Contents

Abstract

A consortium of university and national laboratory partners propose to make a major qualitative improvement in the resolution and accuracy of estimates of the time-evolving global-ocean and sea-ice circulations by harnessing NASA computational resources in order to combine (i) the voluminous new global observations from satellites and in-situ instruments with (ii) a general circulation model that incorporates novel gridding techniques and parallel-computing technologies using (iii) rigorous state estimation methods at hitherto unachievable resolutions. Specifically, a mesoscale-eddy-resolving and sea-ice model will be constrained to all relevant, global-scale, NASA and other data streams to obtain a best-possible synthesis of the global-ocean and sea-ice circulations during the ocean satellite era (1978-present) for applications in climate research and forecasting. The work aims to obtain estimates that satisfy the model time-evolution equations and in which the error field has been propagated through the same model as the state vector.

The proposed work is made possible by recent advances in ocean state estimation technology (ECCO: Estimating the Circulation and Climate of the Ocean), in modeling infrastructure (ESMF: Earth System Modeling Framework), in computational fluid dynamics (MITgcm: Massachusetts Institute of Technology general circulation model), and in automatic differentiation (AD). The following strands will be developed: (i) very high resolution simulations of ocean circulation and sea-ice distribution using the MITgcm, (ii) assimilation of satellite and in-situ global data sets at mesoscale eddy-resolving resolution using rigorous state estimation methods, (iii) detailed evaluation of resulting ocean-circulation and sea-ice estimates, including comparisons with existing coarse-resolution estimates and quantitative statements about residual uncertainties, (iv) demonstration science applications, including analysis of heat budgets, studies of climate variability, sensitivity, and predictability, and estimates of error covariance matrices for the Global Modeling and Assimilation Office (GMAO) seasonal-to-interannual prediction effort, and (v) scientific visualization and distribution of the high-resolution ocean and sea-ice state estimates and of ESMF-compatible modeling and estimation software to the community.

The proposed work will strengthen the connection between NASA's modeling and remote sensing strategies and will enhance the value of NASA satellite retrievals for studies and discussions about climate and climate variability, and many related societal concerns.

1 Introduction and objectives

In the past five years, global ocean state estimation has matured to the extent that estimates of the evolving circulation of the ocean constrained by in-situ and remotely sensed global observations are now routinely available and being applied to myriad scientific problems. These products represent a huge effort by and on behalf of the oceanographic community and have involved scientists, engineers, observationalists, and modelers of many different shades working cooperatively together.

One such effort has been the consortium for Estimating the Circulation and Climate of the Ocean (ECCO), including many of the authors of the present proposal. The ECCO group, funded by the National Oceanographic Partnership Program (NOPP) during the period 1998-2003, has focused on developing the models and methodologies to bring global state estimation into being, as briefly reviewed in Section 2. Much has been achieved, but the existing estimates lack an Arctic Ocean and interactive sea-ice and, because of inadequate resolution, do not resolve geostrophic turbulence (mesoscale eddies), narrow passages and sills in the bathymetry, boundary currents, and jets. The lack of resolution compromises the fidelity of the constrained solutions and is a severe limitation on their usefulness, both in understanding the circulation, and in any attempt to forecast climate. Lack of the Arctic Ocean and sea-ice limits the usefulness of the ECCO products in describing and studying polar-subpolar interactions.

This proposal, ECCO-II, seeks to extend the scope, resolution, and accuracy of the existing ECCO ocean syntheses (similar to meterological reanalyses) via explicit representation of mesoscale eddies and of sea-ice processes. Specifically, the proposed work aims to harness NASA computational resources, advances in computational fluid dynamics and software engineering, and the ability to solve massive control problems in order to:

1. carry out very high resolution simulations of ocean circulation and sea-ice distribution,
2. constrain models of ocean circulation and sea-ice using global satellite and in-situ data sets at much higher resolutions than hitherto possible,
3. evaluate and apply the resulting ocean and sea-ice state estimates to key science problems, and
4. make these state estimates, new modeling tools, and related estimation software, available to the community.

The significance of our proposal is that we are attempting to describe and understand the global ocean circulation in the presence of its vigorous geostrophic eddy field, by constraining a high resolution ocean model with observations. This is a huge engineering challenge but one that we feel can be accomplished by the team that we have assembled and the resources we are requesting. By harnessing the computational resources being put in place by Project Columbia and other NASA initiatives, NASA data streams, in-situ observations and the state-estimation methods developed in ECCO, we are in a position to transform the fidelity of ocean models. Without data assimilation our models drift away from the observed state: without models the data yield only a partial description of the ocean state and often do not allow us to compute what we are really interested in. The estimates will be used for understanding dynamics, for calculating and monitoring climate-related quantities such as meridional fluxes, air-sea property fluxes etc, for initializing ENSO predictions and for study of the role of the ocean in climate.

The objective is to create a synthesis that combines, in an optimal way, most accumulated knowledge about ocean circulation and sea-ice, both dynamical and observational, and that is of sufficient quality so that for most science applications it can be used in lieu of raw data. Preliminary work in preparation for this proposal, reviewed in Section 3, has given a tantalizing glimpse of what the future holds and helped to forge strong links between key players in this consortium. As discussed in Section 4, methods will be used that will enable ocean and sea-ice models to be constrained by all relevant data streams, irrespective of their type, at very high resolution.

2 Previous work: where are we now?

There are several existing data assimilation efforts at eddy-resolving resolutions (for a summary see US GODAE, 2001). For example, the Naval Research Laboratory has a real-time global ocean analysis system (http://www7320.nrlssc.navy.mil/global_nlom/); the HYCOM Consortium (http://hycom.rsmas.miami.edu/) has developed near basin-scale analysis systems for the Atlantic and Pacific; the French MERCATOR project maintains a North Atlantic operational analysis system (http://www.mercator-ocean.fr/). But the focus of these systems is operational ocean nowcast and forecast. Therefore the assimilation methods are mostly based on optimal interpolation (OI) or equivalents and the control problem is restricted to the initial state with no time-dependent components or model parameters. While computationally efficient, these methods introduce temporal discontinuities, rendering the assimilation products physically inconsistent, and do not fully exploit the available information. Moreover, because of operational requirements, the input data streams are limited. Such shortcomings limit the usefulness of these products for climate monitoring and research. The proposed investigation, a continuation of the ECCO project, distinguishes itself from these operational efforts in that it is primarily oriented towards climate research, e.g., to address objectives of the Climate Variability and Predictability (CLIVAR) programme.

The existing ECCO capability was originally constructed to demonstrate the practicality and utility of state estimation for global-scale physical oceanography. At the time of its conception, the oceanographic community was beginning to obtain the first truly global synoptic observations from satellites (principally the TOPEX/POSEIDON and now Jason altimetry, NSCAT/QuikSCAT wind stresses, TRMM/TMI sea surface temperatures, and the ARGO network), as well as those from the global in-situ measurements associated with the World Ocean Circulation Experiment (WOCE). The breadth and depth of the new data bases was clearly only scientifically exploitable in full if these data could be combined, i.e., dynamically linked with each other through an adequate ocean general circulation model (GCM). To this end, the ECCO consortium was formed, involving scientists from the Massachusetts Institute of Technology (MIT), the Scripps Institution of Oceanography (SIO), and the Jet Propulsion Laboratory (JPL). The computational kernel of ECCO is the MITgcm (Marshall et al., 1997a,b) which has been constructed *ab initio* with its use in state estimation a major consideration. The MITgcm incorporates a sophisticated mixed-layer scheme (Large et al., 1994), eddy parameterizations (Gent and McWilliams, 1990), as well as a modified version of the ice model of Zhang and Hibler (1997). Data sets included (eventually) all of the remote and in-situ data collected during WOCE, with NASA data streams being of central importance.

2.1 Combining models with observations

With all data and the model written in discrete form for use on a computer, the state estimation problem reduces to one of ordinary least-squares — see monograph by Wunsch (1996) for a formal and detailed definition of the ocean state estimation problem. Briefly, state estimation applied to large systems seeks the minimum of a cost function, J, of form

$$J = \mathbf{u}^T \mathbf{Q}^{-1}\mathbf{u} + \mathbf{n}^T \mathbf{R}^{-1}\mathbf{n}, \quad (1)$$

where vector $\mathbf{u}$ represents a set of uncertain parameters that can be used as "controls", i.e., adjusted to bring a model simulation closer to observations. For example, vector $\mathbf{u}$ includes terms that represent errors in the initial conditions, in the time-evolving surface forcing fields, and in the model's empirical mixing coefficients. Vector $\mathbf{n}$ represents data errors of all kinds, including model errors that are not already represented by $\mathbf{u}$. $\mathbf{Q}^{-1}$ and $\mathbf{R}^{-1}$ are weight matrices, and superscript T is the transpose operator. Often $\mathbf{u}$ and $\mathbf{n}$ are treated as stochastic variables with zero mean and covariance matrices $\mathbf{Q}$ and $\mathbf{R}$, respectively.

A number of ways are known for the determination of solutions to the above least-squares problem (e.g., Menke, 1989; Wunsch, 1996, and references therein). The ECCO problem is noteworthy primarily for its enormous dimension, rendering conventional or brute force methods impractical. For that reason, the ECCO consortium chose to pursue two (related) methods. The first method is a sequential one that employs a Kalman filter algorithm in the forward time direction and the RTS smoother in the reverse time direction. Because the problem size is so great, the normal approach of time stepping the state vector using a model linearized about the previous estimated state (the extended Kalman filter/smoother) is impractical to implement in full form. Instead, a reduced-state filter/smoother is implemented (Fukumori, 2002). This method, whose development has been focused at JPL, has proved useful in practice and has the virtue of producing approximate error estimates for the state.

The second approach is based upon the least squares method of Lagrange multipliers (also known as the "adjoint method", "4DVAR" and other labels). R. Giering, first as a student at the Max Planck Institute (MPI) in Hamburg and then as a post-doc at MIT, had developed an automatic differentiation (AD) tool capable of (semi-) automatically generating the so-called adjoint model to the MITgcm (see Giering and Kaminski, 1998). The AD tool is now in a commercial version known as TAF (Giering and Kaminski, 2003). The adjoint model serves several purposes, of which the most central is that it produces the partial derivatives of cost function J with respect to control parameters $\mathbf{u}$, which permits a relatively rapid optimization. This methodology has proved practical for the immense-sized ECCO problem and has been documented in a number of places (e.g., Marotzke et al., 1999). A related methodology, which has also been successfully applied to the ECCO problem is the Green function approach — see Menemenlis et al. (2004a). Green functions provide an efficient way to compute the partial derivatives of cost function J when the number of control variables that are included in vector $\mathbf{u}$ is small.

Both adjoint and sequential methods adopted by ECCO are computationally intensive relative to other approximate methods, often based upon nudging, a form of objective mapping, or restricted, e.g., to a Kalman filter step alone. The ECCO goal has been, and remains, to obtain a state estimate satisfying the model time-evolution equation, and in which the error field has been propagated through the same physical model as the state vector. As mentioned above, a system restricted to the filter step alone undergoes non-physical jumps each time the update of the model forecast by the data is used (see, e.g., Fukumori, 2002; Wunsch, 1996). In weather forecasting, where prediction starting from the analysis time is the central goal, the physical jump prior to the prediction time is of no particular concern. The oceanographic focus on analyzing the time-evolving circulation and its properties drives our emphasis on achieving a dynamically consistent evolution. In ECCO, both sequential and adjoint/Green-function methods proved useful. The current JPL implementation of the sequential approach has been focused on adiabatic motions. Because our present focus is on longer-term climate scales, we propose to exploit the adjoint/Green-function approach, for which much of the needed development has already taken place, as described in Section 4.2.

2.2 ECCO results and status

Over 60 publications have originated from ECCO (see http://www.ecco-group.org/publications.html) and we present only a brief summary here. The so-called WOCE synthesis of ECCO was originally carried out at 2° horizontal grid spacing for the 6-year period, 1992-1997, and has evolved to a 12-year period at 1° grid spacing, as described by Stammer et al. (2002, 2003, 2004). Although the results were preliminary (the resolution is coarse, and not all of the model elements were employed), this computation served the primary ECCO purpose of demonstrating the feasibility of the system for climate-scale global problems. Despite its shortcomings, the results of this demonstration solution have proved scientifically useful in the study of a large number of oceanographic and interdisciplinary studies:

1. ocean circulation, e.g., Stammer et al. (2003), Gebbie et al. (2004),
2. ocean biogeochemical cycle and air-sea carbon fluxes, e.g., McKinley et al. (2004),
3. air-sea fluxes, e.g., Stammer et al. (2004),
4. subgridscale parameterization of unresolved processes, e.g., Ferreira et al. (2004), and
5. earth polar motion and rotation studies, e.g., Gross et al. (2004).

As one example, Fig.1 shows a schematic of the data employed in the ongoing ECCO coarse-resolution (1°) optimization. The solution was used to estimate air/sea fluxes of momentum, heat, and fresh water. A large number of comparisons were made between ECCO results and estimates obtained independently. For example, Fig. 2 shows how the optimized ECCO model adjusted the NCEP-NCAR reanalyzed wind stresses and brought them in to closer agreement with NSCAT wind-stress fields.

With support from GODAE the coarse-resolution ECCO state estimates have become quasi-operational and are made available on a public data server. That is, we are committed to maintaining the analysis capability, updating it at convenient intervals, and adding any new data that appear to carry new information. However these "ECCO-GODAE" estimates remain at coarse resolution and do not capture the most energetic part of the spectrum of ocean variability. Furthermore, polar processes, so important to climate, are also not represented. Three major objectives of the present proposal are (i) to handle ocean eddies by resolving them in a global context, (ii) to include polar processes using sea-ice models and novel gridding techniques, and (iii) to exploit new observations, including a variety of hitherto non-assimilated sea-ice data. We now briefly discuss the rationale for our emphasis on eddies and sea-ice.

2.3 The need for eddy-resolving models

Much has been achieved. But, as mentioned in the introduction, the coarse horizontal grid of the ECCO estimates — now 1° — is a severe limitation on their ability to describe the real ocean. Despite the very great progress made in the last 15 years toward parameterizing eddy fluxes in ocean models (see, e.g., Gent and McWilliams, 1990; Visbeck et al., 1997), some problems seem intractable through this route. (i) The eddy-parameterizations are not based upon completely fundamental principles, and fail to adequately account for known anisotropies in their fluxes. Everything we know about eddies suggests that their property fluxes can accumulate in the ocean, changing it measurably and importantly from what it would be if eddies were absent. That is, the long-wavelength, low-frequency features characterizing climate are controlled in part by eddy fluxes. In

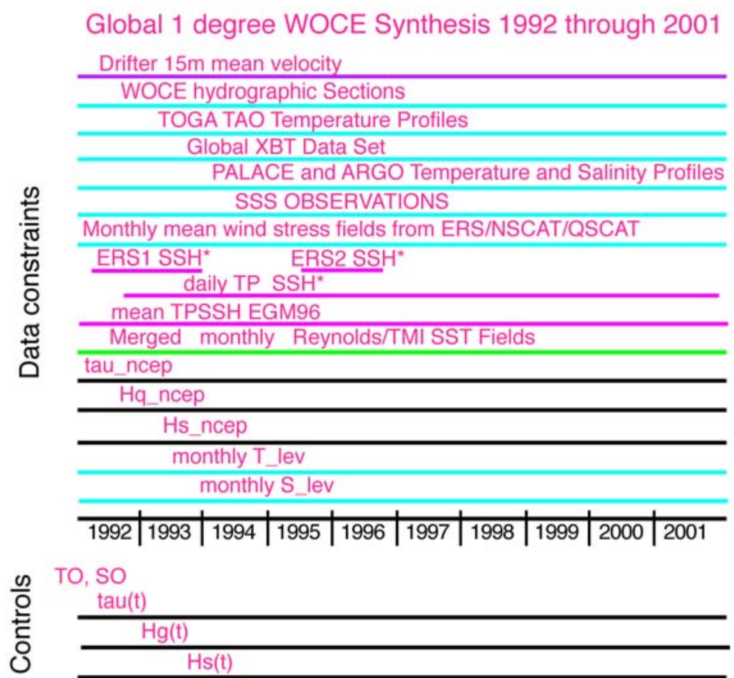

Figure 1: Data used in the ECCO estimates from which air/sea fluxes were extracted by Stammer et al. (2004). "Controls" represent the adjustable parameters of the least squares fit and include initial and surface boundary conditions.

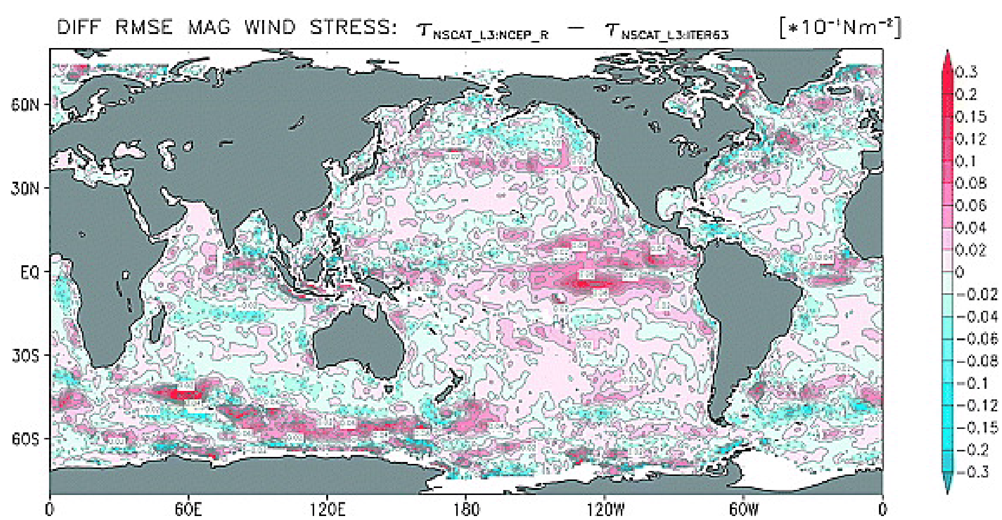

Figure 2: Reduction in the RMS difference between ECCO and NSCAT wind stress time series as compared to the original NCEP-NSCAT differences. Positive values (colored light-pink and red) indicate a respective misfit reduction (in N/m^2). From Stammer et al. (2004). Overall, this is a net reduction in misfit variance.

addition, studies (e.g., Jones and Marshall, 1997) have shown that horizontal grid-spacing of order 2–4 km is required to resolve restratification processes in deep convection, in which stratified fluid in the periphery of the convection patch is drawn over the surface allowing the convected fluid to be 'swallowed' by the ocean. If restratification is not represented adequately, then the water-mass properties of the modeled ocean deteriorate over time, there are model drifts, and the attendant air-sea fluxes become compromised and have to be 'corrected'. This process — restratification of mixed layers — is a ubiquitous feature of the ocean, but is particularly important in strong frontal regions such as the Antarctic Circumpolar Current of the Southern Ocean, the Kuroshio and the Gulf Stream. (ii) In climate, scalar property transports are of central interest (heat, fresh water, carbon, oxygen, etc.) and in the ocean, narrow western and eastern boundary currents make major contributions; these boundary currents are not parameterizable and, until they are resolved, there will always be doubts that the ocean model is carrying its property transports realistically. (iii) Ultimately, water mass properties in the ocean are important to climate and climate change. In the abyssal ocean, the inability to resolve major topographic features (e.g., fracture zones, sills) leads to systematic errors in the movement of deep water masses with consequences for the accuracy of computation of carbon uptake, for example.

The underlying theme of the present proposal is the description, understanding and monitoring of ocean circulation in the presence of fully-resolved eddies and boundary currents. The intention is that the large, climate-relevant scales should be as accurate as possible in response both to large-scale forcing and to the presence of eddies. Although it has been possible to run global-scale forward models with nearly or fully resolved eddy fields and boundary currents for a decade or more (e.g., Semtner and Chervin, 1992), the possibility of rigorous ocean state estimation at this resolution has hitherto been unthinkable. Therefore there is almost no direct experience with state estimation in the presence of active eddies except for comparatively small-area regional models directed at mesoscale variability. Our most immediate experience is from the recently completed PhD thesis of J. Gebbie (2004); Gebbie's results and their implication for the work proposed herein are discussed in Section 3.5. With the computational resources now in place (see Section 3.3) and with the knowledge and experience built-up within the ECCO group, the time seems ripe to launch the ambitious program of rigorous high-resolution state estimation, described herein, that will yield estimates of the evolving state of the ocean at resolutions approaching the limits believed required.

2.4 The need for sea-ice

Sea ice, though only a thin layer between the air and the sea, has strong and numerous influences within the climate system. It affects radiation balance due to its high albedo, surface heat and mass fluxes due to its insulating properties, ocean convection by preconditioning sites of deep water formation, freshwater fluxes during ablation, ice margin processes related to small-scale ocean phenomena, and, of course, human operations (Carsey et al., 1992). In both the north and the south, the seasonal variation in the coverage and thickness of sea ice is appreciable and in addition there are interannual variations and apparent trends in the data (Parkinson et al., 1999). Sea-ice thickness varies widely and depends on ice age, thermal and radiative processes, mechanical deformation, snow cover and ocean mixed-layer processes (e.g., Zhang and Rothrock, 2000).

Sea-ice processes also impact high-latitude oceanic uptake and storage of anthropogenic CO_2 and of other greenhouse gases both directly, by affecting the exchange of these gases across the air-sea interface, and indirectly, by influencing the biological, chemical, and physical processes that transport greenhouse gases from the surface ocean into the interior. A number of modeling studies have shown that a consideration of realistic sea-ice processes often leads to significant improvement in the representation of high-latitude uptake of these transient tracers (Caldeira and Duffy, 2000;

Duffy et al., 2001; Dutay et al., 2002). Furthermore, sea-ice albedo provides a positive feedback mechanism for climate warming or cooling, which enhances the high-latitude response to climate change (Manabe et al., 1991) and which may play a role as an abrupt, climate-switch mechanism (Gildor and Tziperman, 2001).

For practical and computational reasons, existing ECCO analyses exclude the Arctic Ocean and lack an interactive sea-ice model, restricting the use of satellite data over ice-covered regions. As is discussed in Section 3.3, an efficient global-ocean model configuration coupled to a dynamic/-thermodynamic sea-ice model is now in hand, and this coupled ocean and sea-ice model will be an integral component of the next-generation, ECCO-II analyses that are the focus of this proposal. 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.

Improved representation of high-latitude processes in the ECCO estimates will enhance hindcasting and forecasting capability, both of which are needed for climate-change studies, for the design of in-situ measurement campaigns, and for operational purposes, e.g., navigation, drilling activity, wildlife behavior, and dispersion of pollutants. Improved representation of high-latitude processes is also a crucial requirement for reducing uncertainty in predictions of the climate response to prescribed increases in greenhouse gases and in predictions of how the oceanic sink for greenhouse gases might change in the future in response to a changing climate.

3 Prospects for high-resolution state estimation

The key technical challenge of this proposal is to develop rigorous global-ocean and sea-ice state estimation methodologies that can be applied for a decade or more, at eddy-resolving resolution, and for the full ocean depth. This is a huge engineering challenge. Depending on the method and on the approximations that are used, the computational cost of state estimation can be several dozen to several thousand times more expensive than integrating an ocean model without state estimation. Computational cost is further compounded by the chaotic nature of an eddy-resolving model, which complicates the derivation of tangent-linear gradients and their application to the minimization of cost function J, as is done in the existing ECCO solutions. We argue here that a number of recent advances bring rigorous eddy-resolving, decadal-time-scale estimates of the global-ocean and sea-ice circulations within reach: (i) the availability of an increasing number of global, eddy-resolving, satellite data sets, (ii) the establishment of the required computational framework, (iii) the configuration of an efficient mesoscale-eddy permitting model that achieves a throughput approaching ten years of model integration per day of computation, (iv) the demonstration that initial conditions and surface forcing fields estimated at coarse resolution can improve the solution of an eddy-permitting model, and (v) the recent successful ECCO experience with the application of the adjoint method to a regional, eddy-resolving model configuration.

3.1 NASA's eddy-resolving ocean and sea-ice data

NASA's satellite altimetry programs have revolutionized the field of physical oceanography. The decade-long record of the global ocean surface topography produced by the TOPEX/POSEIDON (T/P) mission, now continued by the Jason mission, has led to new understanding of large-scale ocean dynamics (Fu and Cazenave, 2001). As reviewed in Section 2, numerical models of the global ocean constrained by altimetric observations can, for the first time, yield estimates of the evolving physical state of the 3-d ocean circulation. Scientists are now using such analyzed data products

Observations	Source	Spatial Coverage and Sampling	Temporal Coverage and Sampling	Measurement Uncertainty
Sea-Ice Concentration	SMMR SSM/I AMSR-E	Arctic Ocean Southern Ocean	1978-present 2-day	∼4-7%
Sea-Ice Motion	SMMR SSM/I AMSR-E AVHRR	Arctic Ocean Southern Ocean ∼100 km	1978-present 2-day	∼4–9 cm/s
Sea-Ice Temperature and Albedo	AVHRR MODIS	Arctic Ocean Southern Ocean	1978-present monthly	∼2 K
	TOVS	Arctic Ocean ∼100 km	1979-present daily	

Table 1: Sea-ice satellite retrievals to be assimilated.

for studying problems ranging from circulation physics to biogeochemical processes, and temporal changes in the earth's gravity field (see Section 2.2). This new capability — of using numerical models with satellite observations for optimal estimation of ocean state — is an invaluable legacy of NASA's investment in T/P and Jason satellite missions, in scatterometer satellites, and in the ECCO project.

Satellite technology continues to develop, and the ongoing studies of global change are producing demands for ever more data of higher quality. For example, despite its frequent global coverage, T/P or Jason, or even the combination of the two running simultaneously, cannot temporally resolve the mesoscale eddies, which are the most energetic component of ocean circulation (Fig.3). Realizing the importance of ocean eddies, NASA/JPL has developed a new instrument called Wide-Swath Ocean Altimeter (WSOA), based on the technique of radar interferometry (Fu, 2003). WSOA is planned to be launched on the Jason follow-on mission, called Ocean Surface Topography Mission (OSTM). The sampling capability of this new instrument is shown in Fig.4b. The measurement will be made at 15 km resolution with rms accuracy of 4-5 cm. Its sampling capability is better than five coordinated conventional altimetric satellites, allowing the synoptic evolution of the mesoscale eddies to be monitored from space for the first time. We realize that the decision on the flight of WSOA is pending. With the flight of WSOA, our goal is to be able to use the resulting data in ways that will extract all of the new information to achieve the best estimate of oceanic eddy field. Without WSOA, our goal is to use all available altimeter data from multiple satellites. The results will then be useful for future development of similar high-resolution instruments.

Other NASA instruments and data products such as the new generation of scatterometers, high-resolution sea surface temperature (SST), and sea-ice retrievals, greatly increase our ability to resolve mesoscale and high-latitude processes. The QuikSCAT scatterometer, for example, increases the accuracy, precision, and coverage of the wind field forcing of the ocean, with all of its consequences for other air/sea exchanges (heat, freshwater, carbon, etc.). Fig.5 shows the wind field over the western North Atlantic measured by QuikSCAT (Courtesy of D. Chelton), showing many intriguing mesoscale features apparently coupled to ocean currents. To fully understand the ocean circulation, its variability, and its implications for climate change, we will need to understand the mesoscale — both in terms of its coupling to larger oceanic scales and its influence on the atmosphere.

Another example is the tremendous wealth of satellite sea-ice retrievals, with NASA often playing a key role in their analysis. Table 1 is a preliminary list of the sea-ice data products that we plan to assimilate. These include retrievals from the Scanning Multichannel Microwave Radiome-

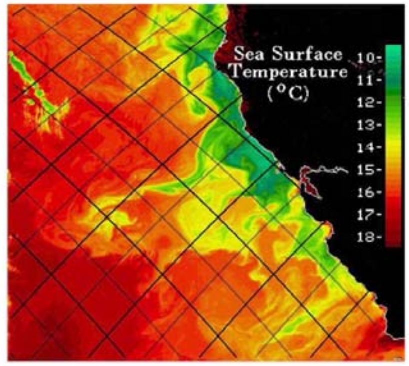

Figure 3: Ground tracks of T/P (thick lines) and Jason (thin lines) superimposed on an image of sea surface temperature (SST) measured by AVHRR off the coast of California. The details of the mesoscale features of SST are not properly resolved by the Tandem Mission of TP and Jason.

ter (SMMR), the Special Sensor Microwave Imager (SSM/I), the Advanced Microwave Scanning Radiometer for EOS (AMSR-E), the Advanced Very High Resolution Radiometer (AVHRR), the Moderate Resolution Imaging Spectroradiometer (MODIS), and the TIROS Operational Vertical Sounder (TOVS):

Sea-ice concentration. Sea ice concentration has been computed routinely for the series of passive microwave radiometers starting with SMMR in 1978 and extending through SSM/I and now AMSR. Recent work by team member R. Kwok (2002), suggesting significant overestimation of open-water coverage, will be utilized in the formulation of error estimates. While the errors indicated by Kwok are problematic, it is clear that the archived passive microwave data set will contribute to our project due to their long duration, and we are optimistic that improvements in the data products will be forthcoming.

Sea-ice motion. Multiple daily sequential satellite images and a maximum cross-correlation algorithms are used to compute sea ice motion for both polar regions. The primary source of ice motion is passive microwave imagery from SMMR and SSM/I. Together these two satellite datasets constitute a series extending from Oct. 1978 to the present. The SMMR imagery produces ice motion vectors with a spatial resolution of about 25 km for the 37-GHz channel starting in 1978. This overlaps with SSM/I, which starts in the fall of 1987. In addition to a 37-GHz channel, SSM/I has an 85-GHz channel with 12.5-km spatial resolution resulting in improved ice motion mapping. Starting in 2002 ice motion is available from AMSR-E data, which adds improved spatial resolution (6 km) from a new 89-GHz channel. For the Arctic, 1982-1999 ice motion vectors are also available from sequential visible, near-infrared, and thermal infrared imagery from the Advanced Very High Resolution Radiometer (AVHRR), with higher spatial resolution (1 km) but limited coverage ($\sim$25%) because of clouds. Thus, there exist ice motion data products starting in Oct. 1978 up through the present with step-wise improvements in spatial resolution as data from new and better satellite radiometers is incorporated.

Sea-ice temperature and albedo. Sea-ice surface skin temperature and summer albedo have been obtained from AVHRR from MODIS and from TOVS as part of the Polar Pathfinder project. These retrievals have been extensively validated with measurements from field programs (CEAREX, LeadEx, SHEBA) and Russian meteorological stations drifting on sea ice.

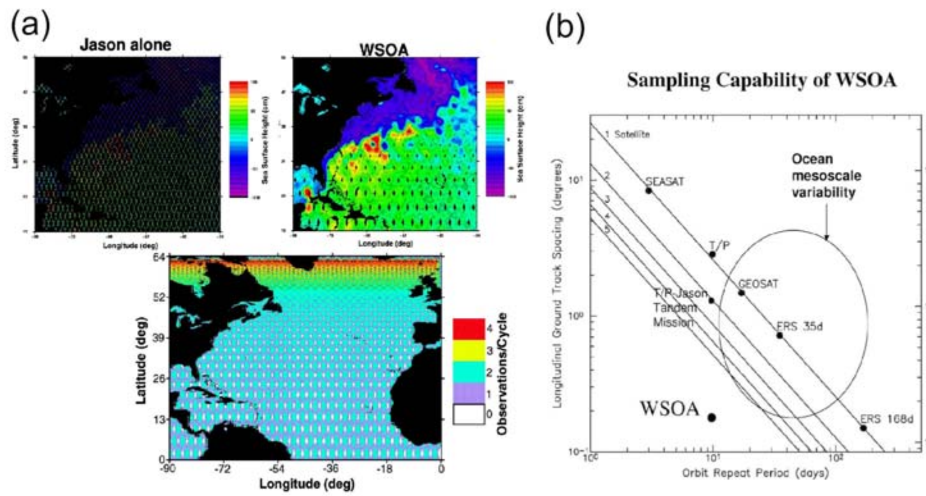

Figure 4: (a) Top right panel: Sea surface heights (SSH) sampled by a single nadir altimeter like Jason. Top left panel: SSH sampled by WSOA. Lower panel: number of observations made by WSOA over a 10-day repeat cycle. At mid and high latitudes, more than one observation are made during a 10-day period. (b) Sampling capability of satellite altimetry as a function of repeat period and ground track spacing. The various straight lines represent the trade-off between repeat period and ground track spacing for optimal satellite configuration consisting of 1-5 conventional nadir altimeters as labeled. The case for various missions such as Seasat, etc., are indicated along the line of a single altimeter. The space and time scales of oceanic mesoscale are represented by the ellipse. The sampling capability of WSOA is better than five nadir altimeters.

3.2 Global eddy-resolving and sea-ice modeling on the cubed sphere

A necessary condition for global, eddy-resolving state estimation is the availability of an efficient model and of significant computational resources. To address the computational challenge of resolving ocean circulation and sea-ice over the globe, new, quasi-isotropic gridding and semi-structured parallel domain decomposition techniques are being employed, which, together with the significant resources made available by Project Columbia (see Section 3.3), make the proposed work possible. We now briefly describe the existing, high-resolution cubed-sphere grid, which is the baseline configuration for the work proposed herein.

High-resolution modeling on the surface of the sphere presents a major challenge but recent advances in numerical methods now offer new avenues for efficient global modeling, which are being exploited by the ECCO group. Historically, ocean climate models were discretized in the horizontal using grids based on geographic coordinates (latitude-longitude) grids. Such grids exhibit a convergence of meridians; that is, in the vicinity of the poles, the grid-spacing in the zonal direction is very small compared to zonal grid-spacing along the Equator. Without special measures, this shortest grid length dictates the longest stable time-step that the model can take and severely limits the utility of the model. Early strategies to overcome this problem involved "polar filters", which stabilize the model near the poles, effectively throwing away the enhanced polar resolution.

Recently, the MITgcm was implemented on an entirely different kind of spherical grid known as the expanded spherical cube. It has much more uniform coverage of the sphere than a latitude-longitude grid and thus does not need any stabilization or grid distortion to be useful for ocean (and indeed atmospheric) modeling. The baseline, global, eddy-permitting model configuration, which

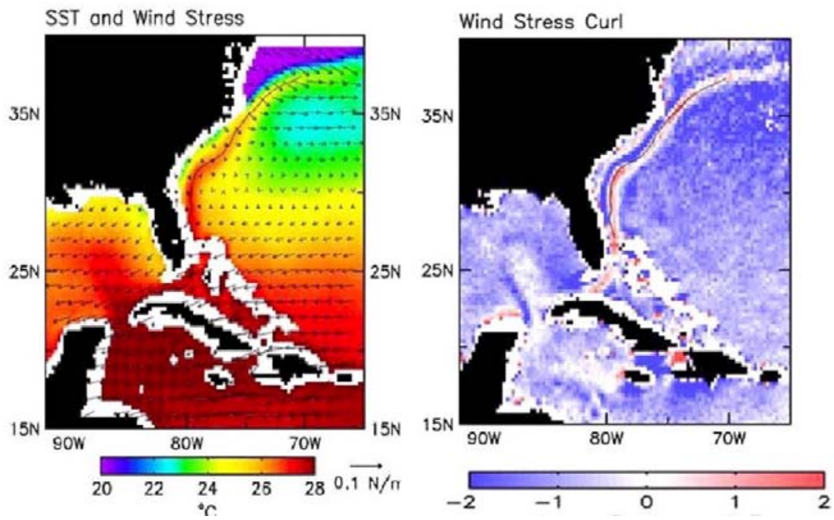

Figure 5: Left: SST (color coded) measured by the microwave radiometer onboard the TRMM satellite and wind stress (vectors) measured by the QuikSCAT scatteromemter. Right: the curl of the wind stress. Note the details of the wind stress curl affected by the Gulf Stream and the Loop Current. (Courtesy of Dudley Chelton of Oregon State University)

will be used for the work proposed herein, is depicted in Fig.6; it uses a conformally expanded spherical cubic grid following Adcroft et al. (2004). The grid is generated algebraically (Purser and Rancic, 1998) and it preserves local orthogonality for efficient and accurate time stepping of the model equations. The cubed-sphere grid is semi-structured in the following sense. When projected on the sphere, each face of the cube can be topologically warped onto a rectangular computational domain and thus a structured grid representation can be used. But unfolding the cube into a plane, as is done in Fig.6, reveals that the pattern of connectivity between the faces is not naturally described by two indices (i, j). Instead, the faces are best described as an unstructured net with arbitrary connectivity at the edges of each cube face, a feature that we propose to further exploit in the generation of yet newer grids, discussed in Section 4.1.2, that will be even more efficient for global high-resolution modeling.

The baseline configuration of Fig.6 has mean horizontal grid spacing of 18 km. That is, each face of the cube is comprised of 510x510 grid points. The configuration has 50 vertical levels, down to 6 km depth, with enhanced near-surface resolution: level thicknesses are 10 m in the top 100 m and gradually increase to 450-m thickness at the bottom. A novel aspect of the vertical grid is the use of shaved cells (Adcroft et al., 1997), which permits a much-improved representation of bottom bathymetry. The model configuration employs the Large et al. (1994) KPP vertical mixing scheme. Horizontal viscosity follows Leith (1968). The ocean model is coupled to an interactive sea-ice model. The thermodynamic component 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, 1997) with an efficient parallel implementation of the line-successive-relaxation solver on the cubed-sphere grid. More information about this particular integration, and others like it, is available at http://ecco.jpl.nasa.gov/cube_sphere/. What matters for the present discussion is not the exact details of mixing parameterizations, forcing methodology, etc., since these will be superseded with forcing fields and parameters estimated using data assimilation, but rather the efficient numerical implementation of this model configuration, discussed next.

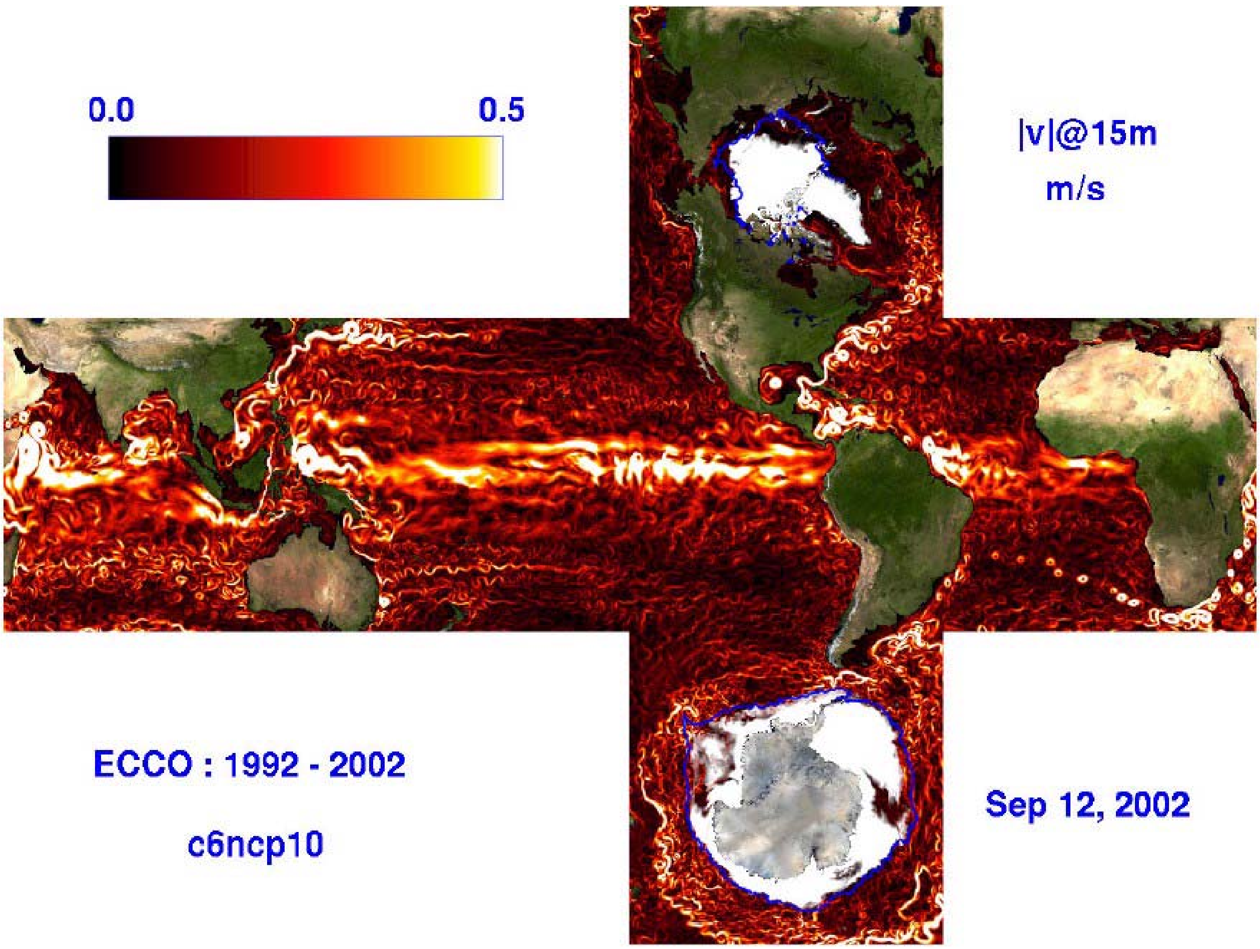

Figure 6: Baseline, eddy-permitting, cubed-sphere grid configuration. The figure shows simulated near-surface (15-m) ocean-current speed and sea-ice cover from a preliminary eddy-permitting integration. Units are m/s. Simulated sea-ice is shown as an opaque, white cover. Land masses and ice shelfs are overlain with NASA satellite imagery. The thin blue line is passive radiometer observations of sea-ice extent (15% concentration). The difference between observed and simulated sea-ice extent, e.g, excessive summer melting in the Arctic and unrealistic open-water winter polynyas in the Ross and Weddell Seas, is one of the signals that we propose to assimilate in order to improve the model representation of high latitude processes. Animation of this figure for the complete 1992-2002 period and more information about this integration, are available at http://ecco.jpl.nasa.gov/cube_sphere/.

3.3 Project Columbia: computational challenges and opportunities

As part of Project Columbia, the NASA Advanced Supercomputing (NAS) group at the Ames Research Center (ARC) has installed a 10240-processor SGI Altix that has a combined peak capacity of 61 teraFLOPS, 50% more capacity than Japan's Earth Simulator. The calculations proposed herein stretch the limits of the available computer technology, even for a system of the size of Project Columbia at NAS-ARC. As has been the case during the first five years of ECCO, a large fraction of the consortium's effort has been, and will continue to be, to scale down the ocean state estimation problem in order to fit within and to take advantage of latest available computer technology. Here, and also in Sections 4.1.3 and 5.2, we describe the computational strategy, which addresses the challenges that the proposed work presents.

Carrying out coupled ocean/sea-ice global data assimilation at eddy-resolving resolutions over the full ocean depth for decades or more is computationally very demanding. It requires iterating over decade-long prognostic simulations containing $\sim 10^8$ grid cells. The calculation foot-print of a single decade of simulation is $\sim 10^{17}$ arithmetic operations. On a modern desktop computer this computation would take several years to complete. But this can be cut to a single day using just 512 Project-Columbia processors.

In preparation for this proposal, ECCO scientists have been working closely with Project Columbia team members during this past year. As noted in the "Networking and Information Technology Research and Development (NITRD) FY 2004 Interagency Coordination Report" the ECCO-NAS collaboration has shown that ECCO-II-type production computations are viable on Project Columbia. The results of Fig.6 were obtained by mapping the cubed sphere on 500 of Project Columbia's 10240 processors. The shared memory architecture of the SGI Altix, the supportive computational resource culture at NAS, and the advanced numerics and parallelization capabilities of the MITgcm have made possible a throughput approaching ten years of model integration per day of computation for this configuration. At this level of throughput, eddy-permitting estimates of the global ocean and sea-ice circulations are now possible. Assuming that a resource equivalent to a 512-processor SGI Altix system continues to be available, the ECCO-II team will be well placed to deliver a preliminary high-resolution estimation product in the early phases of this proposal (see Sect. 4.2.1). This will enable ECCO II to quickly make results broadly available and to accommodate critical feedback from the analysis and application of those results into subsequent production cycles.

3.4 Multi-scale estimation approaches

A first question that has been addressed, using Project Columbia computer resources, is whether the existing, low-resolution ECCO estimates of initial and surface boundary conditions can be used to initialize eddy-resolving estimation efforts. For this purpose two 1992-2002 integrations were conducted using a near-global configuration with 1/4-degree horizontal grid spacing. 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 synthesis (Kistler et al., 2001). Initial conditions and surface fluxes for the second integration are from the ECCO 1°, 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. The NCEP-forced integration requires root-mean-square (rms) temperature relaxation fluxes on the order of 30 W/m^2 while rms temperature relaxation fluxes for the ECCO-forced integration are substantially less, order 10 W/m^2. 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

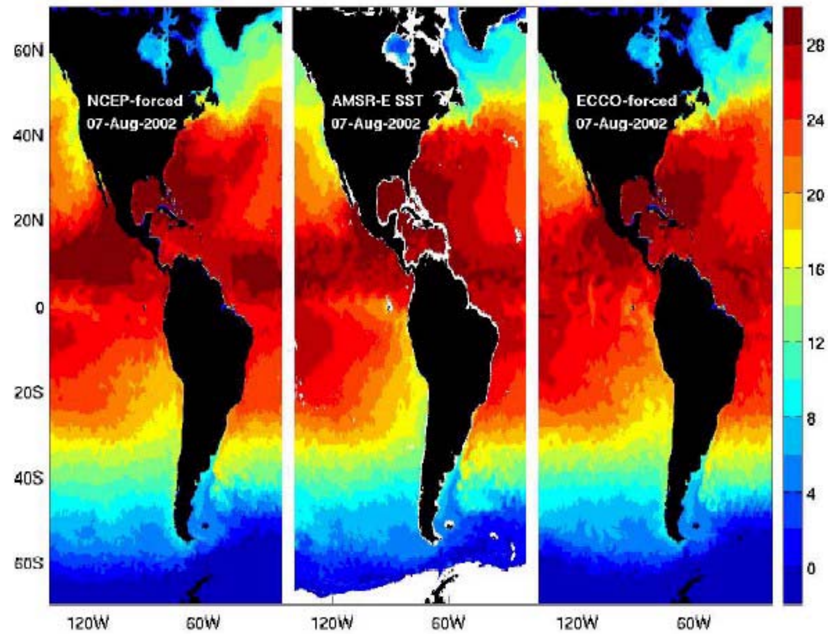

Figure 7: AMSR-E observations of sea-surface temperature (middle) are compared to two, eddy-permitting integrations forced, respectively, by the NCEP meteorological reanalysis (left) and by surface fluxes estimated by the ECCO, 1-degree, adjoint-method optimization (right). Units are °C. The ECCO-forced integration is closer to the observations, both in a global, root-mean-square sense, as well as in the large-scale patterns of the circulation, for example, the separation of the Gulf Stream and upwelling off the West Coast of South America. The figure also displays intricate spatial patterns of oceanic mesoscale activity, both in the observations as well as in the model simulations. This comparison is best viewed as an animation, available at http://ecco.jpl.nasa.gov/cube_sphere/.

eddies and of other physical processes.

The NCEP and the ECCO eddy-permitting integrations were compared to the complete suite of observations used in the coarse-resolution ECCO optimizations (Menemenlis et al., 2004b). 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 variability of upper ocean temperature and salinity. The ECCO-estimated forcing also improves the strength of the Equatorial Undercurrent, the path of the Kuroshio, and the path of the Gulf Stream. For example, Fig. 7 compares sea-surface temperature (SST) observed by the AMSR-E instrument to that of the NCEP and the ECCO-forced model integrations. The root-mean-square difference between AMSR-E data and the NCEP integration is 1.6°C, compared to 1.2°C between AMSR-E and the ECCO integration. The SST data also serves as a tracer, which shows that the ECCO-forced circulation patterns, e.g., the separation of the Gulf Stream and upwelling off the West Coast of South America, are more realistic than those of the NCEP-forced integration. The above results indicate that boundary conditions estimated at coarse resolution can improve the solution of high-resolution simulations; they motivate and justify the multi-scale estimation approaches, i.e., Green functions and approximate adjoint, which are proposed in Sect. 4.2.

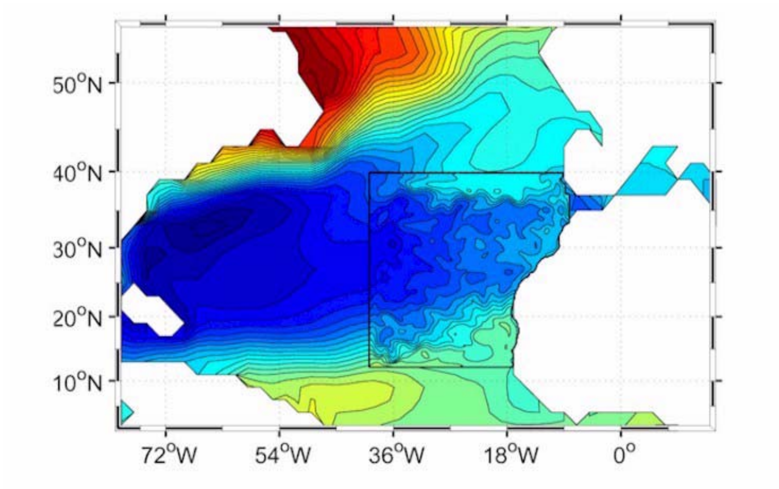

Figure 8: A data-assimilative $\frac{1}{6}^\circ$ model of the subduction region in the North Atlantic embedded in the ECCO 1° global model — from Gebbie et al (2004).

3.5 Adjoint-method optimization in the presence of eddies

Fig.8 displays results from a study by Gebbie et al (2004), whereby the MITgcm is integrated at 1/6° horizontal grid spacing — permitting eddies to develop — in the subduction region of the North Atlantic. The model is embedded within the coarse-resolution global ECCO solution. Gebbie et al (2004) have shown that in this realistic, full eddy-resolving model of the eastern North Atlantic, no particular problems were encountered in practice for assimilation spanning the Subduction Experiment period (2 years). This straightforward behavior is attributed to the quiescence of the region (no boundary currents or strong convection sites), as well as to the availability of adequate sets of observations (provided by the Subduction Experiment) to constrain otherwise unstable model components. The study also found that assimilating statistical properties of the flow field rather than transient features improved the performance of the estimation procedure (a gradient-based iterative optimization). This confirms assertions by Lea et al (2000, 2002) and Köhl and Willebrand (2002; 2003), KW hereafter, that estimation in the presence of strong eddies can be problematic. They show examples of fluid systems in which the local derivatives required by the adjoint method can fail to exist or become inaccurate. Nevertheless, KW demonstrate that focusing the estimation onto statistical characteristics enables the estimation window to be extended far beyond the predictability limit of the transient flow. Besides careful formulation of the objective function in terms of statistical moments, as already implemented by Gebbie et al (2004), KW use a representation of the statistical model in terms of time and space averages of the full forward model to successfully assimilate, e.g., SSH variability into an eddy-permitting model of the North Atlantic. Their idea to approximate the adjoint statistics by a coarse reslution adjoint of the forward model and to force the linearization to (i.e., evaluate the adjoint at) the statistics of the high-resolution forward model through its time-mean can be readily implemented in the present context of exact tangent linear and adjoint code generation, as will be shown in Section 4.2.2.

4 Proposed work

We propose to (i) map very high resolution forward models of the ocean circulation and sea-ice on to Project Columbia at NAS-ARC using a combination of novel gridding techniques and parallel computing technologies, (ii) constrain these models with global data sets using rigorous state estimation methods in order to obtain estimates of the global-ocean and sea-ice circulations, suitable for climate research, at unprecedented resolutions, (iii) evaluate and make use of the estimation products and tools for various science applications, and (iv) make the state estimation products as well as the modeling and estimation tools available to the community.

4.1 High-resolution simulations

The objective of the high-resolution simulations is threefold: (i) to carry out ultra-high resolution simulations, order 1-km horizontal grid spacing, (ii) to configure a next-generation model for global-ocean and sea-ice data assimilation, and (iii) to improve code efficiency and scaling. Before discussing these technical issues, we outline the science questions.

4.1.1 Eddies and ocean circulation (Marshall)

Eddy-mixed-layer interaction As emphasized in Section 2.3, the need for high resolution in ocean modeling is very great. In recognition of their importance to global-ocean circulation and to air-sea interactions, eddy-mixed-layer interactions have become the focus of NSF/NOAA-funded Community Process Team (CPT) EMILIE (see http://cpt-emilie.org/). In support of CPT-EMILIE, the CLIMODE (CLIvar MOde water Dynamic Experiment - see http://www.whoi.edu/science/PO/people/tjoyce/climode_website/index.html) project, recently approved and funded by NSF, is a major field experiment focusing on the role of eddy-mixed-layer interactions. CPT-EMILIE and CLIMODE are motivated by the belief that the representation of lateral eddy transfer through the mixed layer on scales of kilometers may be a critical missing element in coarse-resolution models. Diabatic eddy fluxes are particularly important in regions of frontal instabilities and vigorous air-sea exchange, such as the Gulf Stream extension, the Kuroshio, tropical instability waves, and the Antarctic Circumpolar Current (ACC). In ECCO II we propose to support and leverage ongoing CLIMODE activities by embedding very-high resolution models within the global eddy-permitting baseline integration of Fig.6. In particular we propose to increase resolution in two key regions: (i) in the North Atlantic, including the Gulf Stream, its recirculation, and the Labrador Sea (see Fig. 9) and (ii) in the Southern Ocean. Horizontal grid spacing ~1km is believed needed in order to be able to resolve the balanced part of the motion (which has characteristic scales of a few km), employing the Large et al. (1994) KPP scheme to represent vertical mixing processes. The CLIMODE field experiment will provide the observational setting for evaluating model results in the North Atlantic. Significant eddy-mixed-layer processes also occur all the way around the ACC, motivating the second embedded high-resolution experiment in the Southern Ocean.

Role of eddies in ocean circulation The high-resolution embedded models will also allow us to assess the role of small-scale eddies in the heat, freshwater, and vorticity budgets in regions of critical importance, within the context of a realistic global-scale model constrained by observations. We will carry out comprehensive diagnostics of the model fields to evaluate residual-mean flows, eddy-mixed layer interaction, meridional overturning circulation and ascertain the dominant term balances in the residual-mean momentum, vorticity, thermodynamic, and salinity budgets. Model fields will also be used to compute eddy diffusivities in order to inform parameterization schemes

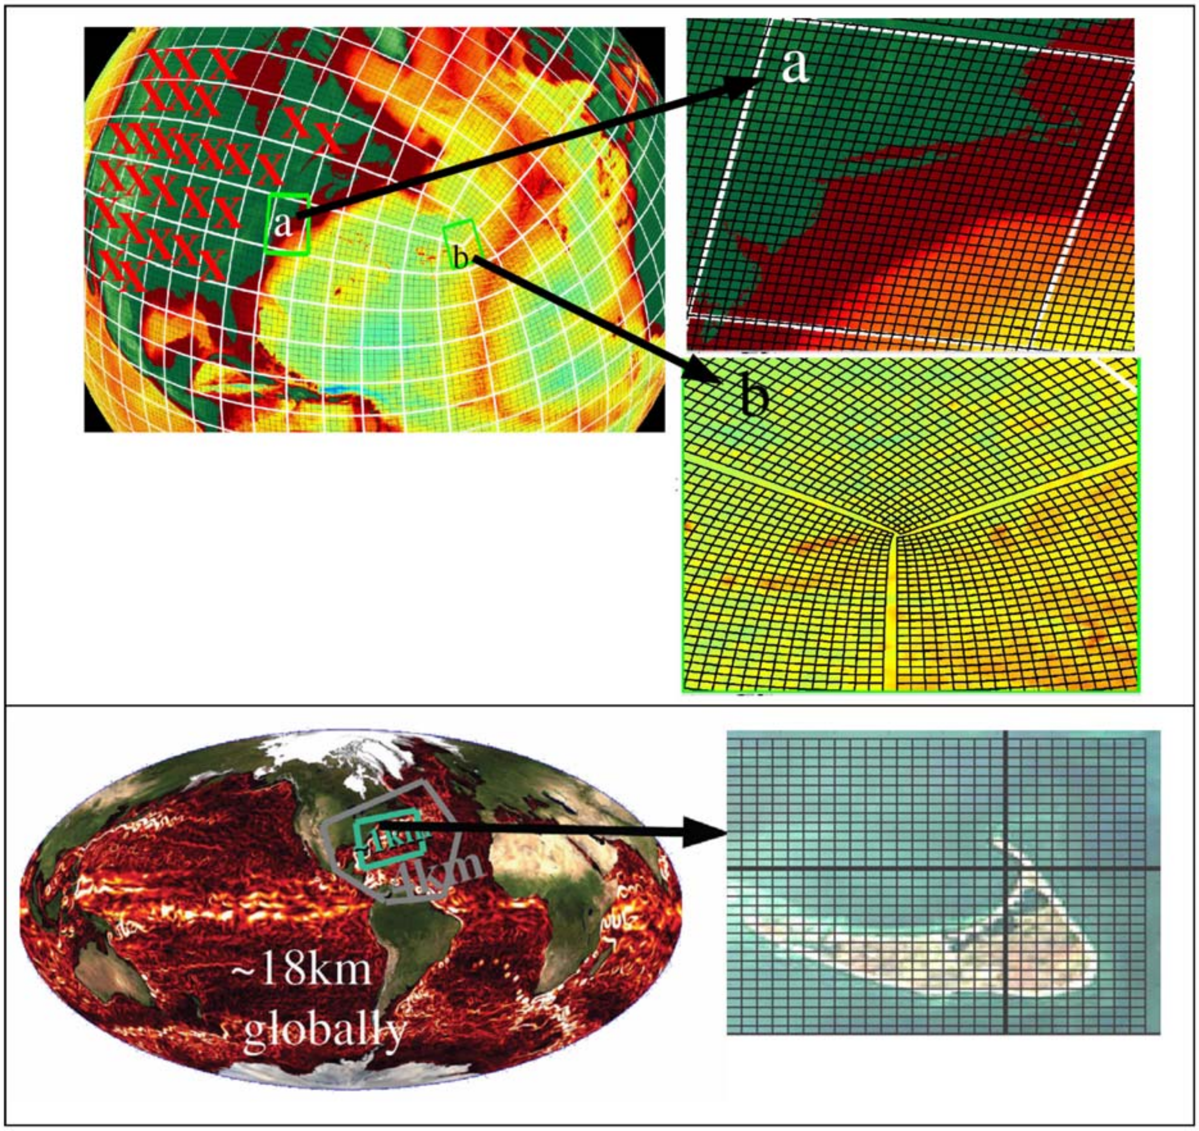

Figure 9: (Top) Baseline cubed sphere grid. The left panel shows the projection of the underlying cube basis onto the sphere. The white squares show 30×30 subdomains that are computed in parallel. The subdomains marked X are dry and are dropped from the computation. Magnified region 'a' shows the configuration grid cells in the region of the North Eastern continental US, the nominal resolution is $\sim$18km. Magnified region 'b' shows the resolution in the vicinity of a cube "corner". The final line of grid cells on each cube face has been deliberately omitted to show how three cube faces meet at each corner. At the cube corner the finest resolution is $\sim$3km. (Bottom) Using grid nesting on Columbia we will undertake simulations with an O(1km) resolution grid two-way nested within the global grid. Two nesting levels (the grey and magenta boxes in the left panel) will be used to "stepdown" from the global grid in factors of four. To understand the impact of this resolution we show on the right panel a 1km grid overlaid on a Modis image of Nantucket Island in New England. The 1km grid is the fine mesh, the thick black lines are 18km grid lines corresponding to our current baseline grid in this region. Preliminary tests suggest that we can perform nested simulations for the magenta region in the left panel using approximately 1000 Project Columbia processors. Process studies have shown that the impact of such resolution on ocean model fidelity is enormous. We will also use our nesting code to remove grid lines close to cube corners, which will permit high resolution isotropic grids to be used for global simulations on the entire Columbia system.

used for coarser-resolution ocean models, in particular the baseline 18-km cubed sphere simulations and synthesis.

Parameterization in coarse resolution models A productive avenue of research followed in ECCO has been to use adjoint models to help study the role of subgridscale processes in ocean circulation. For example, in Ferreira et al. (2004) we adopted a 'residual-mean' dynamical perspective in which 'eddy stresses' appear as a forcing term in the residual mean momentum equation — see, e.g., Marshall and Radko (2003), where residual-mean theory is applied to the ACC. Ferreira et al. (2004) use the adjoint model to adjust these eddy stresses in such a way that the drift of the model from observations is minimized. The resulting patterns of eddy stresses provide a fascinating perspective on the role of eddies in large-scale ocean circulation. The eddy stresses exist in the entire body of the fluid and are often as large as the applied wind-stress, particularly in regions of strong jets such as the Gulf Stream, the Kuroshio, and the ACC. When interpreted in terms of eddy transfer coefficients, K, we discover that the K's are not uniform but vary in space, becoming larger as the sea-surface is approached. These variations in K have now been implemented in global coarse-resolution models and shown to improve the climatology of these models.

The high-resolution forward simulations and synthesis proposed herein afford us a tremendous opportunity to directly compute eddy stresses accomplished by resolved eddies and to compare them to those inferred by the coarse-resolution ECCO estimates. We expect broad agreement in middle-to-high latitudes, but not in the tropics, where the residual mean theory employed by Ferreira et al (2004) breaks down. Comparison between the two models should suggest ways in which approximations in the residual-mean theory can be relaxed to render a theory that remains applicable as the Coriolis parameter goes to zero.

4.1.2 Gridding strategies (Adcroft and Hill)

In Section 3.2 we described the use being made of novel gridding techniques to render quasi-uniform resolution over the globe. The global integration of ocean and sea-ice on the cubed sphere (see Fig. 6) has the great advantage that polar processes are represented in an even-handed way. Nominal horizontal grid-spacing is 18 km, which admits but does not fully resolve eddies. This compromise resolves gross aspects of the geostrophic eddy field and yet still permits adequate throughput (10 years of integration per day of computation on 500 Project Columbia processors) to enable ocean state estimation. The 18-km cubed-sphere configuration will serve as baseline during initial stages of the work proposed herein. But to better address the science questions discussed above, we propose to take advantage of MITgcm non-hydrostatic and isotropic grid capabilities and of ESMF technology (see http://www.esmf.ucar.edu/) in order to carry out ultra-high resolution embedded simulations and in order to obtain a next-generation global grid for ocean state estimation.

Ultra-high resolution O(1km) simulations embedded in the global model To make ultra-high resolution embedded models possible, we will adopt advanced, static, techniques for grid refinement and grid coarsening along the lines demonstrated by Ginis et al. (1998), MacNiece et al, (2000), Colella et al, (2003), Penvan et al, (2004), and Jablonowski et al, (2004). We have started work on a two-way, multi-level nesting strategy in which nested domains with different resolutions are interconnected as ESMF components. The ongoing effort is focused on achieving 1-km resolution in the Gulf Stream extension region (see Fig. 9), a region that is of particular interest for CLIMODE science objectives. In contrast to the 18km baseline configuration, the dark lines in bottom right panel of Fig. 9, the 1km region will have coastlines and bathymetry defined with high precision and will resolve sharp fronts in water mass properties. Preliminary work on the Columbia supercomputer at NAS-ARC suggests that a throughput of three years of integration per day of computation could be achieved using only two of Columbia's twenty 512-processor systems. Much of the general-purpose, flexible technical infrastructure for interfacing between models of different resolution is being developed within ESMF. We are using this ESMF technology to facilitate the MITgcm multi-scale gridding work.

The two-way nesting and refinement illustrated in Fig. 9 bottom has many attractive practical qualities. It allows very general targeted simulations or assimilations to be carried out within a large-scale flow that can itself be constrained to consistency with observations. Subject to requirements of numerical stability, it is possible to concurrently execute a large scale global model and one or more regional fine mesh models. This can enable productive scaling to computational systems with tens of thousands of processors. In addition to the North Atlantic region, our objective is to embed a second high-resolution region within the 18-km cubed-sphere configuration, which will increase resolution in the Southern Ocean as discussed in Section 4.1.1. The ~1km grids are expected to yield a major qualitative change in the nature of the simulation, to shed new light on eddy-mixed-layer interactions, and to improve parameterization of sub-grid-scale processes in coarser-resolution models.

Fully isotropic, next-generation global grid Although they have much improved scaling over latitude-longitude grids, the current generation of expanded spherical cubic grids described in Section 3.2 have regions of locally enhanced resolution in the vicinity of the eight corners of the projected cube (see Fig. 9, top panel). In the baseline 18-km cubed-sphere grid adopted herein, the local resolution enhancement is such that at the cube corners horizontal grid spacing is 3 km. One solution to this problem is to employ nesting techniques in "reverse" in order to coarsen the grid in the vicinity of these poles and to produce a locally reduced grid with horizontal grid spacing of 18km, thereby yielding an almost perfect isotropic grid that is still based on an underlying quadrilateral, semi-structured grid.

A final evolution in our gridding work will be to apply spring dynamics to further refine the algebraically generated grids. The spring dynamics will be similar in principle to the spring meshes applied by Gnofo (2001) and others. However, we will use angle springs in conjunction with (or instead of) linear coil springs. This will provide more control over the resulting grid properties; whereas a mesh of coil springs allows one to optimize the distribution of nodes, a mesh of angle springs allow one to optimize the orthogonality of the grid. Using spring dynamics we will be able to optimize our grid to fit coastlines better, yielding a grid that will be a net of distorted rectangular blocks with unstructured connections.

At this final stage we anticipate that, we would be able to target global non-hydrostatic simulations at close to 1-km horizontal grid spacing and 100 vertical levels using all 10240 processors of the Columbia system. Preliminary estimates suggest that a year of simulation would require around ten days of wall-clock time for this problem. Such a calculation would be a tremendous technical achievement and it would shed considerable light on the difficult sub-grid-scale parameterization problem discussed in Section 4.1. Our progressive deployment of a hierarchy of nesting approaches means that we will be able to maximize science return as we build-up towards this ground-breaking global, 1-km integration.

4.1.3 Parallel implementation, ESMF, and MAP (Hill and Taft)

All the modeling and assimilation tools described here will be deployed as Earth System Modeling Framework (ESMF) components and will therefore be amenable to integration within the Modeling, Analysis and Prediction (MAP) program Modeling Environment. The cornerstone application for this work, the MITgcm, already supports ESMF component interfaces and has been demonstrated to interoperate with other ESMF components. Other key elements of our system include advanced data processing tools and shared-memory optimization, using the MLP library (Taft, 2001). These elements have already enabled the MITgcm code to achieve good scaling and input/output performance for up to 500 processors on the Columbia system. Our work will aim to obtain good scaling for the MITgcm, first on one of the 2048-processor Columbia superclusters, and eventually on the complete 10240-processor system, needed for the global ~1km grand challenge computation.

4.2 State estimation with global eddy-resolving models

The centerpiece deliverable of this proposal will be a best possible global-ocean and sea-ice synthesis at eddy-resolving resolution during the ocean satellite era (1978-present). Initially, an approach based on Green functions and the cubed sphere grid shown in Fig.6 and 9(top) will be used (Section 4.2.1). The Green-function estimation effort will provide a global, eddy-permitting estimation product for distribution early on in the life-cycle of the project. Towards the middle of the 5-year funding period, we plan to start the transition towards an adjoint-model-based method, provided the various bottlenecks, technical as well as computational, have been resolved (see Section 4.2.2). The oceanographic data to be used will involve all of those listed in Fig.1 as well as additional data constraints made possible by the higher resolution and by the inclusion of an interactive sea-ice model (see Section 3.1). Our existing estimation methods permit straightforward use of any data that can be adequately represented as a linear combination of the model state vector and for which a plausible data error estimate is available.

4.2.1 Green-function optimizations (Menemenlis)

When the number of control parameters is small, Green functions provide a simple yet effective methodology to linearize a general circulation model and to minimize a cost function (e.g., Wunsch 1996, Menemenlis et al., 2004a). In some earlier ocean state estimation applications (e.g., Stammer and Wunsch, 1996; Menemenlis and Wunsch, 1997; Fukumori, 2002) Green functions were used to obtain a coarse-scale representation of ocean GCM dynamics, for example, the GCM response to large-scale, geostrophically-adjusted density or sea-surface height perturbations. What we propose to do is different: rather than using Green functions to approximate GCM dynamics, they will be used to calibrate a small number of key GCM parameters and to blend together estimates from existing solutions and data products. This strategy has the advantage of permitting a relatively large impact on the solution from a small number of control variables. Additionally, the representation of GCM dynamics is implicit and exact rather than explicit and approximate. The methodology was successfully used to calibrate the existing ECCO quasi-operational ocean-circulation analysis (Menemenlis et al., 2004a). A total of twenty-six sensitivity experiments were used, resulting in substantial improvements of the solution relative to observations as compared to prior estimates. Overall the cost function was reduced by 43% due to a significant reduction in model bias and drift and to a 10–30% increase in explained variance.

Compared to other rigorous estimation methods, the key advantages of the Green function approach are simplicity of implementation, inherent parallel scalability, and improved robustness relative to the adjoint method in the presence of non-linearities. This is because the Green function approach does not need to rely on an exact, tangent linearization of the model, but can instead be used to linearly interpolate between several, arbitrarily-different forward-model trajectories. This last property is particularly important for eddy-resolving estimation. Green functions also provide an effective method for testing and for calibrating any new model parameterizations that are incorporated in the model, and for studying and quantifying model and data errors.

For the work proposed herein, we expect to compute approximately 100 Green functions, or forward-model sensitivity experiments. The sensitivity experiments will be chosen so as to have maximum impact on cost function reduction. First, they will include experiments that are used to calibrate empirical model parameters, for example, the parameters that control isopycnal and diapycnal mixing, open water and sea-ice albedos, drag coefficients, sea-ice strength, etc. Second, sensitivity experiments will be conducted for various temperature and salinity initial conditions, including estimates obtained from climatologies, from the coarse resolution ECCO optimizations, and from model spin-up solutions. Third, sensitivity experiments will be conducted for different surface boundary conditions, for example, estimates from meteorological analyses, from scatterometer data, and from coarse-resolution adjoint and smoother solutions that have been obtained as part of ECCO or that are being obtained as part of ECCO-GODAE. Finally, the Green-function solution will also make use of improved sub-grid scale parameterizations from the ultra-high-resolution studies of Section 4.1.

The computation of model Green functions for the high-resolution, global, cubed-sphere configuration is already under way on Columbia, e.g., Figs.6 and 7. Therefore, we expect that a preliminary global estimation product will be available within a few months of the beginning of this project. We will continue to upgrade and improve this solution during the first three years of the project so that by the time that the adjoint-model estimation technology is ready, a highly plausible baseline integration, that accurately describes most of the large-scale patterns and variability of the global-ocean and sea-ice circulations, will be available.

4.2.2 Adjoint-model optimizations (Heimbach, Tziperman, and Wunsch)

The adjoint method of state estimation and data assimilation, also known as 4D-Var, the Pontryagin Principle, or the method of Lagrange multipliers, has now been used in oceanography in a range of applications, from idealized model studies (Wunsch, 1988; Tziperman and Thacker, 1989; Thacker and Long, 1988), to coarse resolution GCMs (Tziperman et al., 1992a,b; Marotzke, 1992), to state of the art global models (Marotzke et al.,1999; Stammer et al., 2002, 2003), and to initializing ENSO prediction models (Galanti et al., 2002). There have been preliminary efforts using the adjoint method with eddy-resolving models as well (Schröter et al., 1993; Cong et al., 1998; Lea et al., 2000; Köhl and Willebrand, 2002, 2003; Gebbie, 2004; Gebbie et al., 2004), but the applications have been restricted to either idealized cases, or regions chosen to be only weakly turbulent. No global eddy resolving general circulation model, with its enormous range of physical regimes, has yet been used with the adjoint method. Achieving that capability is the primary goal of this proposal.

The adjoint method minimizes an objective function, J, which represents the mean-square model-data difference, subject to model constraints, which are imposed by the use of Lagrange multipliers. The stationary point of the resulting Lagrangian is found by an iterative search using the gradient of the Lagrangian w.r.t. the set of control variables. This gradient is given by the linearized dynamics about the evolving state of a forward model integration, the model Jacobian, and can be computed very efficiently by means of its transpose, the "adjoint model". An eddy-resolving model is inherently unstable (the instabilities that give rise to the eddy field), as is the corresponding adjoint model. While nonlinearities are able to arrest the instabilities in the forward model and prevent a shift to a completely different mean state, the linear adjoint model has no such non-linear arrest. Because the adjoint solution is used to calculate the gradient of J, an unstable solution would appear to be impossible to use in this way. In practice, however, the system has been rendered stable, and hence solvable, in a number of different ways, for example, approximate adjoints (Schiller and Willebrand, 1995), coarse-scale controls (Gebbie, 2004), and forcing the model to remain within some distance of a climatology (Gebbie, 2004).

The ongoing MITgcm development puts emphasis on maintaining the capability to generate an exact, efficient adjoint and tangent linear model via automatic differentiation (see Marotzke et al., 1999; Heimbach et al., 2002). Key advantages of this approach are: (i) the same strategies for achieving efficient scalable forward code carry over to the adjoint code (Heimbach et al., 2005); (ii) the adjoint code is exact w.r.t. to the state around which it is linearized, i.e., no assumption of equilibrium state is required, enabling the use of configurations where the adjoint code of complicated parameterization schemes is included, e.g., GM in Ferreira et al. (2004) and KPP in Gebbie et al. (2004); and (iii) the adjoint model is kept up-to-date w.r.t. the ongoing model development (Heimbach, 2003). This last point, in particular, makes available the adjoint code for most model components used by the cubed-sphere configuration of Fig. 6. Remaining components are limited to cubed-sphere-specific exchange routines for which hand-written adjoint primitives are currently being derived. The specific ECCO-II eddy-permitting adjoint-optimization strategy will include the following components:

1. Cubed-sphere adjoint: Development of cubed-sphere adjoint will first take place on a global, coarse-resolution setup that includes Arctic Ocean and sea-ice. Each face of the cube comprises 32 by 32 grid cells, for a mean horizontal grid spacing of 300 km, and there are 15 vertical levels. This coarse-resolution cube-sphere configuration is one that we have used extensively for code development work and which is also being used for millennial-scale oceanic and coupled ocean-atmosphere simulations (Campin et al., 2003). The small footprint of this particular configuration will keep computational cost to a minimum and hence speed-up development. Additionally, this effort will benefit the millennial-climate work and provide a good basis for later developing a truly global coupled atmosphere-ocean adjoint.

2. Exact high-resolution adjoint: The full high-resolution adjoint will be attempted for the global, 18-km cubed-sphere configuration and for two regional application, one in the North-Atlantic and one in the Southern Ocean. This will establish baseline time-horizons, over which adjoint sensitivities remain bounded. The sensitivity information provided directly by the adjoint will also be studied to infer potential "choke points" of the model. The regional applications will be useful both to study relevant processes, e.g., western boundary currents, sea-ice, and convection, in detail as well as to reduce computational cost during development of approximate-adjoint technology. The high-resolution exact-adjoint experiments will serve as baseline for implementing, where necessary, approximate adjoint approaches, discussed next.

3. Filtered or approximate adjoints: No approximation assumptions are made in deriving the exact adjoint by means of AD, leaving ample room for smoothing strategies: (i) the adjoint integration could be interrupted periodically (e.g. over an expected "predictability horizon"), and the adjoint snapshot (as well as the required forward state) replaced by time-averaged quantities over the previous time interval; (ii) the adjoint integration could be performed on a coarser resolution than the forward model; (iii) the adjoint model could use simplified numerical schemes compared to the forward model, such as switching from higher-order flux limited advection schemes to lower-order non-limiting schemes. This extends ideas already implemented, such as simplifying/ignoring adjoint code of parameterization schemes (KPP, GM/Redi) or changing parameters in the adjoint sweep (such as viscosity). All strategies can be readily implemented and maintained "automatically" by taking advantage of the adjoint-model's checkpointing scheme which provides a natural interface for inserting spatial and temporal filters.

4. Formulation of the cost function: There will be some experimentation with the formulation of the cost function, for example, such as those that penalize statistical aspects rather than detailed snapshots of the circulation and more emphasis on including covariance information. Note that although the eddy field is sensitive to the precise initial conditions due to its chaotic and turbulent

nature, the long time mean of the ocean circulation is not expected to be sensitive to the details of a specific realization of the eddy field. Furthermore, including model parameters and time-dependent surface forcing fields as control variables relaxes the control problem from a pure initial value to a combined initial and boundary value problem. The mean circulation is only affected by the long time mean of the eddy transport and Reynolds stress terms. Given the insensitivity of the time mean circulation, constraining the long term time mean of the circulation to be near the long term time mean observations may stabilize the adjoint optimization problem.
5. Control parameters: There will also be experimentation with the formulation of control variables, for example, the addition of terms like eddy stresses used in Ferreira et al (2004) that decrease the rigidity of the minimization or of coarse-resolution controls as in Gebbie (2004).

Combinations of the above approximate-adjoint approaches will be developed and tested in the small regional domains first (one in the North-Atlantic and one in the Southern Ocean) before being applied to the global, 18km cubed-sphere model configuration towards the middle of the 5-year funding period.

4.2.3 Open-source automatic differentiation (Utke, Hovland, and Naumann)

In addition to rigorous estimates of the global-ocean and sea-ice circulations, another objective of this proposal is the development and distribution of ESMF-compatible modeling and estimation tools. The central computational engine for adjoint-method optimization, sensitivity, and forecast studies is the automatic differentiation (AD) tool, which generates the adjoint model. At the present time, the only available software capable of generating efficient code for the so-called reverse mode (adjoint) of a state-of-the-art GCM is TAF, a commercial product distributed by FastOpt. Although TAF has made the ECCO computations possible and is being used to derive the adjoint of several other GCMs, including GMAO's global weather forecast model, the generation of efficient adjoint code still requires substantial manual intervention. Because access to the TAF source code is restricted, it is at times difficult to identify problems in the generation of the adjoint code or to implement novel AD strategies. Therefore, it is highly desirable to have an alternative, open-source product available. With support from a NSF ITR grant, the MIT group, along with collaborators at Argonne National Laboratory and Rice University, has been developing an open-source AD tool — which we call OpenAD (see http://www-unix.mcs.anl.gov/∼utke/OpenAD) — that should eventually prove useful as a public domain, open-source tool with widespread applications, including some remote from oceanography. Such tools have wide applications in any branch of science, engineering or economics where complex models are in use, sensitivity tests are required, or where optimization is to be carried out. The oceanographic application, however, remains to this day, the most sophisticated and demanding use to which these AD tools have been put. We are asking support for development of OpenAD in two major directions that are required in the ECCO-II project:
1. Automated selection of transformations, reversal, and checkpointing schemes: OpenAD contains a variety of local code transformation algorithms implementing different levels of preaccumulation. Each of these transformations results in codes which have different complexities with respect to their local run time and storage requirements. For an adjoint model the global run time and storage requirements are affected by the choice of a particular reversal scheme with checkpointing that builds upon the local preaccumulations (a technical term referring to trade-offs between computation and storage). Currently, the selection of efficient local preaccumulation strategies and a global reversal scheme, including the placements of checkpoints, requires a highly educated guess. To manually place a checkpoint at a certain location in the code, the user needs to know which variable values have to be restored to properly restart the computation from that

checkpoint. It is difficult to determine this variable set without data flow analysis. OpenAD already uses some data flow analysis features of OpenAnalysis. OpenAD can be improved to accept user directives for manual checkpoint placement to automatically approximate the minimal set of variables to be stored at the checkpoint. We propose to implement a static evaluation of the codes generated for the various preaccumulation strategies that allows an automatic, usage-specific, best-fit selection.

2. Robustness and canonicalizations: OpenAD uses the Fortran90 front-end of Open64 via the OpenAD Fortran Toolkit (OpenADFortTk) and a language independent code transformation engine to generate the derivative code. Due to the inherent complexity of the semantic transformations for automatic differentiation the application of the OpenAD tools pipeline to new Fortran90 codes can reveal subtle problems with any one of the components. To make OpenAD sufficiently robust for general use these problems have to be analyzed and removed. The implementation of automatic differentiation algorithms in a language-independent fashion requires certain "canonicalizations" to be performed by OpenADFortTk. The basic canonicalizations are currently being implemented in OpenAD. Computationally challenging problems can benefit from advanced distinctions with respect to the control flow, for instance by exploiting explicitly reversible "simple" loops versus generic loops or varying approaches to array operations. OpenAD can provide these advanced canonicalizations to reduce the computational and storage requirements.

4.3 Evaluation, visualization, and science applications

Efforts described below aim to evaluate and to visualize the high-resolution estimation products and to demonstrate their utility via a small number of science applications.

4.3.1 Evaluation and diagnosis of ocean-circulation results (Fu, Lee and Zlotnicki)

The high-resolution ocean state estimation products will be evaluated through comparisons with all available satellite altimetry measurements, and especially the data from the Wide-Swath Ocean Altimeter if it will be demonstrated with OSTM (Ocean Surface Topography Mission). The comparisons will include the eddy energy level, eddy geographic distribution, and frequency-wavenumber spectrum. The vertical structure of the ocean state estimates will be compared to ARGO-float, XBT, XCTD, and other in-situ data. Coarser averages of the model bottom pressure output will be compared against GRACE data to assess deep zonal transport changes in the Antarctic Circumpolar and North Pacific regions. The purpose will be to identify not only overall energy levels, but correct location of currents and water masses.

Existing coarse-resolution ECCO products have been used to elucidate dominant processes responsible for the interannual variations of upper-ocean temperature in the equatorial and mid-latitude Pacific Ocean (e.g., Lee et al., 2004; Kim et al., 2004). A recent study by Roemmich and Gilson (2001), for example, suggests that eddies in the North Pacific may hold the key to interannual and decadal climate variability. We will therefore examine the effects of eddies on upper-ocean heat-content and meridional heat transport. In particular, the contribution of eddies to heat transport convergence in mid-latitude Atlantic and Pacific Oceans will be evaluated, and its impact on climate variability assessed.

4.3.2 Evaluation of sea-ice results (Kwok and Menemenlis)

One objective of this proposal is to evaluate the various data products and to identify those that are most suitable and useful for assimilation in climate models. In addition to the list of sea-ice satellite retrievals that are discussed in Section 3.1 and listed on Table 1, we plan to use a number of other sea-ice data products in order to establish error bars and in order to test the estimation results. These will include data products generated using the Radarsat Geophysical Processor System (RGPS), scatterometer retrievals, a study of Arctic Polynyas, and various in-situ observations.

Sea-ice motion from RGPS, scatterometers, and buoys. In addition to data products listed in Table 1, sea-ice motion is available from RGPS, scatterometer imagery, and from buoy drift. In particular, twice-weekly RGPS coverage of the entire Arctic Ocean has been obtained continuously since the winter of 1996/97, providing fine-scale ice motion fields and derived estimates of deformation, lead openings and closings, ice age, and thin ice volume. These data, processed by R. Kwok, will be used to establish error bars and to evaluate the quality of the estimated sea-ice distribution.

Arctic polynyas. Satellite-derived time series of brine and heat flux from Arctic polynyas are being assembled, under separate NASA funding, by team member R. Kwok, along with B. Holt (JPL) and S. Martin (Univ. Washington). These time series consist of maps of polynya area, ice thickness, and brine flux for the Arctic Ocean from 1987 to the present; they are derived based on empirical sea-ice growth-rate models that regress observed in-situ ice thickness data to Sea-Surface Temperature (SST) retrievals from AVHRR and MODIS and to the 37-GHz polarization ratio of SSM/I and AMSR data. These estimates will also be used for assessing the quality of the estimated sea-ice distribution.

Ice thickness. The most extensive and valuable measurements of Arctic sea ice thickness have been obtained by submarine-mounted upward looking sonar. A series of cruises, called Scientific Ice Expeditions (SCICEX), took place from 1993-1999 and the resulting ice draft measurements are continuing to be released. JPL team members Kwok and Menemenlis will assemble these data and combine them with in-situ observations, upward looking sonar buoys, and with earlier submarine measurements to be used for diagnostic purposes by this project.

4.3.3 Advanced scientific visualization (Henze)

New visualization methods, that permit new scientific insight on the enormous simulation and estimation output fields to be generated, will be developed and applied to the high-resolution ECCO-II results. The emphasis will be (i) on understanding the model sensitivities to various parameters, for example, multifield visualizations of the Green function experiments and (ii) on detailed model-data comparisons, for example animations of sea-surface temperature or of high-resolution altimeter data compared to the ECCO solutions. A key tool in this analysis will be a custom-built array of flat panel screens, each harnessed to its own computer, called the hyperwall, which permits multiple simultaneous viewpoints into the ECCO solutions, for example, the simultaneous display of a model sensitivity response for multiple model fields and depth levels.

4.3.4 Climate sensitivity and prediction (Heimbach, Tziperman, and Wunsch)

The state estimation problem as formulated and solved in this proposal is intimately tied to the prediction problem. We elucidate this in the context of ENSO. In spite of much progress over the past decade in ENSO theory (Neelin et al., 1998) and assimilation and prediction methods (Latif et al., 1998), the operational prediction skill is still not robust, nor satisfactory. Furthermore, the skill of current dynamical prediction models does not significantly exceed those of far simpler statistical approaches. Significant improvements to the prediction skill might be expected through (i) the use of dynamically consistent, high-resolution state estimates for prediction initialization, (ii) the use of a rigorous assimilation technique, such as the adjoint, which produces error estimates through the model dynamics itself and may smoothly link the estimation and prediction phase, and (iii) the use of adjoint tools, closely related to the tool used for state estimate, to investigate the underlying dynamics which set the predictability horizons. All three aspects directly benefit from the proposed work because of the high-resolution model and its adjoint and because of the accurate state estimates.

High-resolution adjoint prediction-type assimilation. ENSO prediction, as with any forecast, is dependent upon the accuracy of the initial state. Existing initial state estimation methods (Chen et al., 1995; Kirtman and Zebiak,1997, Derber and Rosati,1989; Behringer et al., 1998) use practical, but approximate, methods. It is a reasonable expectation that a more rigorous system that uses the full power of the model and known statistics would improve the ENSO-prediction horizon.

Current state of the art GCMs that are used for ENSO prediction have a typical resolution of 1°, with a higher resolution of 1/3° in latitude near the equator. Such a resolution suffices for marginally resolving the main equatorial Kelvin and Rossby modes participating in ENSO dynamics, but does not resolve eddy processes. In particular, a more satisfactory representation of equatorial instability waves which are seen to develop and break at the margins of the tropical Pacific, may provide an important mechanism of heat transport into the cold tongue region, and significantly affect the heat budget there. It is critical therefore, that such eddy processes be well represented in the models. Although some work using adjoints has occurred in the ENSO context (e.g., Bennett and McIntosh, 1982; Thacker and Long, 1988; Tziperman and Thacker, 1989; Bennett et al., 1998; Kleeman et al., 1995; Lee et. al. 2000, Bonekamp et al., 2001; Weaver et al., 2002; Galanti et al., 2002) both for theoretical and for operational purposes, there is only very limited experience in using this method with eddy resolving physics.

The developments outlined in the previous sections enable us to address, as an immediate spin-off, these two major open issues in the ENSO prediction field. A high-resolution state estimate of the tropical Pacific will be readily available studying equatorial processes, based on the global ECCO-II state estimate. The underlying forward and adjoint models can be used to extend the estimation (analysis) phase of the integration into a prediction phase. Linking the estimation and prediction phase in the assimilation using the same model would minimize the so-called initial shock caused by physical inconsistency between the initial state and the forecast-model physics and could thus markedly improve the prediction. The adjoint includes an explicit treatment of model error and is especially formulated for minimizing such an initial shock. Further assimilation studies would focus on the use of satellite observations and on the quantification of their contribution to the predictability skill of ENSO.

Application of adjoint tools for understanding the circulation. Adjoint models are used to define the partial derivatives of objective functions, J, with respect to numerous control variables. They are thus a very general and efficient representation of model sensitivity to any and all of the parameters defining a model, including initial/boundary conditions as well as internal parameters such as mixing coefficients. One can define J to represent not only model/data misfits, but any (scalar) quantity of interest in the ocean circulation or climate (heat flux divergences, for example). A number of preliminary studies have used adjoint models in sensitivity studies (Marotzke et al, 1999; Sirkes and Tziperman, 2001; Van-Oldenborgh et al, 1999; Junge and Haine, 2001; Galanti and Tziperman, 2003; Li and Wunsch, 2004) for a wide variety of purposes. Applications fall under two main categories: (i) sensitivity analysis and (ii) calculating optimal initial conditions and "stochastic optimals" for transient amplification.

Adjoint sensitivities permit a deep understanding of physical processes that is not possible with forward models alone. One example is the discovery of a teleconnection mechanism between the mid latitudes and the equatorial Pacific, based on the amplification of propagating planetary waves by baroclinic instability (Galanti and Tziperman, 2003). Another example is the analysis of the adjoint sensitivities in terms of wave motion ("dynamic" sensitivity) vs. along-isopycnal advection of T,S, that does not affect the density ("kinematic" sensitivity) by Marotzke et al. (1999).

Related, is the calculation of optimal perturbations. In otherwise stable, "non-normal" linear systems, initial perturbations can display powerful transient amplification before ultimate decay (see Farrell and Ioannou, 1996, and references therein). Symbolically, such linear (or linearized) systems can be written,

$$\frac{d}{dt}\mathbf{x}\left(t\right) = \mathbf{A}\left(t\right)\mathbf{x}\left(t\right), \tag{2}$$

where $\mathbf{x}(t)$ is the state vector and $\mathbf{A}$ the non-normal state transition matrix, i.e. $\mathbf{A}^T\mathbf{A} \neq \mathbf{A}\mathbf{A}^T$. Most linearized fluid systems are non-normal, as has been demonstrated in the application of this concept to ENSO (Penland and Sardeshmukh, 1995; Kleeman and Moore, 1997; Moore and Kleeman, 1999 and Thompson and Battisti, 2001), and in a variety of oceanic and atmospheric settings (Moore, 1999; Farrell and Ioannou, 1996; Aiken et al, 2003; Tziperman and Ioannou, 2002; Zanna and Tziperman, 2004).

Deriving optimal perturbations relies heavily on the tangent linear and adjoint model to solve a generalized eigenvalue problem. A natural application spin-off of this proposal, then, is to use this powerful tool in an advanced dynamical framework, e.g. to the study of ENSO. An atmospheric component will be added to the ECCO-II configuration, either in the form of a statistical model, or a model of intermediate complexity which can be readily adjointed. Optimal growth structures will be computed for suitable norms over various time horizons, and will shed light on dominant processes and potential for improving predictability limits. Effort will go to the investigation of the theoretical limits of applicability of the adjoint analysis tools, i.e., to determine when the linearity assumption prevents application of the proposed mechanisms to the actual nonlinear systems. Existing studies were done without mesoscale eddies, including them will require both theoretical and practical development work.

4.3.5 Error covariances for GMAO (Rienecker and Suarez)

As with any least squares estimation procedure, the solution minimizing the objective function J in Eq. 1 is no better than the weight matrices $\mathbf{Q}^{-1}$ and $\mathbf{R}^{-1}$, used to represent the inverse of a priori error covariance matrices for the model and for the data errors, respectively. In particular, the need to improve estimates of $\mathbf{R}$ has been repeatedly identified by various groups including, for example, both the 2001 and 2004 US GODAE Workshops.

In most current estimation schemes, matrix $\mathbf{R}$ is usually prescribed in an ad hoc manner. Observational errors are usually taken to be uncorrelated in space or time (diagonal $\mathbf{R}$), even though significant space/time covariances are known to exist but are rarely quantified. Use of diagonal $\mathbf{R}$ is throwing away important information and inevitably degrades the possible solutions, even when the diagonal elements are accurately specified, which is often not the case. In general terms, $\mathbf{R}$ includes observational errors as well as model errors, which are not already represented in $\mathbf{Q}$. Model error is made up of two parts, the first owing to missing or partially erroneous model physics, the second largely due to inadequate resolution. For example, a non-eddy resolving model can be used with eddy-resolving data by regarding the observed eddy field as part of the observational noise.

Model errors can be estimated for low resolution models by comparison with high resolution ones. The very high resolution simulations and state estimates that will be generated from this study will provide an opportunity to estimate representation errors – including multi-variate covariances – that should be useful for the ocean state estimation system used by the GMAO (Goddard Modeling and Assimilation Office) to initialize coupled seasonal forecasts. The current GMAO global ocean data estimates are made using a 1/3° x 5/8° grid, and the system targets the larger spatial scales of seasonal variability (e.g., Troccoli et al., 2003; Borovikov et al., 2004; Keppenne et al., 2004). Here we propose to analyze the distribution of higher frequency, higher wave number variability of the ECCO-II state estimates and use this information for representation errors in the GMAO estimation system. The estimates will be compared with data from the high resolution XBT sections and from moored arrays. In addition, the high resolution estimates will be evaluated against equatorial Pacific observations (as, e.g., in Borovikov et al., 2004, and Sun et al., 2004) and against the ocean states used to initialize the GMAO coupled forecasts. If the evaluations show improved representation of the seasonal-to-interannual scales, one would like to estimate the representation error variance and covariances by data analyses and comparisons with models. There are a few data sets that are of sufficiently high resolution so as to have potentially useful information to help with the error estimation. The high resolution VOS XBT sections are one example, the time series available from moored buoys is another. However, it is likely the space-time sampling from the high resolution XBT lines is inadequate and variations in ship tracks, which can sometimes be displaced significantly, introduce uncertainties into statistics. Any improvements in $\mathbf{R}$ obtained would obviously also be useful in global, low resolution systems.

5 Expected results and significance

The proposed work will strengthen the connection between NASA's modeling and remote sensing strategies and enhance the value of NASA satellite retrievals for studies and discussions about climate and climate variability. Key deliverables are summarized in Section 5.1. Plans for data distribution are in Section 5.2. Expected significance and relevance to NRA research topics are discussed in Section 5.3.

5.1 Key deliverables

1. Preliminary eddy-permitting, 1992-present, global-ocean and sea-ice synthesis. A preliminary, eddy-permitting synthesis using model Green functions (see Section 4.2.1), will be delivered within six months from the start of this project on the 18-km cubed-sphere grid of Fig. 6. This synthesis will be updated and refined on a routine basis, made available to the community, and serve as an initial baseline for science applications.

2. Detailed evaluation of above solution. Quantitative statements about residual uncertainties will be accompanied by recommendations for additional focused field campaigns, modeling developments, and assimilation studies. Tracer, heat, and freshwater transports will be diagnosed and compared to existing adjoint-method solutions and to other independent estimates.

3. Demonstration science applications. These will include analysis of heat budgets and climate variability, studies of climate sensitivity and predictability, and estimates of prior error covariance for the GMAO seasonal-to-interannual prediction effort. The eddy-resolving estimate will be used in regional focus studies, e.g., the tropical Pacific, and adjoint-based tools (sensitivity and singular vector calculations) will be used to elucidate the dynamics that govern the limit of predictability of ENSO or its dynamical links to slow, long-term variability.

4. High-resolution simulations. As discussed in Section 4.1, two ultra-high-resolution regional models, O(1 km) grid, will be embedded within the 18-km cubed-sphere synthesis, one in the North Atlantic and one in the Southern Ocean, in order to study eddy-mixed-layer interactions and to improve parameterization of sub-grid-scale processes in coarser-resolution models. In addition, a 1-year, grand-challenge integration will be attempted, using an unstructured global grid with 1-km horizontal spacing and making simultaneous use of all Project Columbia processors.
5. ESMF-compatible adjoint-model estimation infrastructure. As per existing MITgcm and ECCO policies, the coupled ocean and sea-ice adjoint model estimation infrastructure will be immediately available to the community through an anonymous CVS server (http://mitgcm.org). OpenAD, the open-source automatic differentiation tool, will continue to be made available at http://www-unix.mcs.anl.gov/∼utke/OpenAD.
6. A best possible, high-resolution, global-ocean and sea-ice synthesis. At the end of year 2, based on input from above activities, feedback from science community and funding agency, and on available computer resources, a decision will be made regarding resolution, model grid configuration, period of synthesis (anticipated to 1978-present), and methodology for the centerpiece deliverable of this proposal, a best-possible, high-resolution, global-ocean and sea-ice synthesis.

5.2 Distribution of ECCO-II synthesis products (JPL PODAAC)

Effective community distribution of the huge outputs generated by ECCO-II computations is a crucial element for maximizing science payoff. The existing 18-km cubed-sphere ECCO computations of Fig. 6, using the current standard of storing three-day-averaged raw fields for the entire estimation period, together with higher resolution "burst" periods, results in approximately 20TB of output per decade of integration. The distribution and analysis of such huge outputs to the early users that have requested results so far, has proved to be a challenging task. As the number of science users as well as the size of the outputs are expected to grow during the five-year period of this project, we request some resources to help with the distribution of ECCO-II synthesis products to the science community.

The ECCO-II data distribution strategy is: (i) to take advantage of existing and planned NASA infrastructure, for example, the existing tape storage archives at NAS-ARC and the planned NASA Research and Education Network (NREN) and National Lambda Rail (NLR) activities, and (ii) to enlist the JPL Physical Oceanography Distributed Active Archive Center (PODAAC). In the recently-completed ECCO project, Live Access Server (LAS) and Distributed Oceanographic Data System (DODS) technology was deployed to provide access to the estimation results (http://www.ecco-group.org/data_server.html). ECCO-II will continue this practice. However, because raw results will now be tens to hundreds of terabytes for each synthesis product, a two-tier system will be deployed. A first-tier output-archive, located at JPL and maintained by JPL-PODAAC, will provide access to a subset of the high-resolution estimation products, including quick-look results that have been averaged in time and in space to reduce their size. A second-tier output system, whose details will be worked out between JPL-PODAAC, NAS-ARC, and ECCO-II science users, will permit access to the raw model output fields.

5.3 Expected significance and relevance to NRA objectives

This proposal addresses Earth Science Enterprise (ESE) questions in the Climate variability and Change focus area: *"How is the global ocean circulation varying on interannual, decadal, and longer time scales?"*, a specific focus of this NRA. The work reinforces GMAO thrusts via development of modeling (Section 4.1.2) and automatic differentiation (Section 4.2.3) tools, and of a next-generation

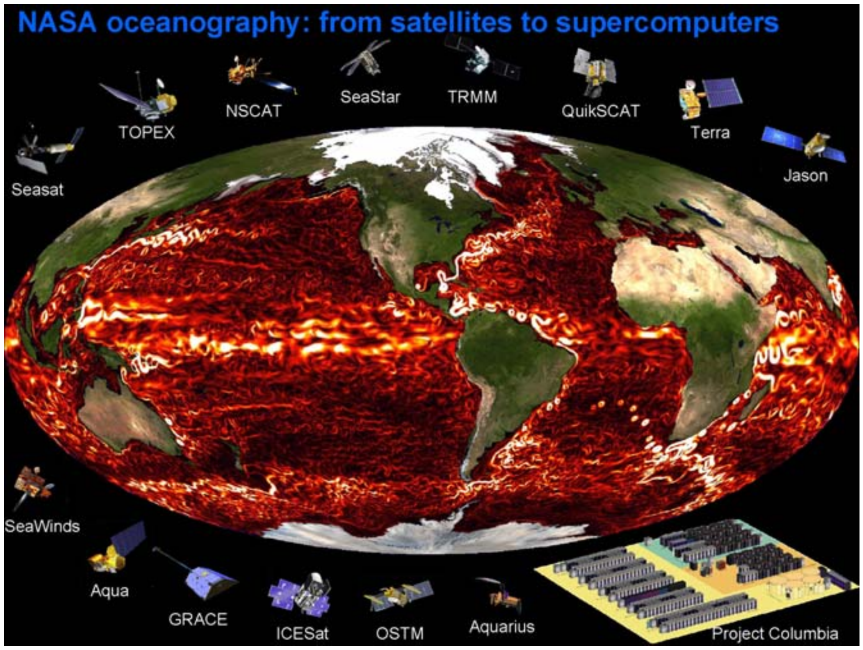

Figure 10: Schematic representation of vision that motivates this proposal: NASA satellite missions and supercomputers, used to deliver the best possible reanalysis of the global-ocean and sea-ice circulation. The proposed work uses sea-surface height from TOPEX and Jason, wind stress from NSCAT, QuiSCAT, and SeaWinds, sea-ice, temperature and albedo from Terra/MODIS, sea-surface temperature from TRMM/TMM and from Aqua/AMSR-E, GRACE geoid information, ICESat sea-ice thickness, computation cycles from Project Columbia. SeaStar/SeaWIFS chlorophyll concentration data may be useful for penetrative solar rradiance. Seasat, OSTM, and Aquarius represent the past and future of NASA oceanography missions.

data assimilation system (Section 3.5), and via estimates of error covariances for the GMAO prediction efforts (Section 4.3.5). The proposed work also supports Goddard Insitute for Space Studies (GISS) via development of new subgridscale ocean model parameterizations (Section 4.1.1), adjoint sensitivities and optimal perturbation studies (Section 4.3.4).The proposal work would provide a definitive demonstration of the impact of NASA's satellite missions on the understanding of global-ocean and sea-ice dynamics and on air-sea interaction when energetic eddy scales are included, a goal hitherto considered unattainable. Fig. 10 represents the vision that motivates this proposal, that is, the use of NASA resources, satellite missions and supercomputers, to deliver a best possible synthesis of the global-ocean and sea-ice circulation. The ocean and sea-ice state estimates that will be provided by this project will be physically consistent, with balanced budgets of mass and energy and with no temporal discontinuities, thus allowing improved understanding of the processes governing the wide spectrum of variabilities. This makes our proposed effort stand apart from the many other high-resolution estimation efforts, whose focus is operational ocean nowcast and forecast, and which admit temporal discontinuities in their solutions. The end result will be a unique, consistent description of the time-dependent, full-depth global-ocean circulation and sea-ice distribution, with unprecedented resolution and accuracy, suitable for adressing some of the most challenging and pressing questions and uncertainties concerning current and future climate system behavior.

6 References

Adcroft, A.J., C. N. Hill, and J. Marshall (1997). Representation of topography by shaved cells in a height coordinate ocean model. Mon. Weather Rev., 125:2293–2315.

Adcroft, A., J.-M. Campin, C. Hill, and J. C. Marshall (2004). Implementation of an atmosphere-ocean general circulation model on the expanded spherical cube. Mon. Weather Rev., in press.

Aiken, C. M., A. M. Moore, and J. H. Middleton (2003). Non-normal perturbation growth in idealised island and headland wakes. Dynamics of Atmospheres and Oceans, 37(3):171–195.

Behringer, D. W., M. Ji, and A. Leetmaa (1998). An improved coupled model for enso prediction and implications for ocean initialization. part i: The ocean data assimilation system. Mon. Weather Rev., 126(4):1013–1021.

Bennett, A. F., B. S. Chua, D. E. Harrison, and M. J. McPhaden (1998). Generalized inversion of tropical atmosphere-ocean data and a coupled model of the Tropical Pacific. J. Clim., 11(7):1768–1792.

Bennett, A. F. and P. C. McIntosh (1982). Open ocean modelling as an inverse problem: tidal theory. J. Phys. Oceanogr., 12:1004–1018.

Bonekamp, H., G. J. Van oldenborgh, and G. Burgers (2001). Variational assimilation of tropical atmosphere-ocean and expendable bathythermograph data in the hamburg ocean primitive equation ocean general circulation model, adjusting the surface fluxes in the tropical ocean. J. Geophys. Res., 106(C8):16693–16709.

Borovikov, A.Y., M.M. Rienecker, C.L. Keppenne, and G.C. Johnson (2004). Multivariate error covariance estimates by Monte-Carlo simulation for assimilation studies in the North Pacific, Mon. Weather Rev., submitted.

Caldeira, K. and P. Duffy (2000). The role of the Southern Ocean in uptake and storage of anthropogenic carbon dioxide. Science, 287:620–622.

Campin, J.-M., A. Adcroft, C. Hill, and J. Marshall (2003). A coupled ocean-atmosphere GCM using a single hydrodynamical kernel. In: International Conference on Earth System Modelling, p. 45. Max-Planck-Institute for Meteorology, Hamburg.

Carsey, F. D., R. G. Barry, and W. F. Weeks (1992). Introduction. In F. D. Carsey, editor, Microwave Remote Sensing of Sea Ice, chapter 1, pages 1–7. AGU, Washington, D.C.

Chen, D., S. E. Zebiak, A. J. Busalacchi, and M. A. Cane (1995). An improved procedure for El Nino forecasting: Implications for predictability. Science, 269:1699.

Colella, P., D. T. Graves, T. J. Ligocki, D. F. Martin, D. Modiano, D. B. Serafini, and B. Van Straalen (2003). Chombo Software Package for AMR Applications Design Document. Applied Numerical Algorithms Group, NERSC Division, Lawrence Berkeley National Laboratory, Berkeley, CA.

Cong, L. Z., M. Ikeda, and R. M. Hendry (1998). Variational assimilation of geosat altimeter data into a two-layer quasi-geostrophic model over the newfoundland ridge and basin. J. Geophys. Res., 103(C4):7719–7734.

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

Derber, J. and A. Rosati (1989). A global oceanic data assimilation system. J. Phys. Oceanogr., 19:1333–1347.

Dickey, J.O., S.L. Marcus, O. de Viron, and I. Fukumori (2002). Recent earth oblateness variations: unraveling climate and postglacial rebound effects. Science, 298(5600):1975–1977.

Duffy, P. B., M. Eby, and A. J. Weaver (2001). Climate model simulations of effects of increased atmospheric CO2 and loss of sea ice on ocean salinity and tracer uptake. J. Clim., 14:520–532.

Dutay, J.-C., J. Bullister, J. C. Orr, R. G. Najjar, M. Follows, R. Matear, S. Doney, E. Maier-Reimer, Y. Yamanaka, H. Drange, A. Yool, J.-M. Campin, M.-F. Weirig, N. Gruber, and K. Caldeira (2002). Comparison of simulations of CFC-11 and CFC-12 during OCMIP. Ocean Modeling, 4:89–120.

Farrell, B. F. and P. J. Ioannou (1996). Generalized stability theory part i: autonomous operators. J. Atmos. Sci., 53:2025–2040.

Ferreira, D., J. Marshall and P. Heimbach (2004). Estimating eddy stresses by fitting dynamics to observations using a residual-mean ocean circulation model. J. Phys. Oceanogr., submitted.

Fu, L.-L., and A. Cazenave, editors (2001). Satellite Altimetry and Earth Sciences: A Handbook of Techniques and Applications. Academic Press, San Diego, 463 pp.

Fu, L.-L. editor (2003). Wide-Swath Altimetric Measurement of Ocean Surface Topography. JPL Publication 03-002. Jet Propulsion Laboratory, Pasadena, CA, 67 pp.

Fukumori, I. (2002). A partitioned Kalman Filter and smoother, Mon. Weather Rev., 130, 1370–1383.

Galanti, E., E. Tziperman, M. Harrison, A. Rosati, and Z. Sirkes (2002). A study of ENSO prediction using a hybrid-coupled model and the adjoint method for data assimilation. Mon. Weather Rev., 131:2748–2764.

Galanti, E. and E. Tziperman (2003). A mid-latitude ENSO teleconnection mechanism via baroclinically unstable long rossby waves. J. Phys. Oceanogr., 33:1877–1888.

Gebbie, G. (2004). 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/.

Gebbie, G., C. Wunsch, and P. Heimbach (2004). Eddy-Resolving State Estimation with a North Atlantic Regional Model. In preparation.

Gent and McWilliams (1990). Isopycnal mixing in ocean circulation models. J. Phys. Oceanogr. 20(1):150–155.

Giering, R. and T. Kaminski (1998). Recipes for adjoint code construction. ACM Transactions on Mathematical Software, 24:437–474.

Giering, R. and T. Kaminski (2003). Applying TAF to generate efficient derivative code of Fortran 77–95 programs. PAMM, 2(1):54–57.

Gildor H., and E. Tziperman (2001). A sea ice climate switch mechanism for the 100-kyr glacial cycles. J. Geophys. Res., 106(C5):9117–9133.

Ginis, I., R. A. Richardson, L. M. Rothstein (1998). Design of a Multiply Nested Primitive Equation Ocean Model. Mon. Weather Rev., 126:1054–1079.

Gnoffo, P. (2001). A vectorized, finite volume, adpative-grid algorithm for Navier-Stokes calculations. Numerical Grid Generation, 819–835.

Gross, R. S., I. Fukumori, D. Menemenlis, and P. Gegout (2004). Atmospheric and oceanic excitation of length-of-day variations during 1980–2000. J. Geophys. Res., 109:B01406.

Heimbach, P., C. Hill and R. Giering (2002). Automatic Generation of Efficient Adjoint Code for a Parallel Navier-Stokes Solver. in: J.J. Dongarra, P.M.A. Sloot and C.J.K. Tan (Eds.), Lecture Notes in Computer Science (LNCS), Vol. 2330, part II, pp. 1019–1028, Springer-Verlag.

Heimbach, P. (2003). Maintaining an up-to-date efficient adjoint for ocean state estimation in the ongoing MIT general circulation model development. In: International Conference on Earth System Modelling, p. 57. Max-Planck-Institute for Meteorology, Hamburg.

Heimbach, P., C. Hill and R. Giering (2005). An efficient exact adjoint of the parallel MIT general circulation model, generated via automatic differentiation. Future Generation Computer Systems, in press.

Jablonowski, C. (2004). Adaptive Grids in Weather and Climate Modeling: PhD Thesis, University of Michigan, Atmospheric and Space Sciences Program.

Jones, H. and J. Marshall (1997). Restratification after deep convection. J. Phys. Oceanogr., 27(10):2276–2287.

Junge, M. and T. Haine (2001). Mechanisms of North Atlantic Wintertime Sea Surface Temperature Anomalies. J. Clim., 14:4560–4572.

Keppenne, C.L., M.M. Rienecker, N.P. Kurkowski and D.D. Adamec (2004). Ensemble Kalman Filter Assimilation of Temperature and Altimeter Data with Bias Correction and Application to Seasonal Prediction, Nonlin. Proc. Geophys., submitted.

Kim, S.-B., T. Lee, and I. Fukumori (2004). The 1997–99 abrupt change of the upper ocean temperature in the northcentral Pacific. Geophys. Res. Lett., in press.

Kirtman, B. P. and S. E. Zebiak (1997). ENSO simulation and prediction with a hybrid coupled model. Mon. Weather Rev., 125(10):2620–2641.

Kistler, R., et al. (2001). The NCEP-NCAR 50-year reanalysis: Monthly means CD-ROM and documentation., Bull. Amer. Meteor. Soc., 82:247–268.

Kleeman, R., A. M. Moore, and N. R. Smith (1995). Assimilation of subsurface thermal data into an intermediate tropical coupled ocean-atmosphere model. Mon. Weather Rev., 123:3103–3113.

Kleeman, R. and A. M. Moore (1997). A theory for the limitation of ENSO predictability due to stochastic atmospheric transients. J. Atmos. Sci., 54:753–767.

Köhl, A. and Willebrand, J. (2002). An adjoint method for the assimilation of statistical characteristics into eddy-resolving ocean models. Tellus, 54(4):406–425.

Köhl, A. and Willebrand, J. (2003). Variational assimilation of SSH variability from TOPEX/POSEIDON and ERS1 into an eddy-permitting model of the North Atlantic. J. Geophys. Res., 108(C3):3092.

Kwok, R. and G. F. Cunningham (2002). Seasonal ice area and volume production of the Arctic Ocean: November 1996 through April 1997. J. Geophys. Res., 107(C10):8038.

Kwok, R., H. J. Zwally, and D. Yi (2004). ICESat observations of Arctic sea ice: A first look. Geophys. Res. Lett., 31:L16401.

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

Latif, M., D. Anderson, T. Barnett, M. Cane, R. Kleeman, A. Leetmaa, J. O'Brien, A. Rosati, and E. Schneider (1998). A review of the predictability and prediction of ENSO. J. Geophys. Res., 103(C7):14,375–14,393.

Lea, D. J., M. R. Allen, and T. W. N. Haine (2000). Sensitivity analysis of the climate of a chaotic system. Tellus, 52A:523–532.

Lea, D., M. Allen, T. Haine, and J. Hansen (2002). Sensitivity analysis of the climate of a chaotic ocean circulation model. Q. J. R. Meteorol. Soc., 128, 2587–2605.

Lee, T., Boulanger, J. P., Foo, A., Fu, L. L., and Giering, R. (2000). Data assimilation by an intermediate coupled ocean-atmosphere model: Application to the 1997–1998 El Nino. J. Geophys. Res., 105(C11):26063–26087.

Lee, T., I. Fukumori, and B. Tang (2004). Temperature advection: internal versus external processes. J. Phys. Oceanogr., 34:1936–1944.

Leith, C. E. (1968). Diffusion approximation for two-dimensional turbulence. Phys. Fluids, 10:1409–1416.

Li, X. and C. Wunsch (2004). An adjoint sensitivity study of chlorofluorocarbons in the North Atlantic, J. Geophys. Res., 109(C1):C01007.

Manabe, S., R. J. Stouffer, M. J. Spellman, and K. Bryan (1991). Transient response of a coupled ocean-atmosphere model to gradual changes of atmospheric CO2, Part I: Annual mean response. J. Clim., 4:785–817.

Marotzke, J. (1992). The role of integration time in determining a steady state through data assimilation. J. Phys. Oceanogr., 22:1556.

Marotzke, J., Giering, R., Zhang, K. Q., Stammer, D., Hill, C., and Lee, T. (1999). Construction of the adjoint mit ocean general circulation model and application to Atlantic heat transport sensitivity. J. Geophys. Res., 104(C12):29529–29547.

Marshall, J., C. Hill, L. Perelman, and A. Adcroft (1997). Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling. J. Geophys. Res., 102(C3):5733–5752.

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

Marshall, J. and T. Radko (2003). Residual mean solutions for the Antarctic Circumpolar Current and its associated overturning circulation. J Phys. Oceanogr., 33(11):2341–2354

McKinley, G., M.J. Follows and J. Marshall (2004). Mechanisms of air-sea CO2 flux variability in the Equatorial Pacific and north Atlantic: Global Biogeochem. Cycles, 18, doi:10.1029/2003GB002179

MacNeice, P., K. M. Olson, C. Mobarry, R. deFainchtein and C. Packer (2001). "PARAMESH : A parallel adaptive mesh refinement community toolkit.", Computer Physics Communications, 126:330–354.

Menemenlis, D., and C.Wunsch (1997). Linearization of an oceanic circulation model for data assimilation and climate studies, J. Atmos. Oceanic Technol., 14(6):1420–1443.

Menemenlis, D., I. Fukumori, and T. Lee (2004a). Using Green's functions to calibrate an ocean general circulation model. Mon. Weather Rev., in press. Available at http://ecco.jpl.nasa.gov/˜dimitri/.

Menemenlis, D., I. Fukumori, A. Köhl, T. Lee, and D. Stammer (2004b). 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. Available at http://www.eccogroup.org/reports.html.

Menke, W., Geophysical Data Analysis: Discrete Inverse Theory, vol 45 of International Geophysics Series. Acadmeic, San Diego, Calif., revised ed., 285 p. 1989

Moore, A. M. (1999). Wind-induced variability of ocean gyres. Dynamics of Atmospheres and Oceans, 29(2–4):335–364..

Moore, A. M. and Kleeman, R. (1999). The nonnormal nature of el nino and intraseasonal variability. Journal of Climate, 12(10):2965–2982.

Neelin, J. D., Battisti, D. S., Hirst, A. C., Jin, F.-F., Wakata, Y., Yamagata, T., and Zebiak, S. (1998). ENSO theory. Special Joint issue of J. Geophys. Res. Atmospheres and J. Geophys. Res. Oceans, 103(C7):14,261–14,290.

Parkinson, C. L., D. J. Cavalieri, P. Gloersen, H. J. Zwally, and J. C. Comiso. Arctic sea ice extents, areas, and trends, 1978–1996. J. Geophys. Res., 104(C9):20837–20856, 1999.

Penland, C. and Sardeshmukh, P. D. (1995). The optimal-growth of tropical sea-surface temperature anomalies. J. Climate, 8(8):1999–2024.

Penven, P., L. Debreu, P. Marchesiello, J.C. McWilliams (2004): Application of the ROMS embedding procedure for the Central California Upwelling System — in preparation

Purser, J. and M. Rancic: 1998, Smooth quasi-homogeneous gridding of the sphere. QJRMS, 124, 637–647.

Roemmich, D., and J. Gilson, 2001: Eddy transport of heat and thermocline waters in the North Pacific: A key to interannual/decadal climate variability? J. Phys. Oceanogr., 31, 675–687.

Schiller, A. and J. Willebrand (1995). A technique for the determination of surface heat and freshwater fluxes from hydrographic observations, using an approximate adjoint ocean circulation model. J. Mar. Res., 53, 433–451.

Schroter, J. (1992). Variational assimilation of GEOSAT data into an eddy-resolving model of the Gulf Stream extension area. J.Phys.Oceanogr., 23:925–953.

Semtner, A. J. Jr. and R. M. Chervin, 1992: Ocean General Circulation from a Global Eddy-Resolving Model, J. Geophys. Res., 97, 5493–5550.

Sirkes, Z. and Tziperman, E. (2001). Identifying a damped oscillatory thermohaline mode in a general circulation model using an adjoint model. J. Phys. Oceanogr., 31:2297–2306.

Stammer, D., and C. Wunsch, The determination of the the largescale circulation of the Pacific Ocean from satellite altimetry using model Green's functions, J. Geophys. Res., 101(C8), 18,409–18,432, 1996.

Stammer, D., C. Wunsch, R. Giering, C. Eckert, P. Heimbach, C. Hill, J. Marotzke, J. Marshall, 2002: The global ocean state during 1992–1997, estimated from ocean observations and a general circulation model. Part I. Methodology and estimated state. J. Geophys. Res., DOI: 10.1029/2001JC000888.

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.

Sun, C., M.M. Rienecker, A. Wittenburt, A. Rosati, M. Harrison, C.L. Keppenne, J. Jacob, A. Borovikov, N. Kurkowski, 2004: Intercomparison of global ocean data assimilation systems in the ODASI experiments (draft ms).

Taft, J.R: Achieving 60 GFLOP/s on the production CFD code OVERFLOW-MLP. Parallel Computing 27(4): 521–536 (2001)

Thacker, W. C. and Long, R. B. (1988). Fitting dynamics to data. J. Geophys. Res., 93:1227–1240.

Thompson, C. J. and Battisti, D. S. (2001). A linear stochastic dynamical model of enso. part ii: Analysis. Journal of Climate, 14(4):445–466.

Troccoli, A., M.M. Rienecker, C.L. Keppenne, and G.C. Johnson, 2003: Temperature Data Assimilation with Salinity Corrections: validations for the NSIPP ocean data assimilation system in the tropical Pacific ocean, 1993–1998, NASA Tech. Memo-2003-104606, Vol.24, 23pp.

Tziperman, E. and Ioannou, P. J. (2002). Transient growth and optimal excitation of thermohaline variability. Journal of Physical Oceanography, 32(12):3427–3435.

Tziperman, E. and Thacker, W. C. (1989). An optimal-control adjoint-equations approach to studying the oceanic general-circulation. J. Phys. Oceanogr., 19(10):1471–1485.

Tziperman, E., W.C. Thacker, R.B. Long and S. Hwang (1992a): Oceanic data analysis using a general circulation model. I: Simulations.J.Phys.Oceanogr., 22:1434–1457.

Tziperman, E., W.C. Thacker, R.B. Long, S. Hwang and S.R. Rintoul (1992b). Oceanic data analysis using a general circulation model. II: A North Atlantic model.J.Phys.Oceanogr., 22:1458–1485.

US GODAE, 2001: Global Ocean Data Assimilation: Prospects and Strategies. A Workshop Report (available from: http://www.usgodae.org/usgodae.html).

Van-oldenborgh, G. J., Burgers, G., Venzke, S., Eckert, C., and Giering, R. (1999). Tracking down the ENSO delayed oscillator with an adjoint ogcm. Mon. Weath. Rev., 127(7):1477–1495.

Visbeck, M., J. Marshall, T. Haine, and M. Spall, (1997) Specification of eddy transfer coefficients in coarse-resolution ocean circulation models, J. Phys. Oceanogr., 27(3), 381–402.

Weaver, A., Vialard, J., Anderson, D., and Delecluse, P. (2002). Three- and fourdimensional variational assimilation with a general circulation model of the tropical pacific ocean. Technical Report 365, ECMWF. 74 pp.

Wunsch, C. (1988). Transient tracers as a problem in control theory. J. Geophys. Res., 93:8099–8119.

Wunsch, C., The Ocean Circulation Inverse Problem. Cambridge Univ. Press, New York, 442 p., 1996.

Zanna, L. and Tziperman, E. (2004). Non normal amplification of the thermohaline circulation. J. Phys. Oceanogr., submitted

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.

Zhang, J. and D. Rothrock. Modeling arctic sea ice with an efficient plastic solution. J. Geophys. Res., 105:3325–3338, 2000.

The NITRD Program: FY 2004 Interagency Coordination Report. Interagency Working Group on Information Technology Research and Development (IWG/ITR&D), September 2004.

7 Management plan

The proposed management of the ECCO II team is set out in Fig.11. Many members of the group have a long track record of working successfully together on various NASA satellite missions and, over the last 5 years, the ECCO project. The JPL team and the ocean modeling group at MIT have worked together for many years: in particular they have worked closely with the NASA Ames group over the past year to target the ECCO/MIT models on to the Project Columbia computer.

Overall responsibility for the project will rest with John Marshall, the head of the Climate Modeling Initiative at MIT, who will also direct the computational strategy. The technical lead will be his long-time collaborator, Chris Hill. Hill is one of the principle authors of the MITgcm, one of three technical leads overseeing the ESMF project and a member of the ACTS project that has been developing the OpenAD system. Lee-Leung Fu at JPL will oversee the global state estimation efforts at JPL. Fu is the project scientist for the Topex/Poseidon and Jason altimiters and a member of the Wide Swath Altimiter design team. He will work with Dimitris Menemenlis as the technical lead. Carl Wunsch, who with Detlef Stammer has directed ECCO efforts over the last 5 years, will take responsibility for adjoint methods and optimization technology with Patrick Heimbach as the technical lead.

The ECCO II group would plan to have meetings twice a year. These would alternate between East coast and West coast locations.

We now describe the connections within each activity.

7.1 Adjoint methods

P. Heimbach will take the lead on implementing and testing the tangent linear (TLM) and adjoints (ADM) for the cubed-sphere setup, the sea-ice model and the high-resolution configuration. Heimbach has many years of experience maintaining the automatic differentiation capabilities of the MITgcm and in maintaining and developing the optimization machinery used in the first ECCO project. He participates in ECCO-GODDAE overseeing the production ECCO runs for that project. He will work in close collaboration with model developers at MIT and D. Menemenlis at JPL. Proven techniques of code sharing through CVS, testing through a well-defined verification suite and model checkpoints will be continued.

P. Heimbach will liaise with J. Utke who will lead OpenAD development, and will provide a relevant test suite of models and benchmarks to ensure prioritization for ECCO-II relevant applications. The major focus during ADM development will be on its application for model/data misfit optimization, i.e. state estimation as outlined in Section 4.2.2.

E. Tziperman will take the lead on the use of the ADM and TLM products for climate sensitivity and predictability applications as outlined in Section 4.3.4. He will work closely with P. Heimbach to make these configurations readily available for the proposed focus studies. Tziperman will also collaborate with M. Rienecker and M. Suarez at GSFC in the context of ENSO prediction.

7.2 Global state estimation

The JPL team is responsible for delivering, maintaining, and distributing results from the two ECCO-II reanalyses, the preliminary 18-km cubed-sphere synthesis and the final, best-possible, eddy-resolving synthesis. In addition, the JPL team is responsible for sea-ice model development, Green-function optimization methodology, and for all the data streams.

D. Menemenlis will be technical lead for global state estimation. He is the developer of the Green function optimization methodology that was used in the first ECCO project and, together

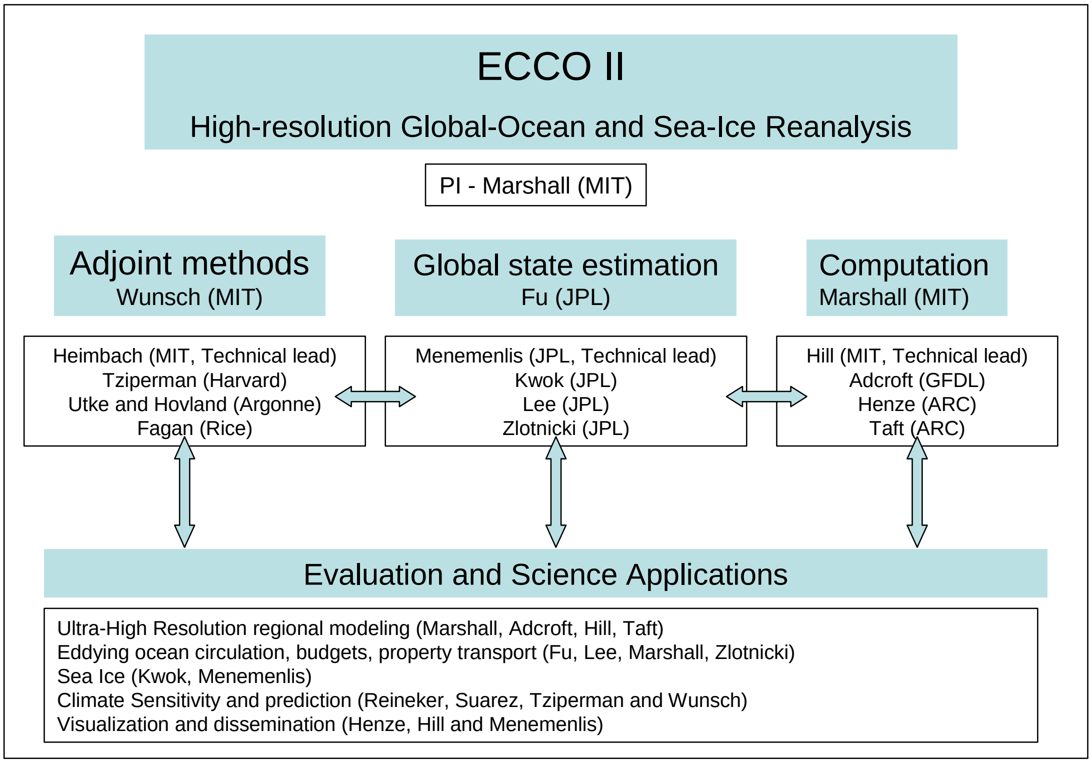

Figure 11: Management plan for the ECCO II project.

with Ron Kwok maintains the sea-ice model that will be used in ECCO. He will liaise with C. Hill and A. Adcroft regarding forward model development, with J. Taft, regarding computational efficiency issues, with P. Heimbach regarding application of approximate-adjoint technology, with C. Henze regarding visualization, with JPL PODAAC regarding data distribution, and with L.-L. Fu, T. Lee, and V. Zlotnicki regarding evaluation and application of results. Strong collaborative links are already in place between the above team members, many established starting more than a decade ago at MIT, others established during the five-year (1998-2003) ECCO project, and finally NAS-ARC contacts established during this past year while carrying out preliminary test integrations on Project Columbia supercomputer, in preparation for this proposal.

7.3 Computation

The computation team will be responsible for:

1. developing optimized code layers that can take full advantage of the Columbia system current and future shared memory architecture
2. developing high-end visualizations of assimilation outputs.
3. developing underlying gridding and numerical strategies.

The key participants — Hill (technical lead), Taft and Henze — have an established collaboration, working closely over the last year with Menemenlis at JPL and with the MITgcm development team at MIT. Jim Taft is the author of the MLP shared memory programming library and an expert on scientific code optimization. Chris Henze is one of the key designers of the Hyperwall system at NASA -AMES. During the last year of working together Hill, Taft and Henze have introduced a shared memory communication layer to boost ECCO assimilation run scalability on the Columbia system and developed visualizations that capture the time-dependent behavior of flow in the ECCO runs. Chris Henze has also advised on the construction of a Hyperwall system at MIT that will be used for visualization of results from this proposal.

Under this proposal the team would further to develop optimized, scalable communication and I/O software (Taft and Hill) to support scaling to thousands of processors and develop visualization techniques (Henze and Hill) that can be used to explore temporal and spatial correlations within multi-variate assimilation solutions containing many terabytes of information.

John Marshall and Chris Hill will work closely with their long-time collaborator Alistair Adcroft in the development of gridding algortihms and their parallelization that make the cubed sphere calculations — and their refinements and elaborations — possible.

Hill will oversee and coordinate effort (together with the NASA ESMF team in which he participates as a technical lead and as a member of the ESMF Joint Specification Team) to ensure that all modeling tools emerging from this proposal are compatible with ESMF and with the MAP modeling environment that the NRA envisages.

7.4 Evaluation and Science Applications

Members from all three organizational groupings will participate in product evaluation and science applications activities. Results from these activities will provide feedback that will guide the efforts of the adjoint methods, global state estimation and computation teams. The activities will include (i) work by researchers with John Marshall to diagnose eddy fluxes, Reynolds stresses, applying these results to the development of theoretical understanding and parameterizations, (ii) work by Tony Lee at JPL to diagnose and monitor heat and water mass fluxes, building on substantial work in this area in the prior ECCO round (iii) work by Ron Kwok, an expert on sea-ice processes and remote sensing at JPL, to evaluate high-resolution Arctic and Anarctic sea-ice estimates, (iv) work by Victor Zlotnicki to assess bottom pressure estimates and transports in the Antarctic Circumpolar and North Pacific regions, (v) work by Michelle Rienecker and Max Suarez to develop and apply new error covariance estimates to the GMAO prediction systems, (vi) work by Eli Tziperman and collaborators to apply ECCO-II results to ENSO prediction and predictability scenarios, (vii) work by Carl Wunsch and collaborators to look at property transports and at global ocean energy budgets. These activities will be coordinated through the twice yearly ECCO-II group meetings. The meetings will include focussed reviews of the evaluation and science results, providing crucial, rapid feedback to the ongoing production and research assimilation activities.

8 Budget summary and timeline

Figure 12 shows the year by year budget breakdown for the project grouped by partner institution. Key areas of responsibility associated with the investigator teams at the different institutions are also shown. It is planned to have a first product available within 6 months of project start (preliminary work on the Columbia systems has already been carried out) so that all groups can begin contributing right from the project start. In addition to soliciting community feedback on the ECCO-II products, several of the ECCO-II participants plan to have active science analyses using the products. Responsibility for "results analysis" therefore lies with many groups.

Institution	Investigators	Key Responsibilites	Yr. 1	Yr. 2	Yr. 3	Yr. 4	Yr. 5	Total
MIT	Marshall Wunsch Hill Heimbach	Prognostic and adjoint model. Advanced non-linear optimization techniques. ESMF integration. Results analysis.	780K	759K	790K	823K	856K	4,008K
HARVARD	Tziperman	Application of ECCO-II to ENSO and climate prediction. Results analysis.	201K	210K	217K	226K	234K	1,088K
PRINCETON	Adcroft	Gridding and numerics.	90K	90K	95K	99K	103K	477K
CHICAGO	Hovland Utke	OpenAD tool.	113K	128K	125K	128K	131K	625K
JPL	Fu Menemenlis Kwok Lee Zlotnicki	Operational assimilation. Sea-ice model. Data and visualization product distribution. Results analysis.	899K	976K	998K	1012K	1104K	4,989K
NASA/AMES	Henze Taft	Project Columbia code optimization. Advanced visualization generation.	100K	102K	104K	107K	109K	522K
NASA/GSFC	Rienecker Suarez	Application of error covariances at GMAO. Results analysis.	104K	113K	117K	121K	125K	580K
Total			2,287K	2,378K	2,446K	2,516K	2,662K	12,289K

Figure 12:

8.1 Project timeline

Figure 13 shows the planned project timeline. Throughout the project a key series of datasets (shown in the first year entries of the figure) will form the basis of the estimation product. These will be continually updated and, as model and assimilation fidelity increase during the project, revised. The software used in this work will continue to be available and updated through repositories maintained at M.I.T and Argonne.

Time from start (years)	Products and key activities.
0.5	First reanalysis 1992-present - including raw velocity, temperature, salinity, sea-surface height (SSH) bottom pressure (BP), vertical mixing, sea-ice thickness, area and velocity datasets available (full spatial resolution, time averaged at 3-days with hourly data for SSH and BP). Downscaled versions available (averaged in space and time at monthly and annual periods). Animations of fields available. Team meeting to review and plan for derived products. Raw products will be updated on an ongoing basis from here on. Supporting software will be updated on an ongoing basis.
1	First derived products available - including climatologies and high frequency time-series of heat and fresh water transports (including ice-export), overturning, surface fluxes (including ice-ocean), eddy fluxes, eddy mixing coefficients, residual circulation, energy budgets, mixed-layer depth and water-mass transformation rates, error covariance estimates. Derived product time series animations available. Team meeting including planning first adjoint synthesis and regional embedded assimilation work. Derived products will be updated on an ongoing basis from here on.
2	Useful preliminary adjoint synthesis available with fields as above plus time series of adjoint sensitivities including heat-transport, overturning, deep-ocean heat content, passive tracer distributions, ice-volume, residual model data misfits. Animations of time dependent adjoint sensitivities available. Adjoint products will be updated on an ongoing basis from here on. Updated OpenAD tool available. Updated Green function analysis evaluation. Raw fields (see above) for embedded regional estimates available. Team meeting to determine precise "best-possible" estimate configuration.
3	Green function raw and derived products for "best-possible" configuration available. Derived products (see above) for embedded regional estimates available. Team meeting to refine "best-possible" configuration and to plan for global ~1km experiments.
4	First generation adjoint "best-possible" raw and derived products available. First global ~1km raw and derived products available. Final OpenAD tool released.
5	Final raw and derived product cycles for "best-possible" and regional embedded estimation available. Final global ~1km raw and derived products available.

Figure 13: Project timeline