misfits as illustrated in Fig. 10. This is equally true for ECCO2 eddying solutions (bottom panels) and for coarser model solutions (top and middle panels). Such broad misfit patterns typically denote spurious model drifts and biases, which are common symptoms of model deficiencies (Stammer, 2005; Ferreira et al., 2005). Similarities in misfit patterns amongst ECCO2 eddying solutions (using a common model set-up, under different sets

of forcing) for example suggest internal ocean model deficiencies. So do similarities in misfit patterns (aside from differences in amplitudes) amongst the four adjoint optimized solutions of comparable resolution that use different adjusted forcing fields (ECCO v2, v3, v4 and GECCO2).

The contrast in misfit amplitude between ECCO v4 and earlier solutions (Figs. 9 and 10) tends to be reduced near the sea surface (not shown), which is encouraging but not entirely surprising since surface forcing fields were already control parameters in earlier solutions. Conversely the contrast in misfit amplitude tends to increase with depth (not shown), where internal model error sources may predominate.

Within ECCO v4, jT and jS are particularly sensitive to estimated turbulent transport parameter adjustments and generally less sensitive to estimated surface forcing adjustments, with the exception of expectedly high salinity sensitivity to precipitation (see Table 8, first two columns). This result is in contrast with the analysis of Liu et al. (2012) who suggest that parameterized physics are only marginally important in this regard, a suggestion consistent with the relative weakness of their turbulent transport parameter adjustments (see Forget, 2015b). A plausible explanation for this contrast lies in the fact that Liu et al. (2012) only estimate the period 1992–2001, whereas ECCO v4 covers 1992–2011. This difference has two important implications: (1) argo largely increased the volume of in-situ constraints, (2) slow models drifts are more prominent in longer unconstrained solutions. One should expect larger turbulent transport parameter adjustments on both counts.

Amongst turbulent transport control parameters in ECCO v4, jT and jS are most sensitive to the K_{gm} adjustments (this result is in agreement with Liu et al., 2012). A caveat should be noted though: parameterized surface and interior fluxes are all interactive so that any control vector adjustment can potentially affect any surface or interior flux. Hence Table 8 should not be mistaken for a precise ranking of the importance of the various controls. It clearly shows, however, that turbulent flux parameter adjustments were instrumental in fitting observed hydrographic profiles in ECCO v4.

5.3 Parametric and structural model error

In this section, the focus is on model uncertainty and controllability, which directly impacts the possibility of fitting a model to observed data. Random data errors and model representation errors are left out of the discussion, which are comparatively well studied (e.g. Forget and Wunsch, 2007; Ponte et al., 2007; Quinn and Ponte, 2008, 2010; Chaudhuri et al., 2013; Forget and Ponte, 2015). Errors associated with computing environment changes (top three rows in Table 3) are generally small enough to be neglected when using the MITgcm.

The interplay of external, structural and parametric ocean model errors has never been tackled in any systematic and quantitative manner. To distinguish amongst model uncertainties associated with ECCO v4 settings, we propose the simple, practical category definitions in Table 9. Clearly the separation between these three categories leaves room for ambiguities. For example, selecting one of the available atmospheric re-analysis products to force the model may fall under "structural", while tuning bulk formulae coefficients may fall under "parametric" and adjusting re-analyzed fields may fall under "external". Nevertheless, as a starting point, the above definitions provide a useful frame of reference. A related discussion can be found in Marzocchi and Jordan (2014), although the focus here is on curve fitting (i.e. interpolation within a time period) rather than on forecasting (i.e. extrapolation forward in time). Relevant discussion can also be found in Danabasoglu et al. (2014) and Balmaseda et al. (2015).

A first assessment of the relative importance of external, parametric and structural model uncertainty in ECCO v4 can then be made from Table 3 (structural sensitivity tests) and Table 8 (external and parametric sensitivity tests). Structural model uncertainty associated with choices of advection, mixed layer and momentum schemes are sizable over 20 years (Table 3). Solutions perturbed by this much are sufficiently distinct from the state estimate to prompt further optimization leading to a different state estimate. The most important result, however, may be that adjusted control parame-

ters generally have a much larger impact (Table 8) than switching amongst numerical schemes (Table 3).

A ratio C of model uncertainty controlled by continuous parameters (external or parametric) to structural model uncertainty is introduced to better illustrate this result (Fig. 11). The adjoint method allows for reduction of parametric and external errors, but it does not lend itself to reduction of structural errors that are fundamentally discontinuous. Hence, C is an index of model controllability, which can be interpreted as a *signal to noise ratio* of sorts, but for model simulations rather than observations. Large values of C are a priori favorable to state estimation.

It is therefore encouraging that $\log_{10}(C) > 0$ for all variables considered (Fig. 11), showing that controlled model uncertainty exceed the *noise level* set by structural model uncertainty. Certain ocean characteristics are particularly prone to structural model uncertainty, whereas others are highly controllable. On the one hand, model-data distances for regional sea level variability and in situ hydrography appear most controllable with $\log_{10}(C) > 1.5$ (top panels). On the other hand, global mean temperature and sea surface salinity appear most prone to structural model uncertainty with $\log_{10}(C) < 0.5$. The high level of structural uncertainty seen in global mean heat uptake (i.e. mT) is cause for concern in the context of climate change monitoring (see also Sect. 5.1).

Increasing model controllability is a priori favorable to state estimation. To this end, one may seek to replace discrete choices and switches with continuous parameter specifications that enable smooth state transitions[11], or simply add adjustable parameters[12]. The replacement of the C-D scheme by optional targeted viscosity, and the replacement of KPP with GGL (Sect. 3.3) thus aim at increasing model controllability. For example, KPP is a very complex and non-linear parameterization that involves many discrete switches and thousands of code lines. GGL yields broadly similar results to KPP over 20 years (Table 3) and is in contrast a very simple code, so that a practical adjoint may be within reach. It is also noteworthy that activating the C-D scheme generally trumps the impact of switching between mixed layer schemes, albeit with the notable exception of global mean characteristics (see Table 3). This result highlights the potential benefits of further extending the inversion problem to viscosity parameters.

[11] At this point it is assumed, for the sake of a simple preliminary discussion, that an *expert consensus* could be reached to exclude certain discrete numerical schemes (see Marzocchi and Jordan, 2014).

[12] If algorithmic differentiation is the method of choice to this end, then schemes that have fewer discrete switches are preferable over other comparable schemes.

5.4 Known issues

State estimation should aim towards universality and completeness (see Wunsch and Heimbach, 2013a, for a review). Thus, its practice always warrants continuous improvement in many respects. In ECCO v4, without trying to be exhaustive, one can distinguish at least three types of issues.

Firstly, the ECCO-Production, release 1 state estimate would benefit from further optimization, with additional data sets, controls, and refined error covariance specifications. The addition of turbulent transport parameters is a step forward, but their specified covariances remain very imprecise. Parametric error in the momentum equations also deserves further attention, since it may limit model controllability. Error covariances between adjustable control parameters (e.g. atmospheric variables) are also neglected. A permanent issue is the need for additional data constraints, particularly in the abyss (Wunsch and Heimbach, 2014). Amongst available data that is not yet in ECCO v4, the growing bio-geochemistry data base is becoming a priority.

Secondly, the lack of "posterior" error estimates is regarded as the most outstanding issue with ECCO-Production, release 1. Producing formal error estimates, at a reasonable computational expense and with acceptable precision, for the full, evolving ocean state would be another major breakthrough. In principle, a number of methods are available to this end. In practice, however, most of them are intractable for problems of size $> 10^8$ (sizes are reported in Tables 5–7). One approach that is being pursued is the use of second derivative (Hessian) information that, under the assumption of Gaus-

sian distribution, can be readily related to the posterior error covariance (see Kalmikov and Heimbach, 2014). Also a possibly useful estimate of uncertainty in ECCO v4 may follow from computing the spread amongst available ocean data syntheses, although it is unclear how such ensemble spreads should be interpreted (Balmaseda et al., 2015).

Thirdly, the ECCO v4 model setup could be extended and improved, with possibly important implications for the state estime. The lack of atmospheric, land, and biogeochemistry components is an obvious limitation of ECCO v4 at this stage. The surface boundary conditions and seaice model settings require further assessment. Issues such as the use of the Boussinesq approximation (in Eqs. 1–5), the omission of geothermal heating (Piecuch et al., 2015), and the lack of a coastal wetting/drying mechanism are matters for further MITgcm development that are also of importance to state estimation.

6 Conclusion and perspectives

This paper emphasizes the synergy between ocean modeling and data analysis. The entanglement of models and observations is nothing new – Ekman (1905), Sverdrup (1947),Munk (1966) and Wunsch (1977) are just a few historical examples. The synergy of ocean modeling and data analysis is further becoming a reality as a growing community engages in ocean state estimation, which in essence is the hybridization of ocean modeling and data analysis. What is different now merely is the level of (in)completeness, complexity, and diversity of the models and observations being employed in modern oceanographic and climate science. The scope and size of the ocean state estimation problem tackled in ECCO v4 requires collaborative research and production activities. This unescapable conclusion leads to this attempt at offering ECCO v4 as a fully integrated framework for non-linear inverse modeling and global ocean state estimation. Along with the MITgcm and its adjoint capability, the ECCO v4 framework currently includes the components listed in Table 10.

Each component of the framework is being (re)designed to be modular and of general applicability, as they all are thought to provide valuable stand-alone pieces to different degrees. Standardized in-situ data sets in particular, while a by-product of carrying out ECCO v4, allow for a variety of scientific analyses in their own right. For example they are used for analyses of observed variance that is never fully represented in numerical model solutions (Forget and Wunsch, 2007; Forget, 2015a), of water masses volumetric census (Forget et al., 2011; Speer and Forget, 2013), and of macro turbulence (McCaffrey et al., 2015) and mixing (Forget, 2015b). A complementary description of the standardized in situ observations and related ECCO v4 components is provided in Appendix, directed towards users of in-situ observations.

As another example, the gcmfaces Matlab framework (Appendix C) is suitable for the analysis of gridded earth variables (whether observational or modeled) beyond the ECCO v4 model setup and state estimate. At this stage it has already been applied to analyze MITgcm simulations on various grids, and to a variety of observational data sets. Interfacing gcmfaces with output from models other than MITgcm would allow for rigorous model intercomparisons without the need to introduce errors through interpolation. As a final example, any interested modeling group should be able to take advantage of the global grids.

The state estimate and the MITgcm are highly integrated with each other. Beyond the few aspects of the solution that have been investigated in some detail, the MITgcm provides numerous prognostic and diagnostic capabilities that remain to be applied to, or employed within, ECCO v4. The “ctrl”, “ecco” and “profiles” packages, are just examples of the many MITgcm packages. The last two diagnose model-data misfits and statistics. In contrast, the “ctrl” package defines control parameters that act upon the forward prognostic equations. It also lends itself to development of new parameterizations. Note that the roles of these packages (diagnosing or acting on the solution) are reversed in the adjoint. Amongst forward prognostic MITgcm packages not yet used in ECCO v4, biogeochemistry and simplified atmospheres (Dutkiewicz et al., 2005; Follows et al., 2007; Marshall et al., 2007a; Ferreira et al., 2011) are worth singling

out, as they offer a great potential for extending ocean state estimation. The adjoint capabilities of MITgcm further allow for computations of sensitivity, Green functions, singular value decomposition, mechanistic attribution of variability, optimal observation design (Marotzke et al., 1999; Köhl and Stammer, 2004; Fukumori et al., 2007, 2015; Heimbach et al., 2011; Zanna et al., 2011).

Furthermore, the MITgcm provides a convenient platform for parallel computing and variational estimation that allows for, but is not limited to, ocean data synthesis and analysis (Hoteit et al., 2013; Goldberg and Heimbach, 2013). Optimal interpolation (OI) of an individual variable, for instance, can readily be carried out using Eq. (12) and its adjoint with $\mathcal{M} = I$ (i.e. the identity operator) as illustrated by Forget (2010). In between OI and full ocean state estimation, and beyond, lie many interesting stages and possibilities. For instance, stand alone bulk formulae configurations (available at mitgcm.org, with or without seaice) could readily allow for assessment and optimization of air–sea fluxes (along the lines of, e.g., Yu and Weller, 2007; Maze et al., 2009). The (re)implementation of Eq. (12) within MITgcm provides a versatile environment for such projects, and for variational estimation purposes most generally (and is complementary to, e.g. Barth et al., 2014; Wilson et al., 2014; Hoppe et al., 2014).

It is expected that any of the ECCO v4 components listed in Table 10 will eventually be replaced. Most immediately, the specifics of the ocean state estimation problem (grid, forcing, ocean and seaice model settings, control parameters, observational constraints) can all be refined or substituted for improved components. Our continued commitment is to make every updated component freely and fully available online as soon as possible. All of the Fortran and Matlab components are already available, and served through the CVS server of MITgcm, where they were added in real time and with free access over the years (Sect. “Code availability”). The state estimate monthly output and the model-data distance (data, model counterparts, and uncertainty) for in situ profiles are also readily available. The rest of the numerical input and output requires additional processing and web interfacing – and is for now instead made available upon email request (Sect. “Code availability”).

Furthermore, at the present time, taking full advantage of the ECCO v4 framework (Table 10) requires two third party commercial tools that are neither free nor open source: Matlab and TAF. The ability to successfully generate efficient adjoint code using alternative open-source tools, such as OpenAD or Tapenade is gaining increasing priority. Despite its limitations, Matlab is one of the most portable, integrated and popular analysis framework, and it is expected to remain as such for the foreseeable future. However, a Python analysis framework similar to gcmfaces is in planning and should better handle massive output from high resolution models (R. Abernathey, personal communication, 2014).

Gridded observational products (such as hydrography climatologies, ocean state estimates, etc.) are commonly used as a practical shorthand to observations. It should be stressed that a gridded field in itself does not provide any information about its errors. Therefore, and since direct observational constraints are unevenly distributed and restricted to a few variables, users of the state estimate are strongly encouraged to consider the underlying observational data base. This being said, and despite the need for continued improvement, the usefulness and scientific value of the ECCO v4 solution is by now largely documented in a number of papers (Speer and Forget, 2013; Wunsch and Heimbach, 2013b, 2014; Buckley et al., 2014, 2015; Forget and Ponte, 2015; Forget, 2015a, b; Liang et al., 2015; Fukumori et al., 2015; Balmaseda et al., 2015).

As compared with earlier ECCO solutions, the ECCO-Production, release 1 state estimate (i.e. the baseline ECCO v4 solution) benefits from an extensive revisit of model settings. The improved fit to in situ hydrography (Argo profiles of T and S in particular) as compared with earlier ECCO solutions may be the defining characteristic of ECCO-Production, release 1. The inclusion of turbulent transport parameters in the set of adjustable control parameters was instrumental in achieving that goal – their inversion from hydrography observations is further assessed in Forget (2015b). Nevertheless, it should not be assumed that broad scale misfits to observations are completely absent (e.g. see Fig. 10). Users of the state estimate are expected to question its realism, while

being provided with capabilities to assess model-data misfits for themselves. More generally, it should not be assumed that all ocean state variables are fully constrained by observations. Integrated transports, global averages, etc. are not directly observed, and it is a priori unclear how well they can be constrained by available observations (see Forget et al., 2008a, b; Heimbach et al., 2009; Forget, 2015b).

Looking to the future, the need for associating formal error estimates with the full, evolving ocean state remains of utmost importance. Aside from this aspect, extensions of the state estimation framework to include other climate components (atmosphere, land, cryosphere) and different variables (biology, chemistry) would be desirable (see, e.g. Blessing et al., 2014; Prinn et al., 2011). By providing ECCO v4 as a fully integrated framework along with a useful baseline solution that any interested investigator should be able to reproduce for the foreseeable future, the authors aim to stimulate independent research along those lines.

The overarching scientific problem (set aside technicalities) to data-model combination lies in the attribution of errors amongst the various elements of Eq. (12). We make no claim to having achieved the proper attribution of errors, but experience gathered in developing ECCO v4 suggests that a paradigm shift, as compared with earlier ECCO publications, is in order. Our results indeed indicate that internal parameters are of first order importance to state estimation, and to fitting the observed hydrography in particular (Table 8). Our assessment is in contrast with that of Liu et al. (2012) who suggest that the importance of internal parameters is of order 10–20 % depending on the model variable of interest. Furthermore the inversion of parameters in the momentum equations, which has received comparatively little attention, emerges as a topic of importance as one gets closer to observed data, and is expected to gain further importance as resolution increases. To provide a frame of reference for future research along those lines, a first attempt at defining and gauging various categories of model uncertainty has been presented.

Alleviating structural model errors is a prerequisite to improved dynamical interpolation of observations. In this regard, the main improvement compared with previous

ECCO estimates may be the extension of the gridded domain to the Arctic, the addition of the non-linear free surface, and the switch to real freshwater flux (Sect. 3). These specific *expert choices* (Marzocchi and Jordan, 2014) should not be controversial. For many other model settings, the situation is not so clear but structural model errors are generally regarded as a more difficult issue than parametric model errors. Indeed, structural model errors by definition consist of fundamentally discontinuous modeling choices that cannot be optimally controlled. Thus, structural model errors fundamentally are a matter of *expert choices* (Marzocchi and Jordan, 2014). In contrast, sensitivity to continuous parameters can readily be probed in adjoint mode (Sect. 4.2) so that they can be estimated objectively under the constraint of observations (Sect. 5.2).

Parametric and external model uncertainty (Table 8) generally appear to dominate over structural model uncertainty (Table 3) as illustrated by Fig. 11. Such a conclusion most likely depends on spatial resolution, the chosen 20 year duration, and the necessarily limited array of model settings being considered in Tables 3–8. In particular, we expect that the choice of momentum schemes would be more important in eddy-resolving models, as kinetic energy overcomes potential energy at the meso-scale. Examples of large structural uncertainty in eddy permitting models can be found in Barnier et al. (2006) and subsequent studies. Here, however, the estimated control parameter adjustments appear to determine the solution beyond the level of structural model uncertainty (Sect. 5.3).

Parametric model uncertainty (associated here with interior turbulent transports) and external model uncertainty (associated here with surface forcing fields) appear to be of comparable magnitude (Table 8). Depending on the characteristic of interest, one predominates over the other. Hence, the importance of including turbulent transport parameters, which are highly uncertain, in the control vector cannot be overstated.

Appendix A: Grid generation method

At high-latitude, the LLC mesh is generated numerically by adapting the two dimensional conformal mapping algorithm developed by Zacharias and Ives in the 1980s (see Ives and Zacharias, 1989; Trefethen, 1989; Wilkin and Hedström, 1998) to spherical geometry. The approach is similar to that used in the SeaGRID package (Denham, 2000), except that here spherical polar coordinate geometry defines sub-domain boundaries. The numerical mesh is generated separately for the Arctic Cap and the transition sector. Each quarter of the transition sector is bounded by: the 57° N parallel (southern edge), two 90° spaced meridians (eastern and western edges), and a small-circle arc that crosses the eastern and western edges at 67° N (northern edge). The four northern edges of the transition sector bound the Arctic cap.

To numerically mesh each sub-domain it is first conformally projected onto a plane, using a polar stereographic transformation. The result is then conformally mapped to a rectangular shape by iteratively applying the so-called "hinge-point" or "power" transformation to each of the four arc segments that make up the sub-domain edges. The transformation works with points (x,y) in the complex plane $x+iy$ and applies the mapping $\omega = (x+iy)^P$. The transformation is applied iteratively to adjacent pairs of discrete line segments that define the sub-domain edges. The transformation adjusts P at each iteration for successive line segment pairs, so that the angle between adjacent segments is adjusted to be $\frac{\pi}{2}$ at corners and π for all intermediate segments.

The result of the transformation is a rectangular shape in a new coordinate space denoted by coordinates ζ and η. The rectangular shape has two edges that are line segments of constant ζ and two edges that are lines of constant η. The points that define the line segments have corresponding mappings to the line segment points in the original (x,y) coordinate system. A set of x and y locations that describe orthogonal grid lines in the sub-domain interior can then be generated numerically; by solving two

Laplace equations (Ryskin and Leal, 1983) of the form

$$\frac{\partial^2 X}{\partial \zeta^2} + \frac{\partial^2 X}{\partial \eta^2} = 0$$
$$\frac{\partial^2 Y}{\partial \zeta^2} + \frac{\partial^2 Y}{\partial \eta^2} = 0$$

over the (ζ,η) rectangular shape and subject to the respective boundary conditions $X = x$ and $Y = y$ on the respective ζ = constant and η = constant rectangular shape edges.

Appendix B: Time stepping

The time-discretized version of Eqs. (1)–(5) and (7) calculate the updated state $(\boldsymbol{v}^{n+1}, w^{n+1}, \eta^{n+1}, \theta^{n+3/2}, S^{n+3/2})$ at time $t+\Delta t$ from the current state at time t $(\boldsymbol{v}^{n}, w^{n}, \eta^{n}, \theta^{n+1/2}, S^{n+1/2})$ following:

$$(\Phi')^{n+1/2} = \frac{g}{\rho_c}\int_z^{\eta^n} (\rho')^{n+1/2} \mathrm{d}z \quad \text{with: } (\rho')^{n+1/2} = \rho(\theta^{n+1/2}, S^{n+1/2}, -\rho_c g z^*) - \rho_c \tag{B1}$$

$$\frac{\boldsymbol{v}^{n+1} - \boldsymbol{v}^{n}}{\Delta t} - \left[\mathbf{G}_{\boldsymbol{v}}^{n}\right]^{\mathrm{AB}} + g\nabla_{z^*}\eta^{n+1} + \nabla_h \Phi'^{(n+1/2)} = \mathbf{D}_{z^*,\boldsymbol{v}}^{n} + \mathbf{D}_{\perp,\boldsymbol{v}}^{n+1} + \mathcal{F}_{\boldsymbol{v}}^{n+1/2} \tag{B2}$$

$$\frac{\eta^{n+1} - \eta^{n}}{\Delta t} + \nabla \cdot \int_{-H}^{\eta^n} \boldsymbol{v}^{n+1} \mathrm{d}z = \mathcal{F}^{n+1/2} \tag{B3}$$

$$\frac{1}{H}\frac{\eta^{n+1} - \eta^{n}}{\Delta t} + \nabla_{z^*}(s^{*n}\boldsymbol{v}^{n+1}) + \frac{\partial w^{n+1}}{\partial z^*} = s^{*n}\mathcal{F}^{n+1/2} \tag{B4}$$

$$\frac{s^{*n+1}\theta^{n+3/2} - s^{*n}\theta^{n+1/2}}{\Delta t} - \mathcal{A}\left(\theta, \boldsymbol{u}^{n+1} + \boldsymbol{u}_{\mathrm{b}}\right) = s^{*n}\left(\mathcal{F}_{\theta}^{n+1} + D_{\sigma,\theta}^{n+1/2} + D_{\perp,\theta}^{n+3/2}\right) \quad \text{(B5)}$$

$$\frac{s^{*n+1}S^{n+3/2} - s^{*n}S^{n+1/2}}{\Delta t} - \mathcal{A}\left(S, \boldsymbol{u}^{n+1} + \boldsymbol{u}_{\mathrm{b}}\right) = s^{*n}\left(\mathcal{F}_{S}^{n+1} + D_{\sigma,S}^{n+1/2} + D_{\perp,S}^{n+3/2}\right) \quad \text{(B6)}$$

where $\boldsymbol{u}$ represents the three components velocity vector (u, v, w), $\boldsymbol{u}_{\mathrm{b}}$ the bolus velocity and $\mathcal{A}(\,)$ denotes the advection term.

Momentum advection and the Coriolis term are evaluated at time t from $\boldsymbol{v}^n, w^n$ in $\mathbf{G}_{\boldsymbol{v}}^n = -(f+\zeta)\hat{\boldsymbol{k}}\times\boldsymbol{v} - \nabla_{z^*}\mathrm{KE} - w\frac{\partial \boldsymbol{v}}{\partial z}$ and the resulting tendency ($\mathbf{G}_{\boldsymbol{v}}^n$) is extrapolated forward in time to $t + \Delta t/2$ using the Adams–Bashforth 3 (AB-3) scheme:

$$\mathbf{G}_{\boldsymbol{v}}^{n+1/2} = \left[\mathbf{G}_{\boldsymbol{v}}^n\right]^{\mathrm{AB}} = (1 + \alpha_{\mathrm{AB}} + \beta_{\mathrm{AB}})\,\mathbf{G}_{\boldsymbol{v}}^n - (\alpha_{\mathrm{AB}} + 2\beta_{\mathrm{AB}})\,\mathbf{G}_{\boldsymbol{v}}^{n-1} + \beta_{\mathrm{AB}}\,\mathbf{G}_{\boldsymbol{v}}^{n-2}$$

Here we use $(\alpha_{\mathrm{AB}}, \beta_{\mathrm{AB}}) = (1/2, 0.281105)$ to improve the stability (Shchepetkin and McWilliams, 2005) compared to the true 3rd order in time Adams–Bashforth $(\alpha_{\mathrm{AB}}, \beta_{\mathrm{AB}}) = (1/2, 5/12)$. The precision of the scheme drops to just 2nd order accuracy with little consequences here since most of the other terms are also 2nd order in time (tracer time-stepping, internal-waves dynamics). Note that the precision is still improved compared to the quasi-AB-2 used in previous ECCO configurations which become only 1rst order accurate with stabilization factor ($\epsilon_{\mathrm{AB}} \sim 0.1$).

Simple eulerian time stepping (1st order, forward in time) is used in $\mathbf{D}_{z^*,\boldsymbol{v}}^n$ for horizontal dissipation (harmonic and bi-harmonic viscosity) and quadratic bottom drag. Using a quasi-AB-2 scheme instead (as in previous ECCO configurations) would reduce the stability limit from 1 to 0.9 (for pure damping term, with $\epsilon_{\mathrm{AB}} = 0.1$). AB-3 would reduce it even further to 0.55 and therefore was not considered here. For stability reason also, a backward time stepping is used for the other dissipation term in Eq. (B2) (i.e. $\mathbf{D}_{\perp,\boldsymbol{v}}^{n+1}$) that represents vertical viscosity effects in the interior, except bottom friction:

$$\mathbf{D}_{\perp,\boldsymbol{v}}^{n+1} = -\partial/\partial z\left(-\nu_{\perp}\partial/\partial z(\boldsymbol{v}^{n+1})\right)$$

This vertical shear term is independent of the vertically integrated pressure gradient contribution ($g\nabla_{z^*}\eta^{n+1}$), so that these two operations commute. This allows to find $\mathbf{D}_{\perp,\boldsymbol{v}^{n+1}}$ even before knowing η^{n+1} by solving a tri-diagonal system in each water column.

The updated η^{n+1} is found by combining Eqs. (B2) and (B3) to form a 2-D elliptic Poisson equation for surface pressure (pressure method) which is solved iteratively using the conjugate-gradient method (Marshall et al., 1997). Solver matrix and preconditioner are updated at each time-step as the water column height changes due to the non-linear free-surface (Campin et al., 2004).

The tracer Eqs. (B5) and (B6) contain several subgrid-scale (SGS) terms within $D_{\perp}, D_{\sigma}$ that can use different time-stepping methods. They represent small scale vertical mixing ($K_{\perp}$) due to a time-invariant background diffusivity field ($\mathcal{K}_d$; Sect. 3.3) and time-variable contributions from GGL (Gaspar et al., 1990), as well as isopycnal diffusion ($\mathcal{K}_{\sigma}$; Sect. 3.3). The effect of unresolved eddies parameterized as a bolus velocity ($\boldsymbol{v}_{\mathrm{b}}$) advecting tracers (Gent and Mcwilliams, 1990) is included in $\mathcal{A}(\theta, \boldsymbol{v} + \boldsymbol{v}_{\mathrm{b}})$. All SGS parameters, including $\boldsymbol{v}_{\mathrm{b}}$, isopycnal slope (α_x, α_y) and vertical diffusivity and viscosity $(K_{\perp}, \nu_{\perp})$ are computed at the beginning of the time-step from the current state.

Isopycnal diffusivity ($\mathcal{K}_{\sigma}$) is discretized as a tensor (Redi, 1982) where all the terms are treated explicitly (i.e., as a function of $\theta^{n+1/2}$ gradient) except for the pure vertical component $\frac{\partial}{\partial z}(|\alpha|^2\mathcal{K}_{\sigma}\frac{\partial(\theta)}{\partial z})$ where $|\alpha| = (\alpha_x^2 + \alpha_y^2)^{1/2}$ denotes the magnitude of the isopycnal slope. The pure vertical component is combined with $K_{\perp}$ and applied to the future tracer field ($\theta^{n+3/2}$) using a backward time-stepping, leading to:

$$s^{*n}\left(D_{\sigma,\theta}^{n+1/2} + D_{\perp,\theta}^{n+3/2}\right) = -\nabla_{\sigma}\left(-\mathcal{K}_{\sigma}\nabla_{\sigma}\theta^{n+1/2}\right) - \frac{\partial}{\partial z^*}\left(-(K_{\perp} + |\alpha|^2\mathcal{K}_{\sigma})\frac{\partial\theta^{n+3/2}}{\partial z}\right)$$

Rather than to evaluate bolus advection $\mathcal{A}(\theta, \boldsymbol{u}_{\mathrm{b}})$ separately from the eulerian advection, the 3 components residual mean velocity is formed $\boldsymbol{u}_{\mathrm{res}}^{n+1} = \boldsymbol{u}_{\mathrm{b}} + \boldsymbol{u}^{n+1}$ and used to

advect tracers, per:

$$\frac{s^{*n+1}\theta^{n+3/2} - s^{*n}\theta^{n+1/2}}{\Delta t} + \nabla_{z^*}(s^{*n}\theta^{n+m/2}\boldsymbol{v}_{\text{res}}^{n+1}) + \frac{\partial(\theta^{n+3/2}w_{\text{res}}^{n+1})}{\partial z^*} = \tag{B7}$$
$$s^{*n}\mathcal{F}_\theta^{n+1} - \nabla_\sigma\left(-\mathcal{K}_\sigma\nabla_\sigma\theta^{n+1/2}\right) - \frac{\partial}{\partial z^*}\left(-(K_\perp + |\alpha|^2\mathcal{K}_\sigma)\frac{\partial\theta^{n+3/2}}{\partial z}\right)$$

Horizontal advection (second term in Eq. B7) uses the 3rd order direct space and time (DST-3) advection scheme (MITgcm Group, 2002; Adcroft et al., 2004b) with the direction splitting method (also called multi-dimensional advection) as described in Adcroft et al. (2004b). The tracer field ($m = 2$) obtained after applying 1-D advection (in X or Y direction) on current tracer ($m = 1$) is used to compute the advective fluxes in the other direction (Y or X) and ensures 2nd order accuracy in space and time. Regarding vertical advection, the backward time stepping (unconditionally stable) is applied with 3rd order advection scheme; this involves solving a penta-diagonal system (with some additional contributions from vertical mixing to the 3 main diagonals) for each column. In particular, this choice alleviates adjoint stability restrictions.

Appendix C: Diagnostics

The MITgcm "diagnostics" package is generally used to generate binary output for offline analysis of the solutions. In the case of the LLC90 grid, a two-dimensional field is thus output as an array of size 90 × 1170. It can easily be re-organized according to Table 11 to match the MITgcm layout of the LLC90 grid (Fig. 12). The state estimate output is made available online in a tiled netcdf format (nctiles) where each tile is a 90 × 90 subdivision of a face (i.e. of f1, f2, f3, f4 or f5 in Table 11) and is written to an individual netcdf file.

The need for nctiles files stems from the fact that there is no simple, robust and general way to re-arrange global model output in a single two-dimensional array. For

LLC fields, it is only the LL sector that can readily be re-assembled as a single two-dimensional array. A simple Matlab script is provided to this end (eccov4_lonlat.m; see Sect. "Code availability"). It is mainly intended for users of earlier non-global ECCO estimates that may want to re-use their old analysis codes. ECCO v4 users are generally advised against interpolating, which introduces errors, and often precludes accurate transport computations. Instead, mimicking the gridded earth decomposition of general circulation models is regarded as the most convenient, robust and general way to carry out offline analyses of the solutions.

This approach is readily implemented in Matlab by the *gcmfaces* toolbox. It defines a class of objects (the gcmfaces class depicted in Table 11) that is a natural extension to the common array class. Basic operators (such as "+") are readily overloaded (i.e. re-defined) for the gcmfaces class. Thus, for example, the addition of two gcmfaces objects can simply be written in the compact and general "fld1 + fld2" form – exactly as if fld1 and fld2 were two array objects. Note that the grid-specific organization of the binary data (e.g. Table 11) does not appear in "fld1 + fld2", which reflects that this compact code is immediately applicable to all supported grids (Fig. 1).

Transport and budget computations are coded with the same degree of generality within gcmfaces. Hard-coding array sizes or exploiting specific grid symmetries (e.g. the zonal symmetry of the LL grid) is excluded, in order to avoid having to re-code the same diagnostics on different grids. Two basic elements are instrumental to the generality of gcmfaces codes, which are worth noting here. Firstly, any transport is computed following a grid line path, as illustrated in Fig. 13. Three types of paths are readily treated in a general fashion: small circles of constant latitude, great circles defined by two points (as shown in Fig. 13), and the edge of a specified subdomain. Secondly, the familiar mechanism[13] by which rows and columns of neighboring faces are appended at the edges of an array (e.g. to f1, f2, f3, f4 or f5 in Table 11) is readily implemented. This yields general codes for gradient, rotational, divergence, etc. computations that are immediately applicable to all supported grids (Fig. 1).

[13] It is commonly called *exchanges* in the parallel computing terminology.

From the state estimate output made available online, users can readily re-compute the gcmfaces standard analysis. The standard analysis document serves as a general documentation of the state estimate, and allows for a direct comparison with other MITgcm simulations regardless of the grid-specifics. It proceeds in two steps:

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

The computational loop (i.e. diags_driver.m) uses model output in "release1/nctiles/", which results are stored to files in "release1/mat/". The display phase (i.e. diags_driver_tex.m) then generates "release1/tex/standardAnalysis.tex".

Diagnosing mass, heat and salt budgets requires snapshots of the ocean + seaice + snow model state (to compute the tendency terms), as well as time averaged fluxes between snapshots (to match the tendency terms). The MITgcm flux output accounts for variations of layer thicknesses in z^* coordinate. Tendency terms are computed after the fact using snapshots of e.g. η and θ (Sect. 3.1). The assembled mass, heat and salt budgets are provided online in the extensive form (in $\mathrm{kg\,s^{-1}}$, $\mathrm{J\,s^{-1}}$, $\mathrm{g\,s^{-1}}$ respectively) and in nctiles format (monthly, three-dimensional). The budgets residuals are less than 10^{-6} times the budget magnitude (a Euclidean norm is used). Here "mass budget" simply denotes the constant Boussinesq density ρ_c times volume – in contrast with the hydrostatic pressure budget that is most directly relevant to diagnosis of sea level variability (Forget and Ponte, 2015).

The full specification of the MITgcm "diagnostics" package ("data.diagnostics") are available online for ECCO v4, along with the gmfaces (Matlab) codes that assemble the budgets and compute the standard analysis. They can be readily applied to re-runs of the state estimates, or to most perturbation experiments. Re-running the state estimate after editing "data.diagnostics" is the re-commended method for users that desire output that is not readily online.

Appendix D: Profiles

The MITgcm "profiles" package subsamples the model solution, while it is being computed, at the locations and times of observed in situ profiles. At model initialization, observed profiles dates and locations are read from MITprof files (see below) and each profile is allocated to the processor corresponding to its sub-domain tile. The latter is generally facilitated by a pre-processing step: observed profiles are collocated with grid points using gcmfaces (see Appendix C) and grid locations added to the MITgcm input files. During model integration, profiles are sampled at time steps and locations closest to observations, vertically interpolated to the MITprof depth levels, and written to file. At the end of the forward model integration, these profiles are re-read from file along with observed and weight profiles, and the normalized distance between modeled and observed profiles is computed (see Sect. 4).

MITprof files contain in situ profiles (prof_T and prof_S) as well as corresponding state estimate profiles (prof_Testim and prof_Sestim) and least square weights (prof_Tweight and prof_Sweight) as illustrated in Fig. 15. Weights are set according to the method of Forget and Wunsch (2007) albeit with updated variance fields. The normalized distance to observations (Eq. 12; Sect. 4.3) is thus readily computed as

$$\begin{aligned} \mathrm{jT} &= (\mathrm{prof\_Testim\text{-}prof\_T})^2 \cdot \mathrm{prof\_Tweight} \\ \mathrm{jS} &= (\mathrm{prof\_Sestim\text{-}prof\_S})^2 \cdot \mathrm{prof\_Sweight} \end{aligned} \tag{D1}$$

from the content of any MITprof file. The intention is to eventually distribute all observed data constraints used in ECCO (e.g. altimetry and SST) in a similarly self sufficient and practical format (i.e. observations, model values and weights all together).

The MITprof format contains a limited amount of ancillary information: profile locations, dates, and an identifying code (prof_descr). This choice, along with the use of standard depth levels, yields data sets that are both more compact and simpler than most data center formats (e.g. the Argo format), providing easy access to vast collections of profiles of various origins (Table 5). The identifying code may be a cruise ID

(e.g. for shipboard CTDs) or an instrument ID (e.g. for Argo profiles). They are informative of the data origin, and used for analyses of transects or time series.

As part of the MITprof Matlab toolbox, the pre-processing of in situ profiles consists of four basic steps: (1) applying relevant data quality flags, if provided by data center, (2) converting in-situ to potential temperature or pressure to depth, if needed, (3) interpolating to standard depth levels[14], (4) resetting weights to 0 for standard levels that are not closest neighbors to observed levels, for S outside the 25–42 range, and when jT (resp. jS) exceeds 50 (i.e. 7 standard errors) when computed for an Argo-based atlas (Forget, 2010). Zero weights thus indicate suspicious data points that users are advised to discard.

Appendix E: Smooth

The MITgcm "smooth" package is an implementation of recipes presented in detail by Weaver and Courtier (2001). At the core of the method, a diffusion equation is time integrated to smooth a field. Applying the smoother directly (without additional factors) to model-data misfits (as part of $\mathcal{P}$ in Eq. 13) yields a practical method to omit scales at which observations and models are not expected to be consistent with each other. This approach is useful, for example, to constrain eddying models to coarse grained climatological fields, or to constrain models with along-track altimetric data (Forget and Ponte, 2015).

When the smoother is applied to uncorrelated grid scale noise, the resulting fields have a Gaussian correlation (Fig. 14) with a e-folding scale L determined by the joint specification of integration time and diffusivity. The noise amplitude reduction by the smoother (Fig. 14, color scale) can be computed exactly or approximately (Weaver and Courtier, 2001). Normalizing the smoother to account for this effect yields a spatial correlation operator that conserves variance (in the case of uncorrelated noise). A spatial covariance operator is then immediately obtained by further multiplying the normalized smoother with a specified error field, and grid cell areas or volumes are used as a preconditioner (in two- and three-dimensional cases respectively), following Weaver and Courtier (2001).

[14] An option also exists to interpolate to standard density levels, which was used in McCaffrey et al. (2015), although the corresponding option is lacking in MITgcm.

This method is used for all control parameter covariances (see Sect. 4.4; $\mathcal{Q}$ in Eq. 15). Key advantages of this method are that it is matrix free, naturally handles coast lines, and easily accommodates a variety of grids. In practice, "smooth" also damps grid scale noise that can arise from the adjoint model, and it thus facilitates optimization.

Appendix F: Benchmarking

While MITgcm evolves continuously its results are tested against benchmarks on a daily basis, with a variety of compilers, on a variety of computing platforms. These tests are carried using the "CVS" and "testreport" capabilities for short runs (a few time steps), on a small number of processors (or just one), and exclude optimization by compilers. This design is suited to detect mistakes in code revisions and distinguish them from false positives (associated with truncation errors). The ECCO v4 model setup (Sects. 2 and 3) takes full advantage of that framework, which makes it both portable and stable (Sect. "Code availability").

Advanced usage of ECCO v4 may include re-running forward model solutions (the state estimate in particular) or its adjoint. Computational requirements are modest – the 20 year forward model integration typically takes between 6 and 12 h on 96 processors. ECCO v4 users can thus easily re-run the state estimate solution to generate additional output and carry out analyses that may not already be covered by the publicly distributed material. Running the adjoint model allows for analyses of processes and mechanisms (e.g. see Fukumori et al., 2015) as well for the possibility of further optimization of the state estimate.

While the "testreport" tool is very useful and practical, it does not directly apply to the state estimate, but rather to the underlying model code and setup. An extension to the benchmarking framework is therefore proposed that is suited for the full state estimate solution. The benchmarking tool is implemented as a self-contained Matlab routine (testreport_ecco.m). It relies upon distances to observations and monthly mean model output (Table 2). It provides a simple mechanism that allows users to verify that their 20 year solution is acceptably close to the released state estimate. The first three lines of Table 3 are reflective of small differences that user should expect when re-running the state estimate using a different computer or an updated MITgcm code. Such slight changes typically result form compiler optimization of slightly different codes and slightly different arithmetic and MPI libraries.

For any given model run, model-data distances are simply read from a summary text file (typically named cost function0011) that MITgcm generates at the end of the model integration. Benchmark values are then read from a Matlab file (typically name testreport_release1.mat) and relative differences are reported as shown in Table 3. The other tests are slightly more computationally intensive as they read binary output of model fields – a subset of the fields that are distributed online as nctiles files (Appendix C). It was chosen to focus on integrated quantities (global means and transports) that are known to be model sensitive (e.g. see Table 3) and of common interest to ECCO users. Computations of monthly global mean free surface height, temperature and salinity illustrate usage of grid cell surfaces and volumes. If gcmfaces is activated in Matlab (see Appendix C and Sect. "Code availability") then a benchmarking of integrated transports can also be performed.

Appendix G: Solution history

The ECCO-Production, release 1 state estimate was produced in several phases over the course of the ECCO v4 development. In total, 45 iterations were performed, and a summary of the different phases is provided below. We should stress that the documented solution history reflects the progressive development of ECCO v4 – as opposed to a systematic or advocated approach to the optimization of model solutions.

The first series of 14 adjoint iterations was carried (with the MITgcm's checkpoint62k) using a non-synchronous time step (3 h for tracers, and 20 min for momentum), sea surface salinity relaxation to climatological values, and the linear free surface method. Revision 1 was the switch to the one hour time step (for both tracers and momentum) and to the non linear free surface, followed by 14 adjoint iterations (with checkpoint62y).

In revision 2, the Duffy et al. (1999) parameterization as implemented for the MITgcm by Nguyen et al. (2009) was added, the solution was extended through 2011, and 13 more adjoint iterations were carried out (with checkpoint63g). In revision 3, the surface salinity relaxation was removed and its effect replaced by an adjustment of precipitation controls, followed by 3 adjoint iterations (with checkpoint63r). The resulting solution is used in Speer and Forget (2013); Wunsch and Heimbach (2013a, b); Buckley et al. (2014); Balmaseda et al. (2015). In revision 4, the adjustment of precipitation from revision 3 was removed, followed by 8 adjoint iterations (with checkpoint64f).

Up to this point (revision 4, iteration 8) time-variable global mean sea level had been omitted from the altimetric constraints – letting the other data constraints (from in situ hydrography, SST and regional altimetry primarily) determine the solution variability. Then, revision 4 iteration 9 consisted in estimating a time variable global mean precipitation adjustment under the sole constraint of fitting the time-variable global mean altimetry. This operation had very little influence on the rest of model-data misfits – consistent with the analysis presented in Sect. 5.1. This solution is used in Forget and Ponte (2015).

Revision 4 iteration 10 consisted in a filtering of atmospheric control parameters adjustments to reduce irregularities in the forcing that had appeared during adjoint iterations. To further reduce dynamical imbalances during the first years of integration, the initial state of 1 January 1992 as adjusted during the adjoint iterations was replaced with the state of 1 January 1995. This solution is used in Wunsch and Heimbach (2014). Finally, revision 4 iteration 11 consisted in a reduction of vertical viscosity

to $5 \times 10^{-5} m^2 s^{-1}$ to reduce a low bias in the Equatorial Undercurrent velocity, albeit with little impact on model-data misfits.

Revision 4 iteration 11 is the ECCO-Production, release 1 state estimate, which originally ran with MITgcm's checkpoint64t. For benchmarking purposes (Appendix F) the 20 year solution is re-run once a month with the up-to-date MITgcm. As of MITgcm's checkpoint65i, it matches the original solution within the precision seen in Table 3 (top three rows).

Code availability

The MITgcm is developed and maintained within Concurrent Versions System (CVS). This framework allows users to download frozen versions of the model code (checkoint65i at the time of writing) or to maintain their local copy up to date. The evolving code is benchmarked on a daily basis using the "testreport" capability (Appendix F). Documentation for the MITgcm itself, the CVS framework, and the "testreport" capability can respectively be found at:

- http://mitgcm.org/public/r2_manual/latest/online_documents/manual.pdf
- http://mitgcm.org/public/using_cvs.html
- http://mitgcm.org/public/devel_HOWTO/devel_HOWTO.pdf

The ECCO v4 model setup (Sects. 2 and 3) exploits the MITgcm CVS and testreport capabilities, to allow any interested user to obtain the up-to-date setup and re-run the short ECCO v4 benchmark (Appendix F). Results of the automated daily benchmarking are reported at:

- http://mitgcm.org/public/testing.html

The Matlab analysis framework (gcmfaces and MITprof; see Appendices C and D) is also developed and maintained within CVS. The ECCO v4 model setup, gcmfaces, and MITprof can be found, along with their respective documentations, at:

- http://mitgcm.org/viewvc/MITgcm/MITgcm_contrib/gael/verification/
- http://mitgcm.org/viewvc/MITgcm/MITgcm_contrib/gael/matlab_class/
- http://mitgcm.org/viewvc/MITgcm/MITgcm_contrib/gael/profilesMatlabProcessing/

The state estimate monthly output, profiles output, budget output, and the standard analysis (see Appendices C and D) can be found at

- http://mit.ecco-group.org/opendap/ecco_for_las/version_4/release1/

that also provides ancillary data (e.g. grid files), and the stand-alone Matlab routine (eccov4_lonlat.m) that extracts the LL sector out of global LLC fields. The ecco-support@mit.edu mailing list provides for collaborative user support (analogous to mitgcm-support@mitgcm.org).

The Supplement related to this article is available online at doi:10.5194/gmdd-8-3653-2015-supplement.

Acknowledgements. Major support for this work was provided through NASA's Physical Oceanography Program. The bulk of the calculations was performed on the NASA Advanced Supercomputing (NAS) division's Pleiades supercomputer at NASA/ARC. The authors wish to acknowledge the various groups that carry out and promote ocean state estimation, at IFM-UH, SIO and JPL in particular. The authors also wish to give much credit to John Marshall for his leadership of, and continued commitment to, the MITgcm ocean model development. To Detlef Stammer for his leadership of the German ECCO project, and his continued commitment to, and promotion of the MITgcm state estimation framework. To Fastopt for providing the TAF algorithmic differentiation tool, and the support that was provided to facilitate adjointing the exch2 package in particular.

References

Adcroft, A. and Campin, J.: Rescaled height coordinates for accurate representation of free-surface flows in ocean circulation models, Ocean Model., 7, 269–284, 2004. 3660, 3661, 3662

Adcroft, A., Hill, C., and Marshall, J.: A new treatment of the Coriolis terms in C-grid models at both high and low resolutions, Mon. Weather Rev., 127, 1928–1936, 1999. 3666

Adcroft, A., Campin, J.-M., Hill, C., and Marshall, J.: Implementation of an atmosphere-ocean general circulation model on the expanded spherical cube, Mon. Weather Rev., 132, 2845–2863, doi:10.1175/MWR2823.1, 2004a. 3660

Adcroft, A., Hill, C., Campin, J.-M., Marshall, J., and Heimbach, P.: Overview of the formulation and numerics of the MITGCM, in: Proceedings of the ECMWF Seminar Series on Numerical Methods, Recent Developments in Numerical Methods for Atmosphere and Ocean Modelling, 139–149, ECMWF, available at: http://mitgcm.org/pdfs/ECMWF2004-Adcroft.pdf (last access: 29 April 2015), 2004b. 3655, 3660, 3695

Andersen, O. B. and Knudsen, P.: DNSC08 mean sea surface and mean dynamic topography models, J. Geophys. Res.-Oceans, 114, C11001, doi:10.1029/2008JC005179, 2009. 3723

Balmaseda M. A., Hernandez, F., Storto, A., Palmer, M., Shi, L., Smith, G., Toyoda, T., Valdivieso, M., Alves, O., Barnier, B., Boyer, T., Chang, Y.-S., Chepurin, G. A., Ferry, N., Forget, G., Fujii, Y., Good, S., Guinehut, S., Haines, K., Ishikawa, Y., Keeley, S., Köhl, A., Lee, T., Martin, M., Masina, S., Masuda, S., Meyssignac, B., Mogensen, K., Parent, L., Peterson, D., Yin, Y., Vernieres, G., Wang, X., Waters, J., Wedd, R., Wang, O., Xue, Y., Chevallier, M., Lemieux, J.-F., Dupont, F., Kuragano, T., Kamachi, M., Awaji, T., Cantalbiano, A., Wilmer-Becker, K., and Gaillard, F.: The Ocean Reanalyses Intercomparison Project (ORA-IP), Proceedings of the Institute of Marine Engineering, Science, and Technology, Journal of operational oceanography, in press, 2015. 3678, 3682, 3685, 3688, 3702

Barnier, B., Madec, G., Penduff, T., Molines, J.-M., Treguier, A.-M., Le Sommer, J., Beckmann, A., Biastoch, A., Böning, C., Dengg, J., Derval, C., Durand, E., Gulev, S., Remy, E., Talandier, C., Theetten, S., Maltrud, M., McClean, J., and De Cuevas, B.: Impact of partial steps and momentum advection schemes in a global ocean circulation model at eddy-permitting resolution, Ocean Dynam., 56, 543–567, 2006. 3690

Barth, A., Beckers, J.-M., Troupin, C., Alvera-Azcárate, A., and Vandenbulcke, L.: divand-1.0: *n*-dimensional variational data analysis for ocean observations, Geosci. Model Dev., 7, 225–241, doi:10.5194/gmd-7-225-2014, 2014. 3687

Blessing, S., Kaminski, T., Lunkeit, F., Matei, I., Giering, R., Köhl, A., Scholze, M., Herrmann, P., Fraedrich, K., and Stammer, D.: Testing variational estimation of process parameters and initial conditions of an earth system model, Tellus A, 66, 22606, doi:10.3402/tellusa.v66.22606, 2014. 3689

Buckley, M. W., Ponte, R. M., Forget, G., and Heimbach, P.: Low-frequency SST and upper-ocean heat content variability in the North Atlantic, J. Climate, 27, 4996–5018, 2014. 3678, 3688, 3702

Buckley, M. W., Ponte, R. M., Forget, G., and Heimbach, P.: Determining the origins of advective heat transport convergence variability in the North Atlantic, J. Climate, doi:10.1175/JCLI-D-14-00579.1, in press, 2015. 3688

Campin, J., Adcroft, A., Hill, C., and Marshall, J.: Conservation of properties in a free surface model, Ocean Model., 6, 221–244, 2004. 3663, 3665, 3694

Campin, J.-M., Marshall, J., and Ferreira, D.: Sea ice–ocean coupling using a rescaled vertical coordinate, Ocean Model., 24, 1–14, 2008. 3663, 3664, 3667, 3675, 3719

Chaudhuri, A. H., Ponte, R. M., Forget, G., and Heimbach, P.: A comparison of atmospheric reanalysis surface products over the ocean and implications for uncertainties in air–sea boundary forcing, J. Climate, 26, 153–170, 2013. 3670, 3677, 3682

Comiso, J.: Bootstrap sea ice concentrations for NIMBUS-7 SMMR and DMSP SSM/I, Digital Media, National Snow and Ice Data Center, doi:10.5067/J6JQLS9EJ5HU, 1999. 3723

Dail, H. and Wunsch, C.: Dynamical reconstruction of upper-ocean conditions in the Last Glacial Maximum Atlantic, J. Climate, 27, 807–823, doi:10.1175/JCLI-D-13-00211.1, 2014. 3656

Danabasoglu, G., Yeager, S. G., Bailey, D., Behrens, E., Bentsen, M., Bi, D., Biastoch, A., Böning, C., Bozec, A., Canuto, V. M., Cassou, C., Chassignet, E., Coward, A. C., Danilov, S., Diansky, N., Drange, H., Farneti, R., Fernandez, E., Fogli, P. G., Forget, G., Fujii, Y., Griffies, S. M., Gusev, A., Heimbach, P., Howard, A., Jung, T., Kelley, M., Large, W. G., Leboissetier, A., Lu, J., Madec, G., Marsland, S. J., Masina, S., Navarra, A., Nurser, A. G., Pirani, A., Salas y Mélia, D., Samuels, B. L., Scheinert, M., Sidorenko, D., Treguier, A.-M., Tsujino, H., Uotila, P., Valcke, S., Voldoire, A., and Wang, Q.: North Atlantic simulations in Coordinated

Ocean-ice Reference Experiments phase {II} (CORE-II). Part I: Mean states, Ocean Model., 73, 76–107, doi:10.1016/j.ocemod.2013.10.005, 2014. 3659, 3660, 3668, 3682
Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., Bauer, P., Bechtold, P., Beljaars, A. C. M., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A. J., Haimberger, L., Healy, S. B., Hersbach, H., Hólm, E. V., Isaksen, L., Kållberg, P., Köhler, M., Matricardi, M., McNally, A. P., Monge-Sanz, B. M., Morcrette, J.-J., Park, B.-K., Peubey, C., de Rosnay, P., Tavolato, C., Thépaut, J.-N., and Vitart, F.: The ERA-Interim reanalysis: configuration and performance of the data assimilation system, Q. J. Roy. Meteor. Soc., 137, 553–597, 2011. 3667
Denham, C. R.: Seagrid orthogonal grid maker for matlab, US Geological Survey, 384, available at: http://woodshole.er.usgs.gov/staffpages/cdenham/public_html/seagrid/seagrid.html (last access: 29 April 2015) 2000. 3691
Duffy, P., Eby, M., and Weaver, A.: Effects of sinking of salt rejected during formation of sea ice on results of an ocean–atmosphere-sea ice climate model, Geophys. Res. Lett., 26, 1739–1742, 1999. 3668, 3702
Dutkiewicz, S., Sokolov, A. P., Scott, J., and Stone, P. H.: A Three-Dimensional Ocean-Seaice-Carbon Cycle Model and its Coupling to a Two-Dimensional Atmospheric Model: Uses in Climate Change Studies, MIT Joint Program on the Science and Policy of Global Change available at: http://hdl.handle.net/1721.1/18091 (last access: 29 April 2015), 2005. 3686
Ekman, V. W.: On the influence of the earth's rotation on ocean currents, Ark. Mat. Astron. Fys., 2, 1–53, 1905. 3685
Fekete, B. M., Vörösmarty, C. J., and Grabs, W.: High-resolution fields of global runoff combining observed river discharge and simulated water balances, Global Biogeochem. Cy., 16, 15.1–15.10, doi:10.1029/1999GB001254, 2002. 3667
Fenty, I. and Heimbach, P.: Coupled sea ice–ocean-state estimation in the Labrador Sea and Baffin Bay, J. Phys. Oceanogr., 43, 884–904, 2013. 3674
Ferreira, D., Marshall, J., and Heimbach, P.: Estimating eddy stresses by fitting dynamics to observations using a residual-mean ocean circulation model and its adjoint, J. Phys. Oceanogr., 35, 1891–1910, doi:10.1175/JPO2785.1, 2005. 3655, 3672, 3680
Ferreira, D., Marshall, J., and Rose, B.: Climate determinism revisited: multiple equilibria in a complex climate model, J. Climate, 24, 992–1012, 2011. 3686

Ferron, B. and Marotzke, J.: Impact of 4D-variational assimilation of WOCE hydrography on the meridional circulation of the Indian Ocean, Deep-Sea Res. Pt. II, 50, 2005–2021, 2003. 3655
Follows, M. J., Dutkiewicz, S., Grant, S., and Chisholm, S. W.: Emergent biogeography of microbial communities in a model ocean, Science, 315, 1843–1846, 2007. 3686
Forget, G.: Mapping ocean observations in a dynamical framework: a 2004–06 ocean atlas, J. Phys. Oceanogr., 40, 1201–1221, 2010. 3656, 3677, 3680, 3687, 3699, 3723
Forget, G.: The observed abyssal variability puzzle, Geophys. Res. Lett., in preparation, 2015a. 3686, 3688
Forget, G., Ferreira, D., and Liang, X.: On the observability of turbulent transport rates by Argo: supporting evidence from an inversion experiment, Ocean Sci. Discuss., submitted, 2015b. 3666, 3677, 3680, 3681, 3686, 3688, 3689
Forget, G. and Ponte, R.: The partition of regional sea level variability, Prog. Oceanogr., under review, 2015. 3670, 3671, 3676, 3678, 3682, 3688, 3697, 3699, 3702
Forget, G. and Wunsch, C.: Estimated global hydrographic variability, J. Phys. Oceanogr., 37, 1997–2008, 2007. 3670, 3676, 3680, 3682, 3686, 3698
Forget, G., Ferron, B., and Mercier, H.: Combining Argo profiles with a general circulation model in the North Atlantic. Part 1: Estimation of hydrographic and circulation anomalies from synthetic profiles, over a year, Ocean Model., 20, 1–16, 2008a. 3679, 3689
Forget, G., Mercier, H., and Ferron, B.: Combining Argo profiles with a general circulation model in the North Atlantic. Part 2: Realistic transports and improved hydrography, between spring 2002 and spring 2003, Ocean Model., 20, 17–34, 2008b. 3656, 3679, 3689
Forget, G., Maze, G., Buckley, M., and Marshall, J.: Estimated seasonal cycle of North Atlantic eighteen degree water volume, J. Phys. Oceanogr., 41, 269–286, 2011. 3675, 3686
Fukumori, I.: A partitioned Kalman filter and smoother, Mon. Weather Rev., 130, 1370–1383, 2002. 3737
Fukumori, I., Menemenlis, D., and Lee, T.: A near-uniform basin-wide sea level fluctuation of the Mediterranean Sea, J. Phys. Oceanogr., 37, 338–358, 2007. 3687
Fukumori, I., Wang, O., Llovel, W., Fenty, I., and Forget, G.: A near-uniform fluctuation of ocean bottom pressure and sea level across the deep ocean basins of the Arctic Ocean and the Nordic Seas, Prog. Oceanogr., doi:10.1016/j.pocean.2015.01.013, in press, 2015. 3672, 3673, 3687, 3688, 3700

Gaspar, P., Grégoris, Y., and Lefevre, J.-M.: A simple eddy kinetic energy model for simulations of the oceanic vertical mixing: tests at station papa and long-term upper ocean study site, J. Geophys. Res., 95, 16179–16193, 1990. 3665, 3673, 3694

Gent, P. and Mcwilliams, J.: Isopycnal mixing in ocean circulation models, J. Phys. Oceanogr., 20, 150–155, 1990. 3655, 3662, 3674, 3694

Giering, R. and Kaminski, T.: Recipes for adjoint code construction, ACM T. Math. Software, 24, 437–474, 1998. 3655, 3672

Giering, R., Kaminski, T., and Slawig, T.: Generating efficient derivative code with TAF: adjoint and tangent linear Euler flow around an airfoil, Future Gener. Comp. Sy., 21, 1345–1355, 2005. 3672

Goldberg, D. N. and Heimbach, P.: Parameter and state estimation with a time-dependent adjoint marine ice sheet model, The Cryosphere, 7, 1659–1678, doi:10.5194/tc-7-1659-2013, 2013. 3687

Griewank, A.: Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation, Optim. Method. Softw., 1, 35–54, 1992. 3655, 3672

Griewank, A. and Walther, A.: Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation, Siam, 2008. 3672

Griffies, S. and Greatbatch, R.: Physical processes that impact the evolution of global mean sea level in ocean climate models, J. Marine Syst., 51, 37–72, 2012. 3676

Hansen, P. C.: Analysis of discrete ill-posed problems by means of the L-curve, SIAM Rev., 34, 561–580, 1992. 3671

Hascoët, L. and Pascual, V.: The Tapenade automatic differentiation tool: principles, model, and specification, ACM T. Math. Software, 39, 20.1–20.43, doi:10.1145/2450153.2450158, 2013. 3672

Heimbach, P., Hill, C., and Giering, R.: Automatic generation of efficient adjoint code for a parallel Navier–Stokes solver, in: Computational Science – ICCS 2002, Springer, 1019–1028, 2002. 3655

Heimbach, P., Hill, C., and Giering, R.: An efficient exact adjoint of the parallel MIT general circulation model, generated via automatic differentiation, Future Gener. Comp. Sy., 21, 1356–1371, 2005. 3655, 3669, 3672, 3673, 3733

Heimbach, P., Forget, G., Ponte, R., Wunsch, C., Balmaseda, M., Awaji, T., Baehr, J., Behringer, D., Carton, J., Ferry, N., Fischer, A., Fukumori, I., Giese, B., Haines, K., Harrison, E., Hernandez, F., Kamachi, M., Keppenne, C., Köhl, A., Lee, T., Menemenlis, D., Oke, P., Remy, E.,

Rienecker, M., Rosati, A., Smith, D., Speer, K., Stammer, D., and Weaver, A.: Observational requirements for global-scale ocean climate analysis: lessons from ocean state estimation, in: Proceedings of the OceanObs09 Conference: Sustained Ocean Observations and Information for Society, Venice, Italy, vol. 2, doi:10.5270/OceanObs09.cwp.42, 2009. 3689

Heimbach, P., Menemenlis, D., Losch, M., Campin, J.-M., and Hill, C.: On the formulation of sea-ice models. Part 2: Lessons from multi-year adjoint sea-ice export sensitivities through the Canadian Arctic Archipelago, Ocean Model., 33, 145–158, 2010. 3674

Heimbach, P., Wunsch, C., Ponte, R. M., Forget, G., Hill, C., and Utke, J.: Timescales and regions of the sensitivity of Atlantic meridional volume and heat transport: toward observing system design, Deep-Sea Res. Pt. II, 58, 1858–1879, 2011. 3672, 3687

Holland, W. R. and Malanotte-Rizzoli, P.: Assimilation of altimeter data into an ocean circulation model: space versus time resolution studies, J. Phys. Oceanogr., 19, 1507–1534, 1989. 3655

Hoppe, C. M., Elbern, H., and Schwinger, J.: A variational data assimilation system for soil–atmosphere flux estimates for the Community Land Model (CLM3.5), Geosci. Model Dev., 7, 1025–1036, doi:10.5194/gmd-7-1025-2014, 2014. 3687

Hoteit, I., Cornuelle, B., Köhl, A., and Stammer, D.: Treating strong adjoint sensitivities in tropical eddy-permitting variational data assimilation, Q. J. Roy. Meteor. Soc., 131, 3659–3682, 2005. 3674

Hoteit, I., Cornuelle, B., Kim, S., Forget, G., Köhl, A., and Terrill, E.: Assessing 4D-VAR for dynamical mapping of coastal high-frequency radar in San Diego, Dynam. Atmos. Oceans, 48, 175–197, 2009. 3656

Hoteit, I., Hoar, T., Gopalakrishnan, G., Collins, N., Anderson, J., Cornuelle, B., Köhl, A., and Heimbach, P.: A MITgcm/DART ensemble analysis and prediction system with application to the Gulf of Mexico, Dynam. Atmos. Oceans, 63, 1–23, 2013. 3687

Ives, D. C. and Zacharias, R. M.: Conformal mapping and orthogonal grid generation, J. Propul. Power, 5, 327–333, 1989. 3691

Jackett, D. R. and McDougall, T. J.: Minimal adjustment of hydrographic profiles to achieve static stability, J. Atmos. Ocean. Tech., 12, 381–389, 1995. 3662

Jiang, Z., Hui, W., and Kamachi, M.: The improvement made by a modified TLM in 4DVAR with a geophysical boundary layer model, Adv. Atmos. Sci., 19, 563–582, 2002. 3673, 3674

Kalmikov, A. G. and Heimbach, P.: A Hessian-based method for uncertainty quantification in global ocean state estimation, SIAM J. Sci. Comput., 36, S267–S295, 2014. 3685

Köhl, A.: Evaluation of the GECCO2 ocean synthesis: transports of volume, heat and freshwater in the Atlantic, Q. J. Roy. Meteor. Soc., 141, 166–181, doi:10.1002/qj.2347, 2014. 3656, 3659, 3737

Köhl, A. and Stammer, D.: Optimal observations for variational data assimilation, J. Phys. Oceanogr., 34, 529–542, 2004. 3687

Köhl, A. and Stammer, D.: Decadal sea level changes in the 50-Year GECCO Ocean Synthesis, J. Climate, 21, 1876–1890, doi:10.1175/2007JCLI2081.1, 2008. 3656

Köhl, A., Stammer, D., and Cornuelle, B.: Interannual to decadal changes in the ECCO global synthesis., J. Phys. Oceanogr., 37, 529–542, doi:10.1175/2513.1, 2007. 3656

Köhl, A., Siegismund, F., and Stammer, D.: Impact of assimilating bottom pressure anomalies from GRACE on ocean circulation estimates, J. Geophys. Res., 117, C04032, doi:10.1029/2011JC007623, 2012. 3656

Large, W. and Yeager, S.: Diurnal to decadal global forcing for ocean and sea-ice models: the data sets and flux climatologies, Technical Report TN-460+STR, NCAR, 2004. 3667, 3668

Large, W., McWilliams, J., and Doney, S.: Oceanic vertical mixing: a review and a model with a nonlocal boundary layer parameterization, Rev. Geophys., 32, 363–403, 1994. 3665

Liang, X., Wunsch, C., Heimbach, P., and Forget, G.: Vertical redistribution of oceanic heat content, J. Climate, doi:10.1175/JCLI-D-14-00550.1, in press, 2015. 3688

Liu, C., Köhl, A., and Stammer, D.: Adjoint-based estimation of eddy-induced tracer mixing parameters in the global ocean, J. Phys. Oceanogr., 42, 1186–1206, 2012. 3681, 3689

Llovel, W., Willis, J., Landerer, F., and Fukumori, I.: Deep-ocean contribution to sea level and energy budget not detectable over the past decade, Nature Climate Change, 4, 1031–1035, 2014. 3680

Locarnini, R., Mishonov, A., Antonov, J., Boyer, T., Garcia, H., and Levitus, S.: World Ocean Atlas 2005 Volume 1: Temperature, NOAA Atlas NESDIS, 2006. 3723

Losch, M., Menemenlis, D., Campin, J.-M., Heimbach, P., and Hill, C.: On the formulation of sea-ice models. Part 1: Effects of different solver implementations and parameterizations, Ocean Model., 33, 129–144, 2010. 3667, 3673

Losch, M., Strass, V., Cisewski, B., Klaas, C., and Bellerby, R. G.: Ocean state estimation from hydrography and velocity observations during EIFEX with a regional biogeochemical ocean circulation model, J. Marine Syst., 129, 437–451, 2014. 3656

Lyman, J. and Johnson, G.: Estimating global ocean heat content changes in the upper 1800 m since 1950 and the influence of climatology choice, J. Climate, 27, 1945–1957, 2014. 3677

Marotzke, J., Giering, R., Zhang, K. Q., Stammer, D., Hill, C., and Lee, T.: Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport sensitivity, J. Geophys. Res.-Oceans, 104, 29529–29547, 1999. 3655, 3687

Marshall, J., Adcroft, A., Hill, C., Perelman, L., and Heisey, C.: A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers, J. Geophys. Res., 102, 5753–5766, 1997. 3655, 3660, 3694

Marshall, J., Ferreira, D., Campin, J., and Enderton, D.: Mean climate and variability of the atmosphere and ocean on an aquaplanet, J. Atmos. Sci., 64, 4270–4286, 2007a. 3686

Marshall, J., Ferreira, D., Campin, J.-M., and Enderton, D.: Mean climate and variability of the atmosphere and ocean on an aqua-planet, J. Atmos. Sci., 64, 4270–4286, 2007b. 3657

Marshall, J., Anderson, A., Dewar, W., Doney, S., Edson, J., Ferrari, R., Forget, G., Fratantoni, D., Gregg, M., Joyce, T., Kelly, K., Lozier, S., Lumpkin, R., Maze, G., Paster, J., Samelson, R., Silverthorne, K., Skyllingstad, E., Straneo, F., Talley, L., Thomas, L., Toole, J., and Weller, R.: Observing the cycle of convection and restratification over the Gulf Stream and the subtropical gyre of the North Atlantic Ocean: preliminary results from the CLIMODE field campaign, B. Am. Meteorol. Soc., 90, 1337–1350, 2009. 3722

Marshall, J., Armour, K. C., Scott, J. R., Kostov, Y., Hausmann, U., Ferreira, D., Shepherd, T. G., and Bitz, C. M.: The ocean's role in polar climate change: asymmetric Arctic and Antarctic responses to greenhouse gas and ozone forcing, Philos. T. R. Soc. A, 372, 20130040, doi:10.1098/rsta.2013.0040, 2014. 3659, 3660

Marzocchi, W. and Jordan, T. H.: Testing for ontological errors in probabilistic forecasting models of natural systems, P. Natl. Acad. Sci. USA, 111, 11973–11978, 2014. 3682, 3683, 3690

Maze, G., Forget, G., Buckley, M., Marshall, J., and Cerovecki, I.: Using transformation and formation maps to study the role of air–sea heat fluxes in North Atlantic Eighteen Degree Water formation, J. Phys. Oceanogr., 39, 1818–1835, 2009. 3687

Mazloff, M. R., Heimbach, P., and Wunsch, C.: An eddy-permitting Southern Ocean state estimate, J. Phys. Oceanogr., 40, 880–899, 2010. 3656

McCaffrey, K., Fox-Kemper, B., and Forget, G.: Estimates of ocean macro-turbulence: structure function and spectral slope from Argo profiling floats, J. Phys. Oceanogr., revised, 2015. 3686, 3699

Menemenlis, D., Hill, C., Adcroft, A., Campin, J., Cheng, B., Ciotti, B., Fukumori, I., Koehl, A., Heimbach, P., Henze, C., Lee, T., Stammer, D., Taft, J., and Zhang, J.: NASA supercom-

puter improves prospects for ocean climate research, EOS T. Am. Geophys. Un., 86–96, doi:10.1029/2005EO090002, 2005a. 3657
Mercier, H.: Determining the general circulation of the ocean: a nonlinear inverse problem, J. Geophys. Res.-Oceans, 91, 5103–5109, 1986. 3655
MITgcm Group: MITgcm Release 1 Manual, Technical report, available at: http://mitgcm.org/public/sealion/ (last access: 29 April 2015), Massachusetts Institute of Technology, Cambridge, MA 02139, USA, 346 pp., 2002. 3660, 3695
Munk, W. H.: Abyssal recipes, in: Deep Sea Research and Oceanographic, vol. 13, Elsevier, 707–730, doi:10.1016/0011-7471(66)90602-4, 1966. 3685
Nguyen, A., Menemenlis, D., and Kwok, R.: Improved modeling of the Arctic halocline with a subgrid-scale brine rejection parameterization, J. Geophys. Res.-Oceans, 114, C11014, doi:10.1029/2008JC005121, 2009. 3668, 3673, 3702
Olbers, D., Wenzel, M., and Willebrand, J.: The inference of North Atlantic circulation patterns from climatological hydrographic data, Rev. Geophys., 23, 313–356, 1985. 3655
Pavlis, N. K., Holmes, S. A., Kenyon, S. C., and Factor, J. K.: The development and evaluation of the Earth Gravitational Model 2008 (EGM2008), J. Geophys. Res.-Sol. Ea., 117, B04406, doi:10.1029/2011JB008916, 2012. 3723
Piecuch, C., Ponte, R., Heimbach, H., and Forget, G.: Sensitivity of Ocean-Model Estimates of Contemporary Global and Regional Sea-Level Changes to Geothermal Heat Flow, in preparation, 2015. 3685
Ponte, R. M.: Oceanic response to surface loading effects neglected in volume-conserving models, J. Phys. Oceanogr., 36, 426–434, doi:10.1175/JPO2843.1, 2006. 3667
Ponte, R. M., Wunsch, C., and Stammer, D.: Spatial mapping of time-variable errors in Jason-1 and TOPEX/Poseidon sea surface height measurements, J. Atmos. Ocean. Tech., 24, 1078–1085, 2007. 3670, 3682
Prinn, R. G., Heimbach, P., Rigby, M., Dutkiewicz, S., Melillo, J. M., Reilly, J. M., Kicklighter, D. W., and Waugh, C.: A Strategy for a Global Observing System for Verification of National Greenhouse Gas Emissions, Tech. rep., MIT Joint Program on the Science and Policy of Global Change, 2011. 3689
Purkey, S. G. and Johnson, G. C.: Warming of global abyssal and deep Southern Ocean waters between the 1990s and 2000s: contributions to global heat and sea level rise budgets*, J. Climate, 23, 6336–6351, 2010. 3679

Purser, R. and Rančić, M.: Smooth quasi-homogeneous gridding of the sphere, Q. J. Roy. Meteor. Soc., 124, 637–647, 1998. 3657
Quinn, K. J. and Ponte, R. M.: Estimating weights for the use of time-dependent gravity recovery and climate experiment data in constraining ocean models, J. Geophys. Res., 113, C12013, doi:10.1029/2008JC004903, 2008. 3670, 3682
Quinn, K. J. and Ponte, R.: Uncertainty in ocean mass trends from GRACE, Geophys. J. Int., 181, 762–768, doi:10.1111/j.1365-246X.2010.04508.x, 2010. 3670, 3682
Rančić, M., Purser, R., and Mesinger, F.: A global shallow-water model using an expanded spherical cube: gnomonic versus conformal coordinates, Q. J. Roy. Meteor. Soc., 122, 959–982, 1996. 3657
Redi, M. H.: Oceanic isopycnal mixing by coordinate rotation, J. Phys. Oceanogr., 12, 1154–1158, 1982. 3674, 3694
Reynolds, R. W., Rayner, N. A., Smith, T. M., Stokes, D. C., and Wang, W.: An improved in situ and satellite SST analysis for climate, J. Climate, 15, 1609–1625, 2002. 3723
Risien, C. and Chelton, D.: A global climatology of surface wind and wind stress fields from eight years of QuikSCAT scatterometer data, J. Phys. Oceanogr., 38, 2379–2413, 2008. 3668
Roquet, F., Charrassin, J.-B., Marchand, S., Boehme, L., Fedak, M., Reverdin, G., and Guinet, C.: Delayed-mode calibration of hydrographic data obtained from animal-borne satellite relay data loggers, J. Atmos. Ocean. Tech., 28, 787–801, 2011. 3722
Ryskin, G. and Leal, L.: Orthogonal mapping, J. Comput. Phys., 50, 71–100, 1983. 3692
Scharroo, R., Lillibridge, J., Smith, W., and Schrama, E.: Cross-calibration and long-term monitoring of the microwave radiometers of ERS, TOPEX, GFO, Jason, and Envisat, Mar. Geod., 27, 279–297, doi:10.1080/01490410490465265, 2004. 3723
Schott, F. and Zantopp, R.: On the effect of vertical mixing on the determination of absolute currents by the beta spiral method, Deep-Sea Res., 27, 173–180, 1980. 3655
Shchepetkin, A. F. and McWilliams, J. C.: The regional oceanic modeling system (ROMS): a split-explicit, free-surface topography-following-coordinate ocean model, Ocean Model., 9, 347–404, 2005. 3693
Siedler, G., Griffies, S., Gould, J., and Church, J.: Ocean Circulation and Climate: a 21st Century Perspective, International Geophysics, Elsevier Science, 2013.
Speer, K. and Forget, G.: Chapter 9 – Global distribution and formation of mode waters, in: Ocean Circulation and Climate: a 21st Century Perspective, 211–226, doi:10.1016/B978-0-12-391851-2.00009-X, 2013. 3656, 3678, 3680, 3686, 3688, 3702

Stammer, D.: Adjusting internal model errors through ocean state estimation, J. Phys. Oceanogr., 35, 1143–1153, doi:10.1175/JPO2733.1, 2005. 3680
Stammer, D., Wunsch, C., Giering, R., Eckert, C., Heimbach, P., Marotzke, J., Adcroft, A., Hill, C., and Marshall, J.: The Global ocean circulation during 1992–1997, estimated from ocean observations and a general circulation model, J. Geophys. Res.-Oceans, 107, 1.1–1.27, doi:10.1029/2001JC000888, 2002. 3655
Stammer, D., Ueyoshi, K., Köhl, A., Large, W., Josey, S., and Wunsch, C.: Estimating air–sea fluxes of heat, freshwater, and momentum through global ocean data assimilation, J. Geophys. Res, 109, 8691–8702, 2004. 3656
Steele, M., Morley, R., and Ermold, W.: PHC: a global ocean hydrography with a high-quality Arctic Ocean, J. Climate, 14, 2079–2087, 2001. 3723
Stommel, H.: The delicate interplay between wind-stress and buoyancy input in ocean circulation: the Goldsbrough variations*, Tellus A, 36, 111–119, 1984. 3664
Sverdrup, H. U.: Wind-driven currents in a baroclinic ocean; with application to the equatorial currents of the eastern Pacific, P. Natl. Acad. Sci. USA, 33, 318–326, 1947. 3685
Thacker, W. C. and Long, R. B.: Fitting dynamics to data, J. Geophys. Res.-Oceans, 93, 1227–1240, 1988. 3655, 3671, 3672
Toole, J., Krishfield, R., Timmermans, M.-L., and Proshutinsky, A.: The ice-tethered profiler: Argo of the Arctic, Oceanography, 24, 126–135, 2011. 3722
Trefethen, L.: SCPACK user's guide, Tech. rep., Technical Report 89-2, MIT Numerical Analysis Report, 1989. 3691
Tziperman, E. and Thacker, W. C.: An optimal-control/adjoint-equations approach to studying the oceanic general circulation, J. Phys. Oceanogr., 19, 1471–1485, 1989. 3655
Tziperman, E., Thacker, W. C., Long, R. B., and Hwang, S.-M.: Oceanic data analysis using a general circulation model. Part I: Simulations, J. Phys. Oceanogr., 22, 1434–1457, 1992a. 3655
Tziperman, E., Thacker, W. C., Long, R. B., Hwang, S.-M., and Rintoul, S. R.: Oceanic data analysis using a general circulation model. Part II: A North Atlantic model, J. Phys. Oceanogr., 22, 1458–1485, 1992b. 3655
Utke, J., Naumann, U., Fagan, M., Tallent, N., Strout, M., Heimbach, P., Hill, C., and Wunsch, C.: OpenAD/F: a modular open-source tool for automatic differentiation of Fortran codes, ACM T. Math. Software, 34, 18, doi:10.1145/1377596.1377598, 2008. 3672

Verdy, A., Mazloff, M. R., Cornuelle, B. D., and Kim, S. Y.: Wind-driven sea level variability on the California coast: an adjoint sensitivity analysis, J. Phys. Oceanogr., 44, 297–318, 2014. 3673
Weaver, A. and Courtier, P.: Correlation modelling on the sphere using a generalized diffusion equation, Q. J. Roy. Meteor. Soc., 127, 1815–1846, 2001. 3671, 3677, 3699, 3700, 3734
Wells, D.: Prime Numbers: the Most Mysterious Figures in Math, John Wiley & Sons, 2011. 3659
Wilkin, J. and Hedström, K.: User's manual for an orthogonal curvilinear grid-generation package, Institute of Marine and Coastal Sciences, Rutgers University, available at: http://www.marine.rutgers.edu/po/tools/gridpak/grid_manual.ps.gz (last access: 29 April 2015), 1998. 3691
Wilson, C., Chipperfield, M. P., Gloor, M., and Chevallier, F.: Development of a variational flux inversion system (INVICAT v1.0) using the TOMCAT chemical transport model, Geosci. Model Dev., 7, 2485–2500, doi:10.5194/gmd-7-2485-2014, 2014. 3687
Wunsch, C.: Determining the general circulation of the oceans: a preliminary discussion, Science, 196, 871–875, 1977. 3654, 3655, 3685
Wunsch, C.: Acoustic tomography and other answers, in: It's the water that makes you drunk. A celebration in Geophysics and Oceanography – 1982. In honor of Walter Munk on his 65th birthday, Scripps Institution of Oceanography Reference Series 84-5, Scripps Institution of Oceanography of California, San Diego, La Jolla, CA, 47–62, 1984. 3655
Wunsch, C.: Discrete Inverse and State Estimation Problems: with Geophysical Fluid Applications, Cambridge University Press, 2006. 3654
Wunsch, C. and Heimbach, P.: Practical global oceanic state estimation, Physica D, 230, 197–208, 2007. 3737
Wunsch, C. and Heimbach, P.: The global zonally integrated ocean circulation, 1992–2006: seasonal and decadal variability, J. Phys. Oceanogr., 39, 351–368, doi:10.1175/2008JPO4012.1, 2009. 3656, 3737
Wunsch, C. and Heimbach, P.: Dynamically and kinematically consistent global ocean circulation and ice state estimates, in: Ocean Circulation and Climate: a 21st Century Perspective, 103, 553–579, doi:10.1016/B978-0-12-391851-2.00021-0, 2013a. 3655, 3684, 3702
Wunsch, C. and Heimbach, P.: Two decades of the Atlantic meridional overturning circulation: anatomy, variations, extremes, prediction, and overcoming its limitations, J. Climate, 26, 7167–7186, 2013b. 3678, 3688, 3702

Wunsch, C. and Heimbach, P.: Bidecadal thermal changes in the Abyssal Ocean, J. Phys. Oceanogr., 44, 2013–2030, 2014. 3678, 3684, 3688, 3702

Wunsch, C. and Minster, J.-F.: Methods for box models and ocean circulation tracers: mathematical programing and nonlinear inverse theory, J. Geophys. Res.-Oceans, 87, 5647–5662, 1982. 3655

Wunsch, C., Ponte, R., and Heimbach, P.: Decadal trends in sea level patterns: 1993–2004, J. Climate, 20, 5889–5911, 2007. 3656

Yu, L. and Weller, R. A.: Objectively analyzed air–sea heat fluxes for the global ice-free oceans (1981–2005), B. Am. Meteorol. Soc., 88, 527–539, 2007. 3687

Zanna, L., Heimbach, P., Moore, A. M., and Tziperman, E.: Optimal excitation of interannual Atlantic meridional overturning circulation variability, J. Climate, 24, 413–427, 2011. 3687

Table 1. Interior and boundary model parameters (as of MITgcm's checkpoint65h; see Appendix G). File names where run time model parameters are set are given in square brackets under "MITgcm parameter name". A more exhaustive list of model parameter settings is available within the model standard output (text file).

description	value	MITgcm parameter name
[momentum equation]		[data]
time step	3600 s	deltaTmom
harmonic vertical viscosity	$5 \times 10^{-5}\,m^2\,s^{-1}$	viscAr
harmonic horizontal viscosity (see text)	2×10^{-2}	viscAhGrid
[tracer equations]		[data, data.gmredi]
time step	3600 s	deltaTracer
GM intensity	$10^3\,m^2\,s^{-1}$	GM_background_K
diapycnal diffusivity	$10^{-5}\,m^2\,s^{-1}$	diffKrT,diffKrS
isopycnal diffusivity	$10^3\,m^2\,s^{-1}$	GM_isopycK
[sea floor boundary]		[data]
quadratic bottom drag	10^{-3}	bottomDragQuadratic
[open-ocean surface boundary]		[data.exf]
ocean albedo	0.10	exf_albedo
[ice-covered surface boundary]		[data.seaice]
ice albedo (α_i)	$0.66 < \alpha_i < 0.84$	α_i = SEAICE_*IceAlb*
snow albedo (α_s)	$0.70 < \alpha_i < 0.90$	α_i = SEAICE_*SnowAlb*
maximum ice concentration	0.95	SEAICE_area_max

Table 2. Ocean state characteristics (model-data distances defined in Sect. 4, global mean time series, and averaged meridional transports) used to benchmark 20 year model solutions (Appendix F) and gauge their sensitivity (Tables 3 and 8, Figs. 7 and 11). The first seven rows denote model-data distances. The corresponding observational estimates are listed under "description" where T, S, SST, SSS, SLA, and MDT respectively stand for potential temperature, salinity, sea surface temperature, sea surface salinity, sea level anomaly, and mean dynamic topography. The last six rows are model diagnostics: θ, S, and $\eta + \eta_{\mathrm{ips}}$ denotes model temperature, salinity and free surface height (including η_{ips}; the weight of sea ice plus snow per unit area divided by ρ_{c}; see Campin et al. 2008), and V denotes volume. The governing equations for θ, S, and η are provided in Sect. 3.1.

variable	description
jT	2008–2010 Argo T
jS	2008–2010 Argo S
jTs	1992–2011 Reynolds-SST
jSs	1992–2011 climatological-SSS
jIs	1992–2011 ice-cover fraction
jHa	1992–2011 large scale SLA
jHm	1992–2011 MDT
mH	Monthly global mean $\eta + \eta_{\mathrm{ips}}$
mT	Monthly global mean θ
mS	Monthly global mean S
tV	2008–2010 meridional V transport
tT	2008–2010 meridional θ transport
tS	2008–2010 meridional S transport

Table 3. Benchmarking of (first three rows) and sensitivity experiments with (subsequent rows) the model configuration that produces the state estimate (Sect. 5) being used here as the baseline 20 year solution. The sensitivity experiments pertain to tracer advection schemes, momentum equation settings, and boundary layers. Ocean characteristics that are used to gauged the sensitivity of ocean simulations are listed in Table 2. Departures are computed relative to the state estimate, and normalized by the standard deviation of the state estimate result (for mH, ..., tS) or the state estimate-data distance (for jT, ..., jHm). Positive numbers denote percentages (for differences above 1 %) whereas parenthesized negative numbers are powers of ten (for differences below 1 %). The DST-3 and other advection schemes are discussed in the text.

experiment	jT	jS	jTs	jSs	jIs	jHa	jHm	mH	mT	mS	tV	tT	tS
computer update	(−6)	(−6)	(−7)	(−6)	(−5)	(−6)	(−7)	(−5)	(−5)	(−5)	(−6)	(−6)	(−5)
model update (65 g)	(−7)	(−6)	(−6)	(−5)	(−6)	(−4)	(−4)	(−5)	(−5)	(−5)	(−6)	(−6)	(−5)
24 proc. clusters	(−6)	(−8)	(−6)	(−5)	(−5)	(−4)	(−4)	(−4)	(−5)	(−5)	(−6)	(−6)	(−5)
explicit vert. DST-3	(−3)	(−2)	(−3)	(−2)	(−3)	(−3)	(−2)	60	50	37	(−3)	(−2)	4
3rd order upwind	(−4)	(−3)	(−3)	(−3)	(−4)	(−4)	(−3)	(−2)	(−2)	(−2)	(−4)	(−3)	(−3)
flux limited DST-3	3	6	1	(−2)	(−3)	(−2)	13	98	93	62	1	3	22
C-D scheme	40	52	17	7	2	25	64	69	13	56	2	5	53
added viscosity	6	7	2	6	(−2)	3	6	40	28	31	(−2)	1	22
added bottom visc.	4	5	1	6	(−2)	2	3	18	11	16	(−2)	1	17
KPP instead of GGL	4	11	7	10	11	4	3	148	149	95	(−2)	(−2)	22
added geo. heating	(−3)	(−3)	(−3)	(−3)	(−4)	(−3)	(−3)	(−2)	47	(−2)	(−3)	(−2)	1

Table 4. Generic model-data comparison capabilities provided by the "ecco" package, as of MITgcm's checkpoint65h. The corresponding terms in Eqs. (12)–(14) (m_i, o_i, $\mathcal{D}$, $\mathcal{S}$, $\mathcal{P}$, $\mathbf{W}_i$) are reported here in parentheses, and defined in the text.

generic capability	MITgcm parameter	usage
variable choice (m_i, $\mathcal{D}$)	gencost_barfile	"thetamon", "saltmon", "etaday", etc. for potential temperature ("theta"), salinity ("salt"), see surface height ("eta"), etc.
	gencost_is3d	.TRUE. or.FALSE.
time average (m_i, $\mathcal{S}$)	gencost_avgperiod	"month", "day", or "step" to form monthly, daily, and time step averages, respectively
input data files (o_i)	gencost_datafile	file name root (e.g. "some_sst") for files that may be yearly (e.g. "some_sst_1992" etc.) or otherwise
time average ($\mathcal{P}$)	gencost_preproc gencost_preproc_i	"clim" to form a climatological average misifts e.g. 12 for monthly mean climatologies
spatial smoother ($\mathcal{P}$)	gencost_posproc gencost_posproc_c gencost_posproc_i	"smooth" to apply diffusive smoother to misifts e.g. "some_scale.bin" for the smoothing scale e.g. 10 for the diffusive smoother time stepping
weighing ($\mathbf{W}_i$)	gencost_errfile	squared inverse is used in least squares distance, except in places where a zero value is specified
masking ($\mathbf{W}_i$)	gencost_spzero	data points set to specified value will be omitted
diagnostics	gencost_outputlevel	> 0 to output model data misfits maps to disk

Table 5. Sets of in-situ profiles that have been imposed as constraints on the state estimate (ECCO-Production, release 1) for the period from 1992 to 2011. XBT, CTD and ITP stand for expendable bathythermograph, Conductivity–Temperature–Depth sensors, and Ice-tethered Profiler respectively. SEaOS is data collected by Southern Ocean elephant seals. The CLIMODE field campaign focused on the North Atlantic subtropical gyre (Marshall et al., 2009). A grand total of 1911983 T profiles (resp. 1 239 247 S profiles) were used. Note however that only CTD profiles extend below 2000 m (26 285 for T, 26 220 for S).

Data Set	T profiles	S profiles	origin
Argo	833 033	800 269	IFREMER
CTD	379 012	333 266	NODC, WOA09
XBT	597 009	0	NODC, WOA09
ITP	18 033	17 745	Toole et al. (2011)
SEaOS	103 117	87 806	Roquet et al. (2011)
bobbers	7894	0	D. Fratantoni, CLIMODE
CTD	161	161	L. Talley, CLIMODE

Table 6. Gridded observational data sets that have been imposed as constraints on the state estimate (ECCO-Production, release 1) for the period from 1992 to 2011. *T*, *S*, SST, SLA, MDT, and ICF respectively stand for potential temperature, sea surface temperature, sea surface salinity, sea level anomaly, mean dynamic topography, and ice cover fraction.

variable	description	period	size	origin
MDT	DNSC08 mean SSH minus EGM2008 geoid model	1993–2004	6.2×10^4	Andersen and Knudsen (2009), Pavlis et al. (2012)
T, *S*	blended monthly climatology		$2 \times 5.7 \times 10^8$	
	OCCA	2004–2006		Forget (2010)
	WOA 2005	unclear		Locarnini et al. (2006)
	PHC 3.0	unclear		updated Steele et al. (2001)
SLA	daily bin average of along-track altimetry	1992–2011	7.7×10^7	Scharroo et al. (2004)
SST	monthly maps	1992–2011	1.5×10^7	Reynolds et al. (2002)
ICF	monthly maps	1992–2010	1.4×10^7	Comiso (1999)

Table 7. Control parameters that have been adjusted as part of the state estimate (ECCO-Production, release 1) for the period from 1992 to 2011.

description	frequency	size
Initial condition for temperature	N/A	2.4×10^6
Initial condition for salinity	N/A	2.4×10^6
Diapycnal diffusivity	time mean	2.4×10^6
Isopycnal diffusivity	time mean	2.4×10^6
GM intensity	time mean	2.4×10^6
Atmospheric temperature at 2 m	bi-weekly	3.2×10^7
Specific humidity at 2 m	bi-weekly	3.2×10^7
Precipitation	bi-weekly	3.2×10^7
Downward longwave radiation	bi-weekly	3.2×10^7
Downward shortwave radiation	bi-weekly	3.2×10^7
Zonal wind stress	bi-weekly	3.1×10^7
Meridional wind stress	bi-weekly	3.1×10^7

Table 8. Sensitivity of ocean state characteristics (Table 2) to various sets of control parameters. In each sensitivity experiment one subset of the control vector adjustments (Table 7) is omitted. Departures are computed relative to the state estimate, and normalized by the standard deviation of the state estimate (for mH, ..., tS) or the state estimate model-data distance (for jT, ..., jHm). Positive numbers denote percentages (for differences above 1 %) whereas parenthesized negative numbers are powers of ten (for differences below 1 %). In each sensitivity experiment one subset of the control vector adjustments is reset to zero. "Parameterized ocean" denotes GM intensity, isopycnal and diapycnal diffusivity jointly. "Sea surface forcing" denotes freshwater flux, wind stress and buoyancy flux controls jointly. "Buoyancy flux" denotes atmospheric temperature and humidity, as well as downward shortwave and longwave flux controls jointly.

experiment	jT	jS	jTs	jSs	jIs	jHa	jHm	mH	mT	mS	tV	tT	tS
all controls	369	1027	160	56	17	242	313	7925	99	5295	46	29	396
parameterized ocean	212	317	56	15	12	72	163	329	272	233	4	15	96
sea surface forcing	63	437	87	27	17	117	112	7665	252	5114	44	12	234
GM intensity	121	136	31	6	14	44	116	42	27	26	2	8	70
isopycnal mixing	44	66	6	(−2)	3	9	17	58	62	36	2	6	28
diapycnal mixing	23	44	13	4	(−2)	6	6	437	360	292	2	5	30
freshwater flux	14	182	(−3)	12	3	40	43	5140	68	3418	38	6	163
wind stress	13	22	11	(−4)	(−2)	32	26	21	28	16	2	5	24
buoyancy flux	9	30	69	5	13	9	16	2553	167	1712	5	6	50

Table 9. Model error categories as discussed in this paper.

– structural: settings that are controlled by discrete choices and switches.
– external: initial conditions, boundary conditions, and external forcing fields.
– parametric: other settings that are controlled by continuous parameters.

Table 10. List of the ECCO v4 framework components, which are fully integrated with(in) the MITgcm and its adjoint.

– the 2 km CS and LLC grids, as well as descending resolution grids
– the MITgcm estimation framework implementation charted in Fig. 5
– the model setup subjected to forward and adjoint daily regression tests
– the state estimate output, including model-data misfits
– the observational input, including weights, to the state estimate
– the forward model input needed to re-compute the 1992–2011 solution
– the testreport_ecco.m benchmarking tool to verify re-runs of the solution
– the gcmfaces Matlab framework to analyze global, gridded solutions
– the MITprof Matlab framework to process and analyze in situ profiles
– the solution's standard analysis produced by gcmfaces and MITprof

Table 11. Gridded earth variable (two dimensional) represented in Matlab as a gcmfaces object (a set of connected arrays) when the LLC90 grid is used. See also Fig. 12.

fld =		
	nFaces:	5
	f1:	[90 × 270 double]
	f2:	[90 × 270 double]
	f3:	[90 × 90 double]
	f4:	[270 × 90 double]
	f5:	[270 × 90 double]

Figure 1. Three possible approaches to gridding the Globe. Left: LL maps the earth to a single rectangular array (one *face*). Center: CS (six faces). Right: LLC (five faces). The faces of CS and LLC are color-coded; LL is only split four ways for rendering.

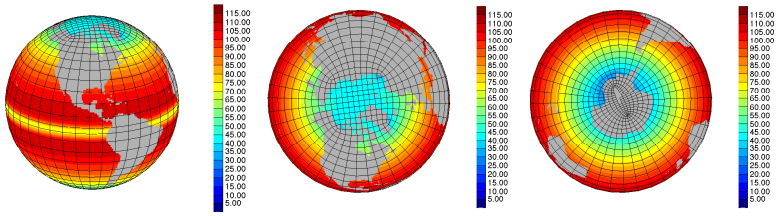

Figure 2. Average grid spacing for LLC90 (in km) computed as the square root of grid cell area. LLC90 denotes the LLC grid with 90 grid points as the common face dimension (one quarter of earth circumference at the equator).

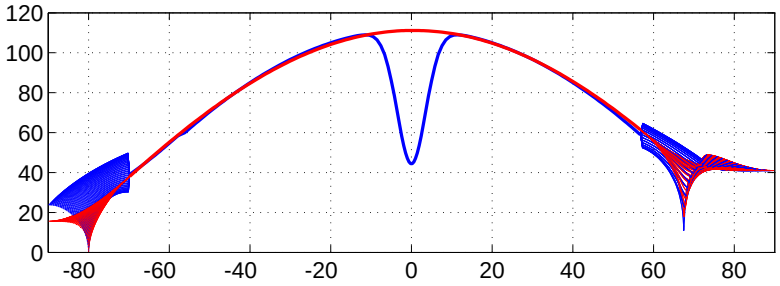

Figure 3. Grid spacing details for LLC90 as function of latitude, in km. Between 70° S and 57° N sector, blue and red curves show meridional and zonal grid spacing, respectively. Poleward of 70° S and 57° N, grid lines deviate from meridians and parallels, and LLC becomes zonally asymmetric (see Fig. 2), leading to the depicted grid spacing ranges.

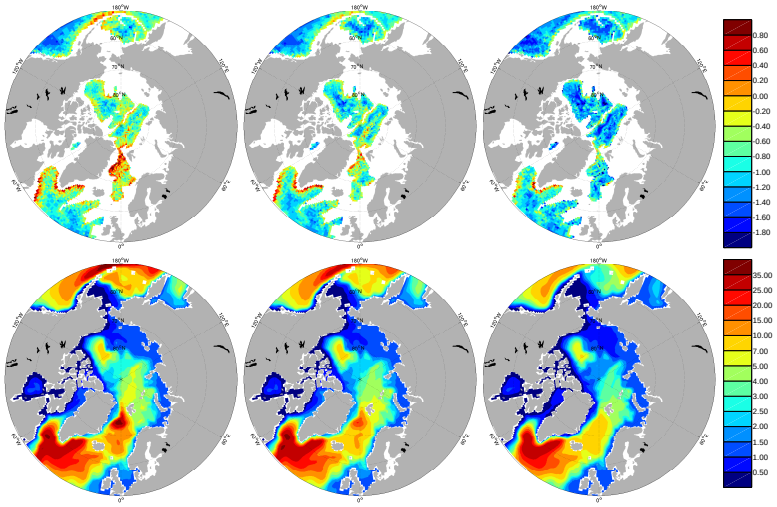

Figure 4. Root mean squared vertical velocity at 2000 m depth (top, in mm day^{-1} log scale) and in horizontal streamfunction (bottom, in Sv) in three solutions. Left panels: ECCO-Production, release 1. Middle panels: same as left panels but with increased horizontal viscosity near the ocean floor. Right panels: same as left panels, but with C-D scheme.

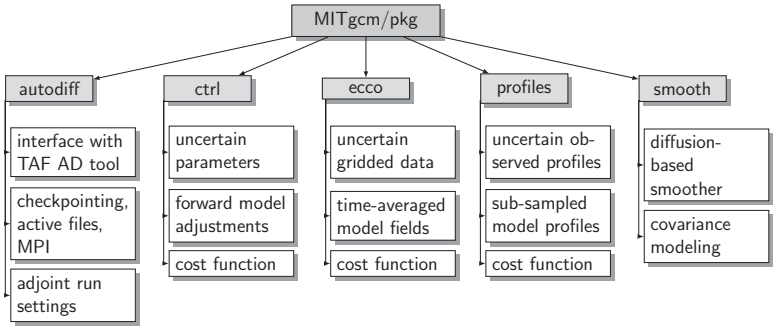

Figure 5. Organization and roles of MITgcm estimation modules. A more complete presentation of MITgcm modules (or packages) can be found in the manual. The algorithmic differentiation (AD) tool being currently used is TAF. The handling of checkpoints and active files is described in Heimbach et al. (2005). The roles of ctrl, ecco, profiles and smooth in forward mode (as reported here) and in adjoint mode are further described in the text. In preparation for this paper, much redundancy was eliminated through generalization of "ecco" and "ctrl" features (see, e.g. Table 4). This process reduced by about 30 % the overall volume of adjointed code in ECCO v4 (counting the entire ocean-seaice model).

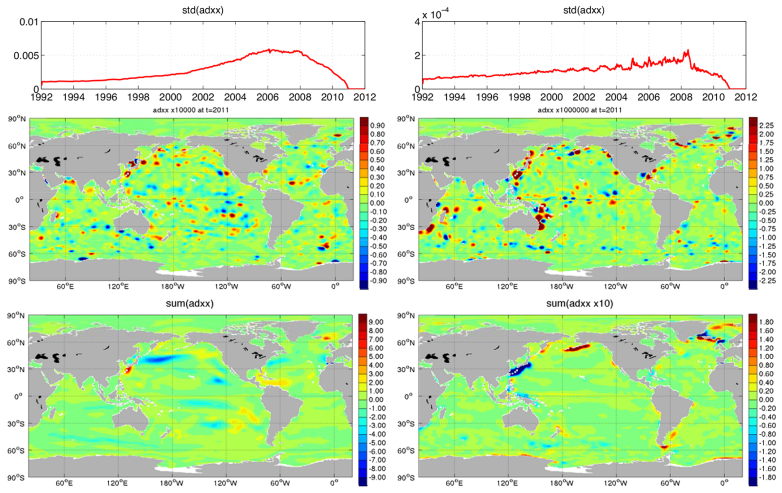

Figure 6. Non-dimensional adjoint sensitivity ($\frac{\partial J}{\partial u}$ per Eq. 12 with $\alpha_i = 1$, $\beta_j = 0$) of the Argo penalty over 2008–2010 (jT + jS) to biweekly GM intensity (left) and Laplacian horizontal viscosity (right) at 300 m. Computations are carried out with added bottom viscosity (as done in Table 3) to improve stability of the adjoint. Top panels show the standard deviation (over all ocean grid points) of biweekly adjoint sensitivities, reflecting their forcing by model-data misfits over 2008–2010 and their propagation backward to 1992 (propagation here may involve persistence, advection, waves, dissipation, etc.). Middle panels shows the first biweekly period of adjoint forcing (in late 2010) and reflects the most patchy, short term sensitivity to model parameters. The patches smoothness is mostly set by the Weaver and Courtier (2001) spatial correlation model, with a specified scale of 3 times the grid scale. Bottom panel show the total sensitivity to time mean model parameters, obtained by summing over all biweekly period. Their broad patterns primarily reflect the aggregation of adjoint sensitivities due to the 20 year integration of Eqs. (1)–(5).

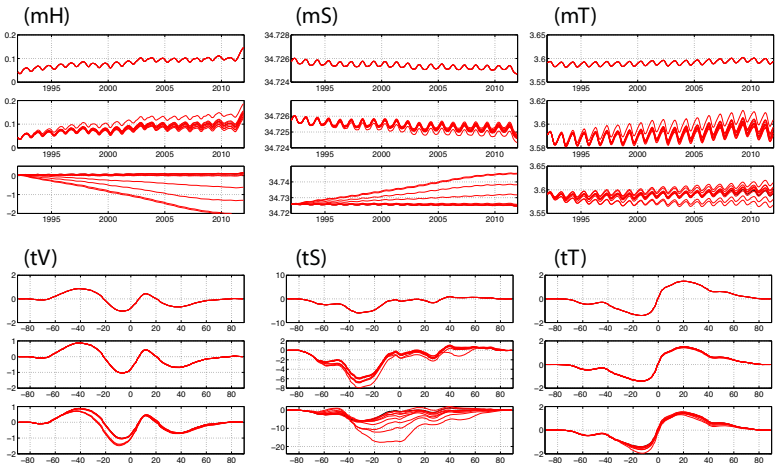

Figure 7. Depiction of ocean state characteristics (defined in Table 2) that reflect the 20 year solution sensitivity to various model settings (experiments in Tables 3 and 8). The respective units are: m (mH), g kg^{-1} (mS), °C (mT), Sv (tV), g kg^{-1} Sv (tS) and PW (tT). For each variable, the top, middle and bottom panel respectively depict results of regression tests (top three rows of Table 3), structural model sensitivity (bottom eight rows of Table 3) and parametric model sensitivity (Table 8). The state estimate itself is shown as a thick black curve. All other runs, which are only discussed collectively, are shown as thin red curves. Note that mH,mS, and mT are global mean time series, whereas tV, tS, and tT are 2008–2010 averages and are functions of latitude.

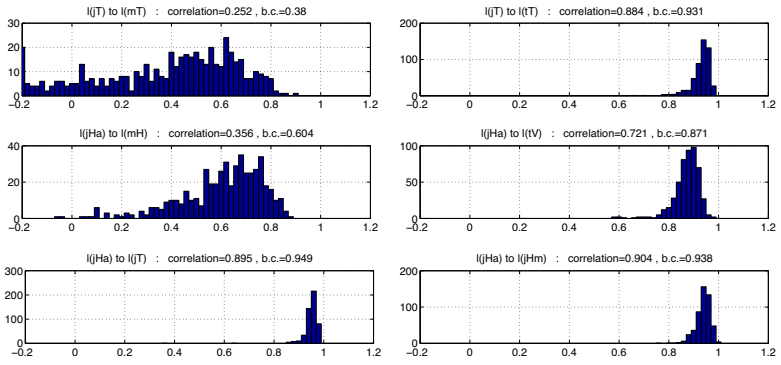

Figure 8. Bootstrap distribution of a cosensitivity index between ocean state characteristics. For two characteristics a and b, the cosensitivity index is defined as the correlation between $\log_{10}(a)$ and $\log_{10}(b)$ and is computed for a set of 15 elements (the last 9 rows of Table 3 and the last 6 rows of Table 8 values). Here $\log_{10}(a)$ and $\log_{10}(b)$ denotes the $\log_{10}$ of the distance between model solutions for characteristic a and b, respectively. Values reported in Table 3 as negative integers are examples of $\log_{10}(a)$ and $\log_{10}(b)$. Bootstrap resampling (500 members) gives the displayed distributions.

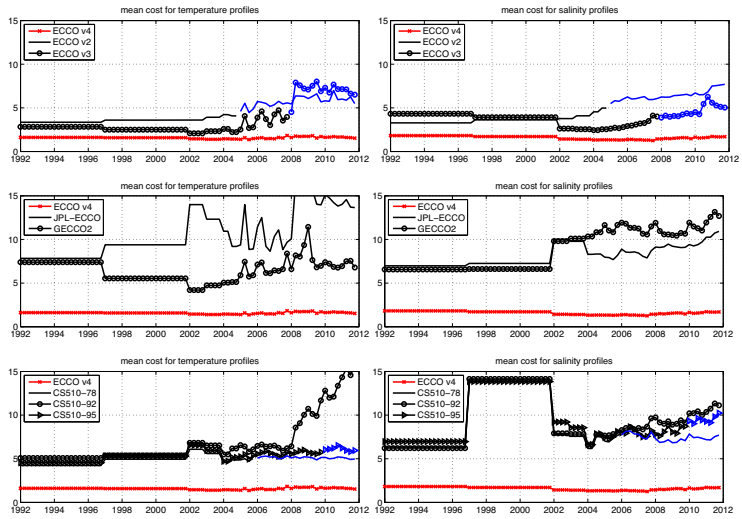

Figure 9. Distance to in situ observations (Table 5; Sect. 4; Appendix D) for various MITgcm solutions; for temperature (left; jT) and salinity (right; jS); as a function of year. The red line with star symbols (in each panel) shows the baseline ECCO v4 solution (the "ECCO-Production, release 1" state estimate). Also shown in top panels: ECCO v2 (Wunsch and Heimbach, 2007, 1992–2004) and ECCO v3 (Wunsch and Heimbach, 2009, 1992–2007); middle panels: JPL-ECCO (Fukumori, 2002, 1992–Present) and GECCO2 (Köhl, 2014, 1948–2011); bottom panels: three ECCO2 eddying solutions using different forcing fields (courtesy of H. Zhang). For each solution, monthly mean output was subsampled at data locations. For solutions that do not extend through 2011, state of the last full year was replicated afterwards (shown in blue). Temporal resolution in jT, jS is reduced in years before 2005, when the Argo array reached near-global deployment. The pre-2005 values, while still a useful indication of skill, may be characteristic of limited regions. The Argo period values are mostly representative of the upper 2000 m of the global ocean.

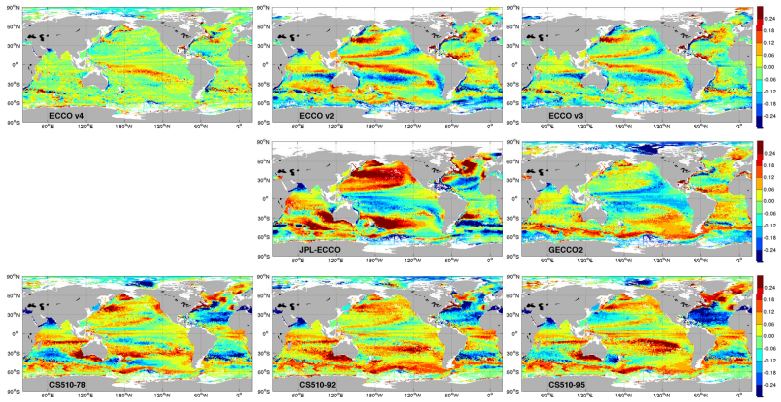

Figure 10. Model-data misfits for salinity at 300 m depth (sample average over all times) for ECCO v4 (top left panel), ECCO v2 (top center), ECCO v3 (top right), JPL-ECCO (middle center), GECCO2 (middle right) and three ECCO2 simulations (bottom panels). Additional computational details are reported in the Fig. 9 caption, along with references for the various solutions.

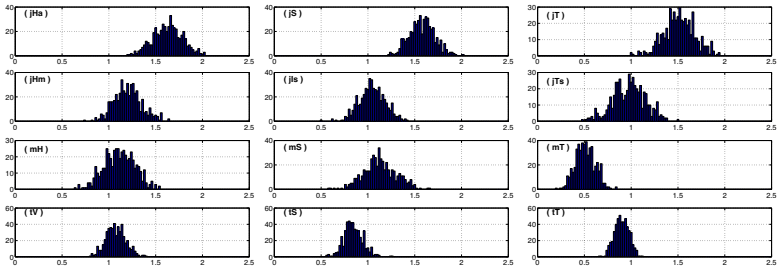

Figure 11. Bootstrap distribution for the controllability index (defined as C hereafter) of 1992–2011 ocean state characteristics. For a characteristics b, the quantity $C = \log_{10}(b_i/b_j)$ is computed for a set of 54 experiment pairs formed from the last 9 rows of Table 3 values (b_j) and the last 6 rows of Table 8 values (b_i). Values reported in Table 3 as positive integers and parenthesized negative integers are examples of b_j and $\log_{10}(b_j)$, respectively. Bootstrap resampling (500 members) gives the displayed distributions.

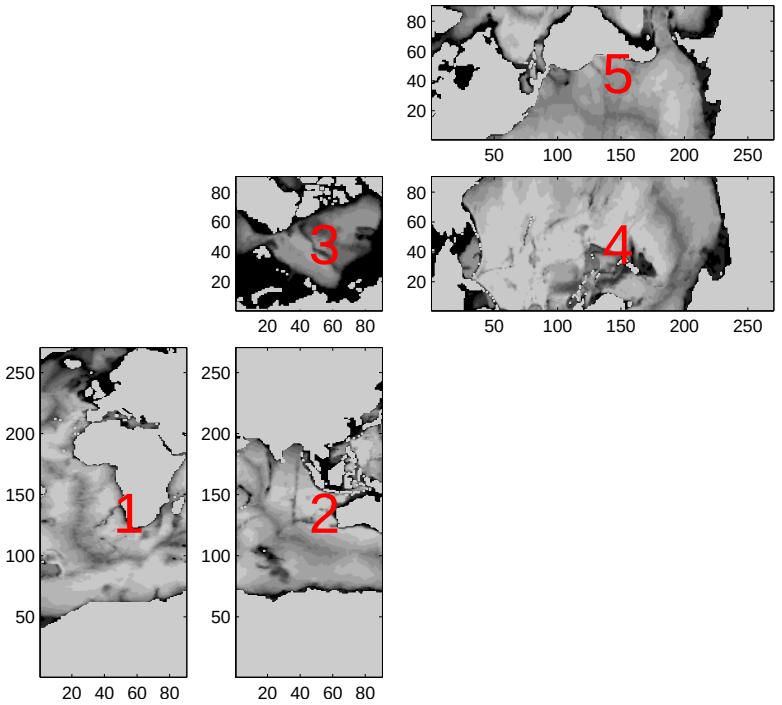

Figure 12. Example of a field (ocean bathymetry) mapped to the LLC90 grid and displayed in a way that reflects the MITgcm layout of the LLC90 grid (Fig. 1, right panel). The five grid "faces" number are indicated in red, and their dimensions are shown in black. See also Table 11.

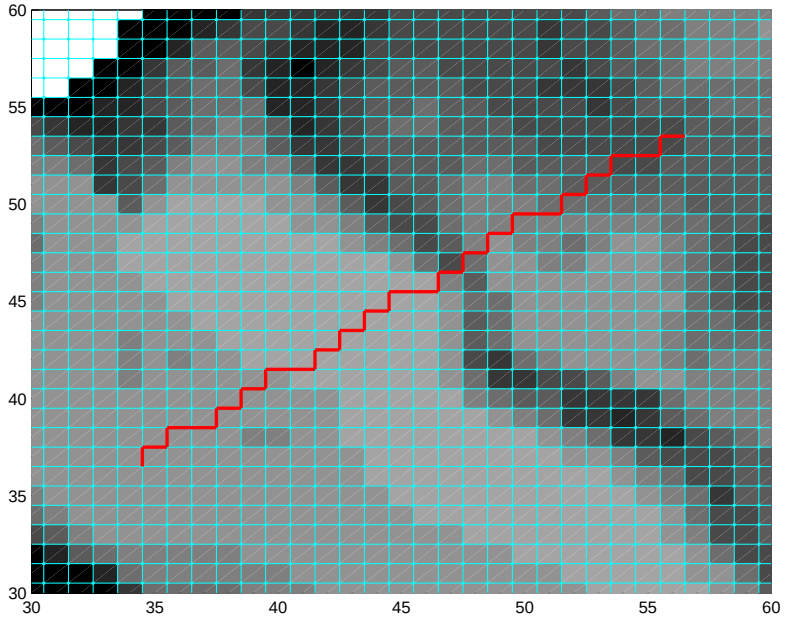

Figure 13. Example of a grid line path (in red) that approximates a great circle between 45° E, 85° N and 135° W, 85° N (a meridian crossing the north pole). Location: central part of face 3 from Fig. 12. Shading: ocean bottom depth. Blue lines: grid cell edges.

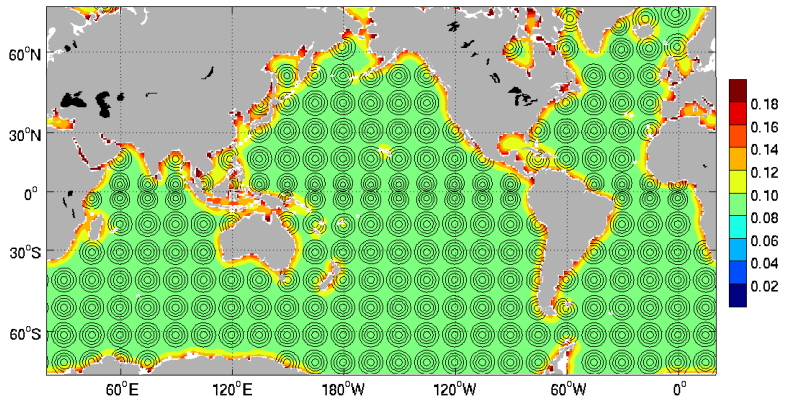

Figure 14. Diffusion applied to grid scale noise (set to unit variance) introduces correlation (contours, drawn for select points) and yields a reduced noise variance (color shading). The smoothing scale was set to three grid points.

```
netcdf argo_feb2013_2008_to_2010 {

...
        double prof_T(iPROF, iDEPTH) ;
                prof_T:long_name = "potential temperature" ;
                prof_T:units = "degree C" ;
        double prof_Tweight(iPROF, iDEPTH) ;
                prof_Tweight:long_name = "least-square weight" ;
                prof_Tweight:units = "(degree C)^-2" ;
        double prof_Testim(iPROF, iDEPTH) ;
                prof_Testim:long_name = "pot. temp. estimate" ;
                prof_Testim:units = "degree C" ;
...
        double prof_depth(iDEPTH) ;
        double prof_YYYYMMDD(iPROF) ;
        double prof_HHMMSS(iPROF) ;
        double prof_lon(iPROF) ;
        double prof_lat(iPROF) ;
        char prof_descr(iPROF, lTXT) ;
                prof_descr:long_name = "profile description" ;
...
}
```

Figure 15. Netcdf file header illustrating the MITprof format used in MITgcm/pkg/profiles.