Towards eddy permitting estimates of the global-ocean and sea-ice circulations

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

The formal combination of global-ocean and sea-ice models with remotely-sensed and in-situ data can yield significant insight into the climate system. Making such state estimates for decadal or longer periods with models that can resolve the energetic ocean meso-scale motions and the dominant sea-ice dynamics is both mathematically and computationally challenging. In this article we describe three recent advances that bring eddy-permitting, decadal time-scale estimates of the three-dimensional global ocean and sea-ice circulations within reach: 1) the configuration of an efficient meso-scale-eddy permitting global ocean and sea-ice model that achieves a throughput approaching ten years of model integration per day of computation, 2) the demonstration that initial conditions and surface forcing fields estimated at coarse resolution can improve the solution of an eddy-permitting model, and 3) the development and deployment of a hierarchy of methods for assimilating observations in a mathematically rigorous way.

Satellite and in-situ observations are now routinely combined with numerical models in order to estimate the time-evolving oceanic circulation, addressing a wide variety of operational and research problems. For climate research, what is required is a dynamically consistent synthesis of all available observations over the last several decades, or longer, with the best possible numerical model. Rigorous low-resolution estimates of ocean circulation are already possible using the existing data base, modeling, and estimation capabilities. But the low-resolution estimates lack the ability to resolve many important oceanic processes, for example, flow over narrow sills, western boundary currents, regions of deep convection, and meso-scale eddies, that are important both for operational applications and for developing fundamental understanding of the climate system.

Similar to atmospheric reanalyses, which are used by climate scientists in lieu of raw data, ocean state estimation provides analyses of the ocean circulation that combine the accuracy of observations with the completeness of numerical models. On the one hand, models are imperfect because of errors in the representation of unresolved processes, in the discretization algorithms, in the initial and boundary conditions, etc. On the other hand, in-situ observations made from ships, moored arrays, and autonomous underwater devices are sparse, and satellite observations only provide information about near-surface or depth-integrated variables such as sea-surface height, temperature, color, wind stress, and the gravity field. Therefore, a complete description of ocean circulation can only be obtained by constraining these imperfect models with available satellite and in-situ observations.

Conceptually, the problem of constraining numerical ocean models in a mathematically rigorous way with observations is akin to fitting data to a curve using, for example, a least-squares approach. Of course the oceanic problem is huge and more complicated because of the large numbers and complex relationships of the model variables. Nevertheless, during the past few years, the consortium for Estimating the Circulation and Climate of the Ocean (ECCO) has demonstrated that it is feasible to carry out these computations in a routine manner and that the resulting ocean circulation estimates possess significant skill [*Stammer et al.*, 2003]. A unique feature that distinguishes the ECCO analyses from others is their physical consistency on decadal and longer time scales. Model errors are explicitly ascribed to initial conditions, boundary conditions, diffusivity coefficients, and other model parameters. By comparison, atmospheric reanalyses and most other ocean-circulation estimates exhibit temporal discontinuities every time new data are assimilated. Solutions that contain discontinuities are adequate for prediction and for operational objectives. But long-time-scale planetary monitoring and scientific discovery applications are better served by rigorous circulation estimates, such as those produced by ECCO, because property budgets are closed and because information from the entire observational record are utilized at each estimation time.

The ECCO ocean circulation analyses are freely available on-line and they are used to address a wide variety of science questions (http://www.ecco-group.org/). The results demonstrate the potential of long-term, full-ocean depth state estimates as a primary tool both for monitoring internal processes as well as for estimating time dependent fluxes of heat, freshwater, momentum, and biogeochemical tracers into and out of the ocean. But the computational demands of ocean state estimation are enormous, limiting the existing ECCO analyses to horizontal grid spacings of order 100 km. Also for computational efficiency reasons, current generation ECCO analyses exclude the Arctic Ocean and lack an interactive sea-ice model, which restricts the utilization of satellite data over polar regions.

At first glance, the objective of rigorous global-ocean and sea-ice state estimation, for a decade or more, at eddy-permitting resolutions, and for the full ocean depth seems impossible. Depending on the method and on the approximations that are used, the computational cost of state estimation is several dozen to several thousand times more expensive than integrating the ocean model without state estimation. Therefore a necessary condition for global, eddy-permitting state estimation is the availability of an efficient model and of significant computational resources.

Cubed-Sphere Model Configuration on a Parallel Supercomputer

To address the above challenges and to more fully utilize the available satellite data, NASA has committed significant new computational support to the ECCO project. The objective is to develop a next-generation, high-resolution ocean state estimation system that 1) is truly global in scope, encompassing polar ocean regions with extensive sea-ice coverage and extending over the full depth of the ocean, 2) can exploit the full span of observational measurements, both in-situ and remotely sensed, that are available today and that will become available in the future, and 3) is algorithmically and computationally constructed to permit routine decadal and longer state estimates at horizontal and vertical scales that explicitly represent the energetic ocean meso-scale and key water mass formation and transformation processes. Below, we describe some early capabilities of this new initiative.

Gridding a sphere completely presents a challenge for time dependent numerical simulation. Polar singularities of the conventional latitude-longitude grids result in unacceptably small grid cell spacings near the Poles. The ECCO ocean state estimation infrastructure is based on the Massachusetts Institute of Technology General Circulation Model (MITgcm [*Marshall et al.*, 1997]), which supports unstructured, curvilinear horizontal grids. For the work discussed herein, a novel, semi-structured cubed-sphere grid projection is employed (see Fig. 1). This projection permits relatively even grid-spacing throughout the model domain, it preserves local orthogonality for efficient and accurate time stepping of the model equations, and it avoids the polar singularities [*Adcroft et al.*, 2004]. Each cube face on Fig. 1 has 510 elements on a side for a nominal horizontal grid spacing of 19 km. Cube faces may be divided into any number of sub-domains that fit into the grid point dimensions of the cube faces. These sub-domains can run on separate processors individually or in groups, which provides for manual compile-time load balancing across a relatively arbitrary number of processors. Finally, sub-domains that do not contain any wet points can be removed from the computation domain, hence providing optimal mapping of the irregular computational domain to the underlying hardware.

The ocean model is coupled to an interactive sea-ice model (see Fig. 2). The sea-ice model includes a thermodynamic component that simulates ice thickness, ice concentration, and snow cover, as in *Zhang et al.* [1998]. Sea-ice dynamics are modeled using a viscous-plastic rheology [*Zhang and Hibler, III*, 1997]. New for this work is an efficient parallel implementation of the line-successive-relaxation solver on the cubed-sphere grid. The inclusion of an interactive sea-ice model provides for more realistic surface boundary conditions in polar regions and allows the model to be constrained by satellite observations over ice-covered oceans. The sea-ice model also provides the ability to estimate the time-evolving sea-ice thickness distribution and to quantify the role of sea ice in the global ocean circulation.

The results of Figs. 1 and 2 were obtained on a 512-processor, shared-memory SGI Altix system operated by the NASA Advanced Supercomputing group at the Ames Research Center (NAS/ARC).

At the time of writing, twenty such systems are being clustered together at NAS/ARC as part of Project Columbia, for a combined peak capacity of 61 teraFLOPS, 50% more capacity than Japan's Earth Simulator. The shared memory architecture of the SGI Altix, the supportive computational resource culture at NAS/ARC, and the advanced numerics and parallelization capabilities of the MITgcm have allowed the ECCO team to configure an eddy-permitting global-ocean and sea-ice model that achieves a throughput approaching ten years of model integration per day of computation. With this fast throughput, eddy-permitting estimates of the global ocean and sea-ice circulations are within reach and the focus of attention can now shift towards estimation methodology. Below we present some early results of the eddy permitting estimation effort.

Low-Resolution Surface Flux Estimates

A first question that has been addressed is whether the existing, low-resolution ECCO estimates of initial and surface boundary conditions can be used to initialize eddy-permitting estimation efforts. For this purpose two 1992-2002 integrations were conducted using a near-global configuration with 1/4-degree horizontal grid spacing [*Menemenlis et al.*, 2004a]. The first integration is initialized from the World Ocean Database [*Conkright et al.*, 1999] and forced by surface fluxes (wind stress, heat, and freshwater) from the NCEP meteorological reanalysis [*Kistler et al.*, 2001]. Initial conditions and surface fluxes for the second integration are from the ECCO 1-degree, adjoint-method optimization [*Stammer et al.*, 2004]. In addition to the specified surface fluxes, both integrations also include surface relaxation terms to observed sea-surface temperature and salinity. On average, the NCEP-forced integration requires time-mean temperature relaxation fluxes on the order of ±30 W/m^2 while the time-mean temperature relaxation fluxes for the ECCO-forced integration are substantially less, order 10 W/m^2. Overall, the magnitude of the time-mean surface salinity relaxation fluxes are also smaller for the ECCO integration than they are for the NCEP integration. The smaller surface relaxation fluxes demonstrate the accuracy and the robustness of the ECCO estimates, in spite of differences in the representation of meso-scale eddies and of other physical processes.

The NCEP and the ECCO eddy-permitting simulations were also compared to the complete suite of observations that were used in the coarse-resolution ECCO optimizations [*Menemenlis et al.*, 2004a]. While the ECCO forcing seems to degrade the skill in estimating observed sea-surface height variability in some regions, it generally improves the time-mean and the variability of upper ocean temperature and salinity (see Fig. 3). The assimilated forcing also improves the paths of the Gulf Stream and of the Kuroshio, and the strength of the Equatorial Undercurrent. These results indicate that boundary conditions estimated at coarse resolution can improve the solution of eddy-permitting models. Next we sketch a preliminary strategy towards rigorous, eddy-permitting estimates of the global-ocean and sea-ice circulations.

Towards Eddy-Permitting Estimates

Low-resolution ECCO ocean circulation analyses have been obtained using three rigorous estimation approaches: the adjoint-model method [*Stammer et al.*, 2003], an approximate Kalman filter [*Fukumori*, 2002], and an approach based on the computation of model Green functions [*Menemenlis et al.*, 2004b]. There is some limited experience in applying the adjoint method to a regional eddy-permitting model configuration [*Gebbie*, 2004] and work is underway to extend the adjoint method to global coarse-resolution and to regional high-resolution model configurations that include sea-ice. Some preliminary estimation results have also been obtained in applying the *Fukumori* [2002] filter to the eddy-permitting cubed-sphere configuration. But at present this filter is limited to the estimation of adiabatic, time-varying model errors. While work continues in developing the adjoint-method and approximate Kalman filter approaches, preliminary, eddy-permitting estimates can be obtained using a Green function approach, described next.

Green functions provide a simple yet effective method to test and to calibrate GCM parameterizations, to study and to quantify model and data errors, to correct model biases and trends, and to blend estimates from different solutions and data products. At the most basic level, the Green function approach involves the computation of model sensitivity experiments followed by a recipe for constructing a solution that is the best linear combination of these sensitivity experiments. Compared to other methods, the key advantages of Green function approaches are simplicity of implementation, inherent parallel scalability, and robustness in the presence of non-linearities. A Green function approach has been applied to one of the ECCO configurations, using a total of twenty-six sensitivity experiments, and resulting in substantial improvements of the solution relative to observations as compared to prior estimates [*Menemenlis et al.*, 2004b]. Overall model bias and drift were substantially reduced and there was a 10% to 30% increase in explained variance. This solution is the backbone of the ECCO quasi-operational, ocean-circulation analysis, which is updated every ten days using an approximate Kalman filter. See http://ecco.jpl.nasa.gov/external/ for a description and for output.

In conclusion, the key ingredients for eddy-permitting estimates of global-ocean and sea-ice circulations are now in place. This includes the modeling and computational infrastructure as well as a range of estimation methodologies. The focus of ocean state estimation during the past five years has been to demonstrate the feasibility and utility of rigorous, global, sustained estimates, with considerable success for upper ocean and for equatorial processes. The scientific challenges during the next five years, for example, quantifying the role of the ocean in the global carbon cycle, understanding polar-subpolar interactions, and quantifying the time-evolving term balances within and between different components of the Earth system, require much improved accuracy in the estimation of water mass formation and transformation rates, mixed layer depths, and high-latitude processes. The accurate monitoring of these processes in turn requires developing state estimation machinery, of the sort we have described in this article, that can fully capitalize on advances in computational and observational technologies.

Acknowledgments

The authors acknowledge support of the ECCO project through the National Oceanographic Partnership Program, with funding from NASA, NSF and ONR. The ECCO consortium also acknowledges computational support through NSF (NPACI, NCAR), ONR and NASA. We thank W. Brooks, E. Lindstrom, T. Lee, J. Marshall, and W. Thigpen for their support. We thank C. Wunsch for helpful discussions and comments on the draft article.

References

Adcroft, A., J.-M. Campin, C. Hill, and J. C. Marshall, Implementation of an atmosphere-ocean general circulation model on the expanded spherical cube, *Mon. Weather Rev.*, Available at http://mitgcm.org/˜adcroft/papers/adcroft_et_al_MWR_2004.pdf, 2004.

Conkright, M. E., et al., World ocean database 1998, Internal Report 14, National Oceanographic Data Center, 1999.

Fukumori, I., A partitioned Kalman filter and smoother, *Mon. Weather Rev.*, *130*, 1370–1383, 2002.

Gebbie, G., Subduction in an eddy-resolving state estimate of the northeast atlantic ocean, Ph.D. thesis, Massachusetts Institute of Technology and Woods Hole Oceanographic Institution, Available from http://www.mit.edu/˜gebbie/, 2004.

Kistler, R., et al., The ncep-ncar 50-year reanalysis: Monthly means cd-rom and documentation., *Bull. Amer. Meteoro. Soc*, *82*, 247–268, 2001.

Marshall, J., A. Adcroft, C. Hill, L. Perelman, and C. Heisey, A finite-volume, incompressible Navier-Stokes model for studies of the ocean on parallel computers, *J. Geophys. Res.*, *102*(C3), 5753–5766, 1997.

Menemenlis, D., I. Fukumori, A. Köhl, T. Lee, and D. Stammer, The impact of ECCO surface forcing fields on the results of a global 1/4-deg model, The ECCO Report Series 30, Scripps Instituttion of Oceanography, 9500 Gilman Dr., La Jolla, CA 92093-0239, Available at http://www.ecco-group.org/reports.html, 2004a.

Menemenlis, D., I. Fukumori, T. Lee, and B. Tang, Using Green's functions to calibrate an ocean general circulation model, *Mon. Weather Rev.*, Available at http://ecco.jpl.nasa.gov/~dimitri/articles/, 2004b.

Stammer, D., C. Wunsch, R. Giering, C. Eckert, P. Heimbach, J. Marotzke, A. Adcroft, C. N. Hill, and J. Marshall, Volume, heat, and freshwater transports of the global ocean circulation 1993-2000, estimated from a general circulation model constrained by World Ocean Circulation Experiment (WOCE) data, *J. Geophys. Res.*, *108*(C1), 3007–3029, 2003.

Stammer, D., K. Ueyoshi, A. Köhl, W. G. Large, S. A. Josey, and C. Wunsch, Estimating air-sea fluxes of heat, freshwater, and momentum through global ocean data assimilation, *J. Geophys. Res.*, *109*, C05023, 2004.

Zhang, J., and W. D. Hibler, III, On an efficient numerical method for modeling sea ice dynamics, *J. Geophys. Res.*, *102*, 8691–8702, 1997.

Zhang, J., W. D. Hibler, III, M. Steele, and D. A. Rothrock, Arctic ice-ocean modeling with and without climate restoring, *J. Phys. Oceanogr.*, *28*, 191–217, 1998.

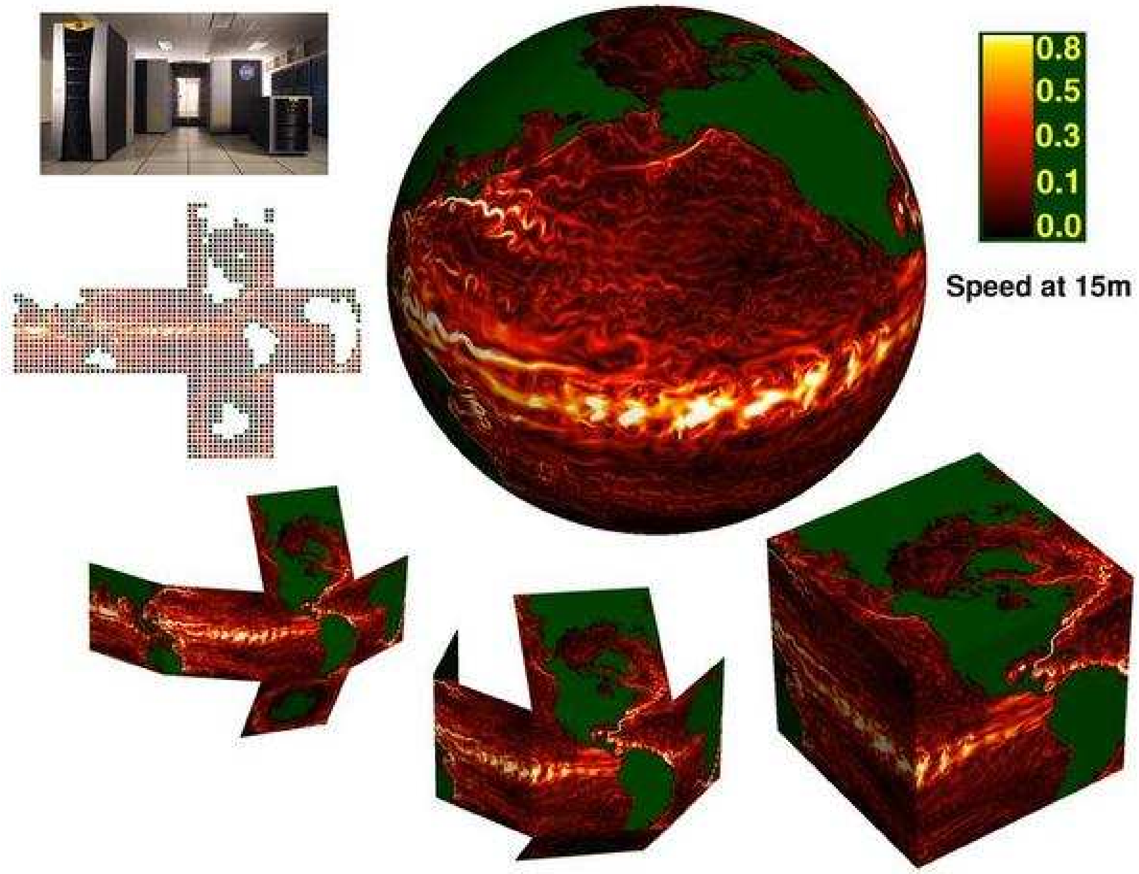

Figure 1: Cubed-sphere ocean model configuration. The figure shows near-surface (15-m) ocean-current speed from an eddy-permitting integration of the cubed-sphere. Units are m/s. Each face of the cube comprises 510 by 510 grid cells for a mean horizontal grid spacing of 19 km. The integration was carried out at the NASA Ames Research Center on a 512-CPU SGI Altix, which is shown on the upper-left panel. The second panel illustrates the innovative tiling strategy, which excludes dry tiles from the computation domain. The remaining panels illustrate the cubed-sphere grid configuration. Cubed-sphere animations, and more information about this integration, are available at http://ecco.jpl.nasa.gov/cube_sphere/.

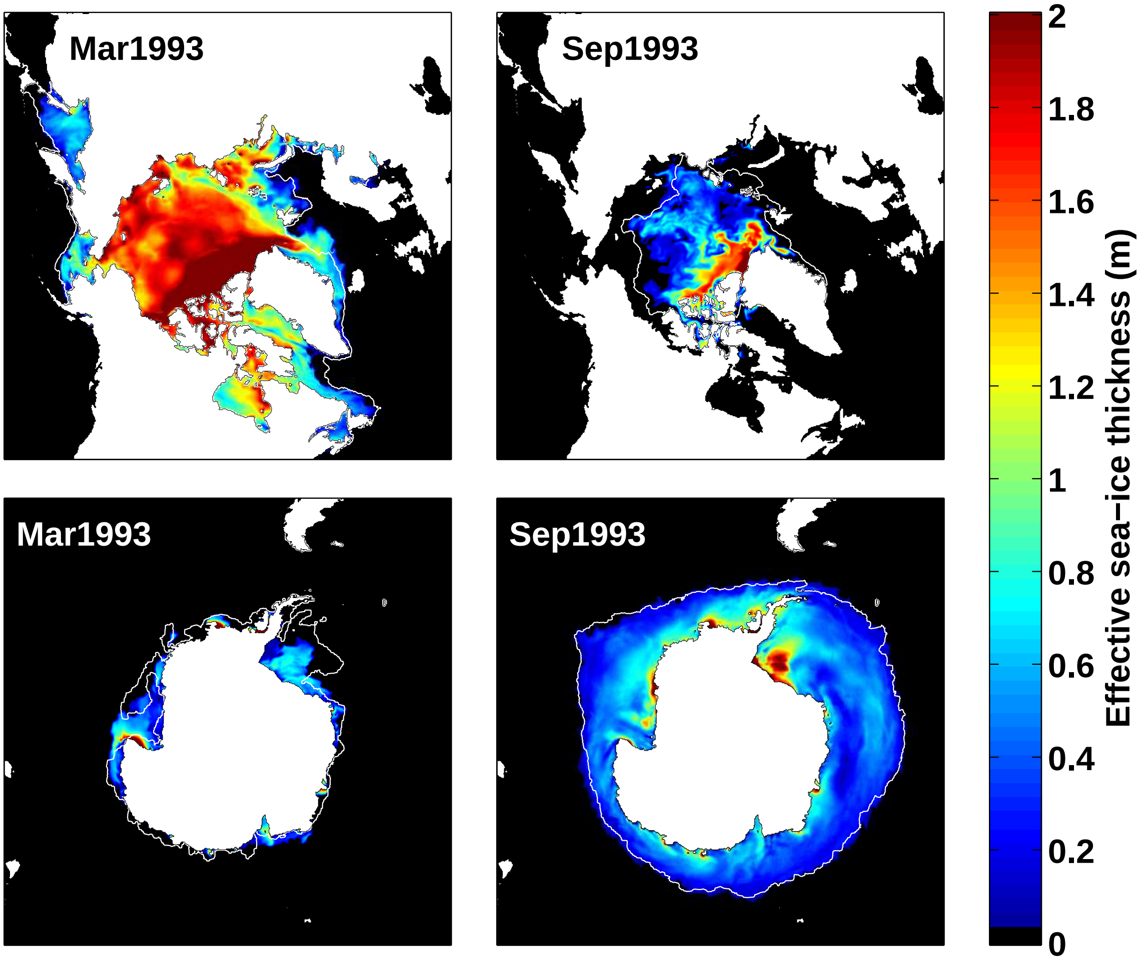

Figure 2: Sea-ice on the cubed-sphere. The figure shows snapshots of simulated effective ice thickness (ice thickness times concentration) for the same cubed-sphere integration that is depicted in Fig. 1. The top and bottom panels of the figure are the Northern and Southern faces of the cube, respectively. The thin white line represents observed sea-ice extent (15% concentration) from passive microwave radiometers. Animations of these panels for the complete integration period are available at http://ecco.jpl.nasa.gov/cube_sphere/IceCube/. The difference between observed and simulated sea-ice extent, for example, excessive summer melting and unrealistic open water polynyas in the Ross and Weddell Seas during the later years of this integration, is one of the signals that we propose to assimilate in order to improve the model representation of high latitude processes.

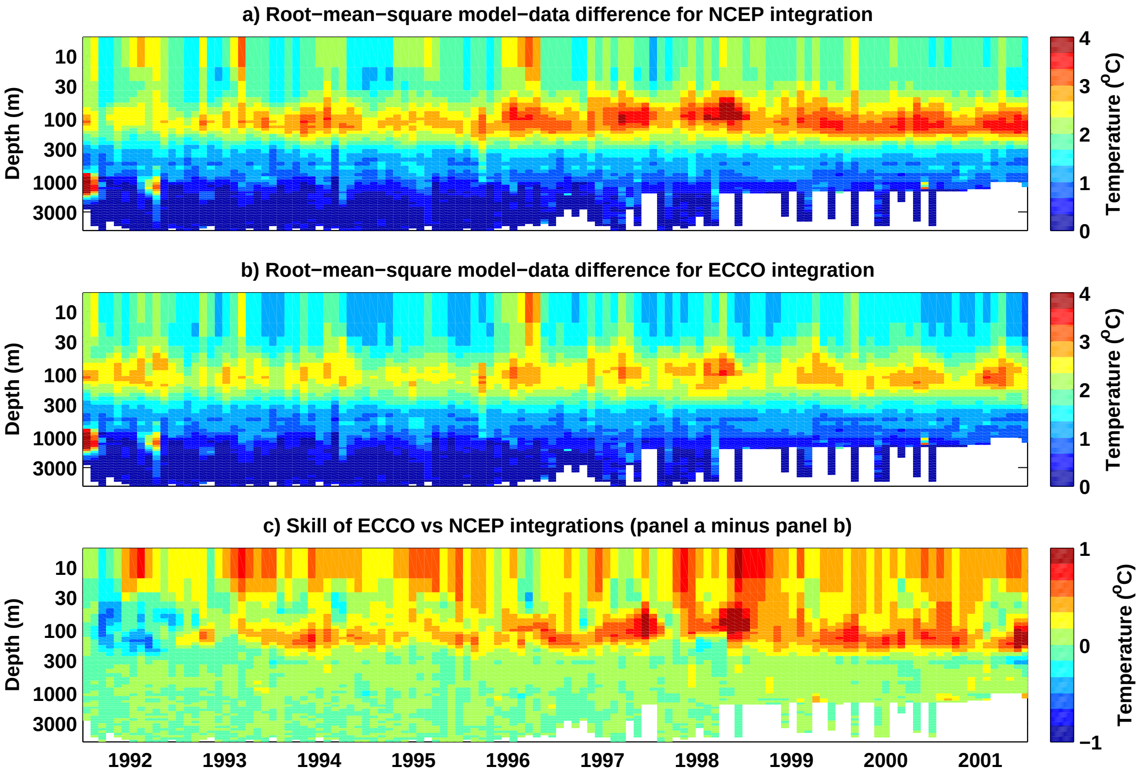

Figure 3: Globally averaged root-mean-square (rms) difference between simulation results and observations of temperature. The data are a compilation from CTD, XBT, moored-array, and autonomous-float measurements. The top panel shows rms difference between the data and an eddy-permitting integration forced by NCEP reanalysis surface fluxes. The middle panel shows rms difference between the data and an eddy-permitting integration forced by the ECCO surface fluxes. The bottom panel shows the difference between the first two panels, positive numbers indicating that the ECCO forced simulation has more skill in simulating the observed temperature than the NCEP forced simulation.