On the formulation of sea-ice models. Part 2: Lessons from multi-year adjoint sea ice export sensitivities through the Canadian Arctic Archipelago.

Patick Heimbach [a,1], Dimitris Menemenlis [b], Martin Losch [c],

Jean-Michel Campin [a] and Chris Hill [a]

[a] *Department of Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA*

[b] *Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA*

[c] *Alfred-Wegener-Institut für Polar- und Meeresforschung, Postfach 120161, 27515 Bremerhaven, Germany*

Abstract

The adjoint of an ocean general circulation model is at the heart of the ocean state estimation system of the *Estimating the Circulation and Climate of the Ocean* (ECCO) project. As part of an ongoing effort to extend ECCO to a coupled ocean/sea-ice estimation system, a dynamic and thermodynamic sea-ice model has been developed for the Massachusetts Institute of Technology general circulation model (MITgcm). One key requirement is the ability to generate, by means of automatic differentiation (AD), tangent linear (TLM) and adjoint (ADM) model code for the coupled MITgcm ocean/sea-ice system. This second part of a two-part paper describes aspects of the adjoint model. The adjoint ocean and sea ice model is used to calculate transient sensitivities of solid (ice & snow) freshwater export through Lancaster Sound in the Canadian Arctic Archipelago (CAA). The adjoint state provides a complementary view of the dynamics. In particular, the transient, multi-year sensitivity patterns reflect dominant pathways and propagation timescales through the CAA as resolved by the model, thus shedding light on causal relationships, in the model, across the Archipelago. The computational cost of inferring such causal relationships from forward model diagnostics alone would be prohibitive. The role of the exact model trajectory around which the adjoint is calculated (and therefore of the exactness of the adjoint) is exposed through calculations using free-slip vs no-slip lateral boundary conditions. Effective ice thickness, sea surface temperature, and precipitation sensitivities, are discussed in detail as examples of the coupled sea-ice/ocean and atmospheric forcing control space. To test the reliability of the adjoint, finite-difference perturbation experiments were performed for each of these elements and the cost perturbations were compared to those "predicted" by the adjoint. Overall, remarkable qualitative and quantitative agreement is found. In particular, the adjoint correctly "predicts" a seasonal sign change in precipitation sensitivities. A physical mechanism for this sign change is presented. The availability of the coupled adjoint opens up the prospect for adjoint-based coupled ocean/sea-ice state estimation.

Key words: NUMERICAL SEA ICE MODELING, VISCOUS-PLASTIC RHEOLOGY, COUPLED OCEAN AND SEA ICE MODEL, STATE ESTIMATION, ADJOINT MODELING, CANADIAN ARCTIC ARCHIPELAGO, SEA-ICE EXPORT, SENSITIVITIES

[1] corresponding author, email: heimbach@mit.edu,

ph: +1-617-253-5259, fax: +1-617-253-4464

1 Introduction

This is the second part of a two-part paper (see Losch et al., 2010, for part 1) describing the development of a sea-ice model for use in adjoint-based regional and global coupled ocean/sea-ice state estimation and sensitivity studies. It has been shown (e.g., Marotzke et al., 1999, Galanti et al., 2002, Galanti and Tziperman, 2003, Köhl, 2005, Bugnion et al., 2006a,b, Losch and Heimbach, 2007, Moore et al., 2009, Veneziani et al., 2009) that adjoints are very valuable research tools to investigate sensitivities of key model diagnostics with respect to a wide variety of model inputs. Furthermore, increasing sophistication of global-scale as well as regional, polar state estimation systems, which attempt to synthesize observations and models (e.g., Miller et al., 2006, Duliere and Fichefet, 2007, Lisaæter et al., 2007, Stark et al., 2008, Stoessel, 2008, Panteleev et al., 2010) call for adequate representation of sea-ice in the model so as to represent relevant processes and to incorporate sea-ice observations in constraining the coupled system. The estimation system developed within the *Estimating the Circulation and Climate of the Ocean* (ECCO) consortium is based on the adjoint or Lagrange multiplier method (LMM) (e.g., Wunsch, 2006). It thus relies heavily on the availability of an adjoint model of the underlying general circulation model (Stammer et al., 2002a, Wunsch and Heimbach, 2007, Heimbach and Wunsch, 2007, and references therein).

Collectively, the lack, until recently, of an interactive sea-ice component in the ECCO approach, the experience gained (and the success) with the ocean-only problem, the importance of representing polar-subpolar interactions in ECCO-type calculations, and the need to incorporate sea-ice observations, make a compelling case for the development of a new sea-ice model. While many of its features are “conventional” (yet for the most part state-of-the-art), the ability to generate efficient adjoint code for coupled ocean/sea-ice simulations by means of automatic (or algorithmic) differentiation (AD: Griewank and Walther, 2008) sets this model apart from existing models. Whereas a few existing models (Kim et al., 2006a,b) allow for the generation of tangent linear code for sea-ice-only model configurations by means of the so-called *forward-mode* AD, until very recently none of these were capable of producing efficient adjoint code by means of *reverse-mode* AD, let alone in a coupled ocean/sea-ice configuration, which can propagate sensitivities back and forth between the two components. Such coupled sensitivity propagation is highly desirable as it permits sea-ice *and* ocean observations to be used as simultaneous constraints on each other, yielding a truly coupled estimation problem.

In addition to the coupled ocean and sea ice system described here, one other coupled adjoint system has recently become available for an Arctic configuration and was used to isolate dominant mechanisms responsible for the 2007 Arctic sea-ice minimum (Kauker et al., 2009). The availability of two adjoint modeling systems holds the prospect (for the first time) to compare *adjoint* calculations for a specific regional setup using different models. This is a proposed future objective within the Arctic Ocean Model Intercomparison Project (AOMIP).

The MITgcm sea ice model was described in detail in Part 1. It borrows many components from current-generation sea ice models, but these components were reformulated on an Arakawa C grid in order to match the MITgcm oceanic grid, and they were modified in many ways to permit efficient and accurate automatic differentiation. Part 1 provided a detailed discussion of the effect on the solution of various choices in the numerical implementation, in particular related to sea-ice dynamics. Such sensitivities are structural or configuration-based, rather than exploring a continuous space of control variables, and are best assessed in separate forward calculations. Special emphasis was put on aspects of the sea-ice dynamics, such as the use of different solvers for sea-ice rheology, the formulation of these solvers on an Arakawa B vs C grid, and the use of free-slip vs no-slip lateral boundary conditions. These scenarios provide important baseline trajectories for the adjoint calculations presented here, as they underscore the importance of the underlying state, around which the model is linearized.

Part 2 focusses on the adjoint component, its generation by means of AD, its reliability, and on the interpretability of adjoint variables. We investigate sensitivities of sea-ice transport through narrow straits, for which rheology configurations become crucial, and the dependence of adjoint sensitivities on the choices of configuration elements described in Part 1. The power of the adjoint is demonstrated through a case study of sea-ice transport through the Canadian Arctic Archipelago (CAA) measured in terms of its export through Lancaster Sound. Thereby we complement a recent study by Lietaer et al. (2008) that focused on the role of narrow straits in this region in setting the sea-ice mass balance in the Arctic. While Part 1 of the present paper showed that different grids, different rheologies, and different lateral boundary conditions lead to considerable differences in the computed sea-ice state, here we show that adjoint sensitivities may differ substantially depending on the baseline trajectory, around which the model is linearized. The present analysis provides important complementary information to the configuration sensitivities of Part 1: it enables us to extend analysis to continuous parameters, it demonstrates the degree of detail the adjoint variables contain, and it exposes causal relationships.

The remainder of Part 2 is organized as follows: Section 2 provides some details of the adjoint code generation by means of AD. Multi-year transient sensitivities of sea-ice export through the Canadian Arctic Archipelago are presented in Section 3. Extending the analysis of Part 1, we assess the consequences of the choices of lateral boundary conditions on the ensuing model sensitivities for various control variables. Discussion and conclusions are in Section 4.

2 MITgcm adjoint code generation

There is now a growing body of literature on adjoint applications in oceanography and adjoint code generation via AD. We therefore limit the description of the method to a brief summary. For discrete problems as considered here, the adjoint model operator (ADM) is the transpose of the Jacobian or tangent linear model operator (TLM) of the full (in general nonlinear) forward model (NLM), in this case, the MITgcm coupled ocean and sea ice model. Consider a scalar-valued model diagnostics, referred to as objective function, and an m-dimensional control space (referred to as space of independent variables) whose elements we may wish to perturb to assess their impact on the objective function. In the context of data assimilation the objective function may be the least-square model vs. data misfit, whereas here, we may choose almost any function that is (at least piece-wise) differentiable with respect to the control variables. Here, we shall be focusing on the solid freshwater export through Lancaster Sound.

Two- and three-dimensional control variables used in the present study are

Table 1

List of control variables used. The controls are either part of the oceanic (O) or sea-ice (I) state, or time-varying elements of the atmospheric (A) boundary conditions.

component	variable	dim.	time
O	temperature	3-D	init.
O	salinity	3-D	init.
O	vertical diffusivity	3-D	const.
I	concentration	2-D	init.
I	thickness	2-D	init.
A	air temperature	2-D	2-day
A	specific humidity	2-D	2-day
A	shortwave radiation	2-D	2-day
A	precipitation	2-D	2-day
A	zonal windspeed	2-D	2-day
A	merid. windspeed	2-D	2-day

listed in Table 1. They consist of two- or three-dimensional fields of initial conditions of the ocean or sea-ice state, ocean vertical mixing coefficients, and time-varying surface boundary conditions (surface air temperature, specific humidity, shortwave radiation, precipitation, zonal and meridional wind speed). The TLM computes the objective functions's directional derivatives for a given perturbation direction. In contrast, the ADM computes the the full gradient of the objective function with respect to all control variables. When combined, the control variables may span a potentially high-dimensional, e.g., O(10^8), control space. At this problem dimension, perturbing individual parameters to assess model sensitivities is prohibitive. By contrast, transient sensitivities of the objective function to any element of the control and model state space can be computed very efficiently in one single adjoint model integration, provided an adjoint model is available.

Conventionally, adjoint models are developed "by hand" through implementing code which solves the adjoint equations (e.g., Marchuk, 1995, Wunsch, 1996) of the given forward equations. The burden of developing "by hand" an adjoint model in general matches that of the forward model development. The substantial extra investment often prevents serious attempts at making available adjoint components of sophisticated models. Furthermore, the work of keeping the adjoint model up-to-date with its forward parent model matches the work of forward model development. The alternative route of rigorous application of AD tools has proven very successful in the context of MITgcm ocean modeling applications.

Certain limitations regarding coding standards apply. Although they vary from tool to tool, they are similar across various tools and are related to the ability to efficiently reverse the flow through the model. Work is thus required initially to make the model amenable to efficient adjoint code generation for a given AD tool. This part of the adjoint code generation is not automatic (we sometimes refer to it as semi-automatic) and can be substantial for legacy code, in particular if the code is badly modularized and contains many irreducible control flows (e.g., GO TO statements, which are considered bad coding practice anyways).

It is important to note, nevertheless, that once the tailoring of the model code to the AD code is in place, any further forward model development can be easily incorporated in the adjoint model via AD. Furthermore, the notion of *the adjoint* is misleading, since the structure of the adjoint depends critically on the control problem posed (a passive tracer sensitivity yields a very different Jacobian to an active tracer sensitivity). A clear example of the dependence of the structure of the adjoint model on the control problem is the extension of the MITgcm adjoint model to a configuration that uses bottom topography as a control variable (Losch and Heimbach, 2007). The AD approach enables a much more thorough and smoother adjoint model extension than would be possible via hand-coding.

The adjoint model of the MITgcm has become an invaluable tool for sensitivity analysis as well as for state estimation (for a recent overview and summary, see Heimbach, 2008). AD also enables a large variety of configurations and studies to be conducted with adjoint methods without the onerous task of modifying the adjoint of each new configuration by hand. Giering and Kaminski (1998) discuss in detail the advantages of AD.

The AD route was also taken in developing and adapting the sea-ice component of the MITgcm, so that tangent linear and adjoint components can be obtained and kept up to date without excessive effort. As for the TLM and ADM components of the MITgcm ocean model, we rely on the AD tool "Transformation of Algorithms in Fortran" (TAF) developed by Fastopt (Giering and Kaminski, 1998) to generate TLM and ADM code of the MITgcm sea ice model (for details see Marotzke et al., 1999, Heimbach et al., 2005). Note that for the ocean component, we are now also able to generate efficient derivative code using the new open-source tool OpenAD (Utke et al., 2008). Appendix

A provides details of adjoint code generation for the coupled ocean and sea ice MITgcm configuration.

Since conducting this study, further changes to the thermodynamic formulation have been implemented, which improve certain aspects of forward and adjoint model behavior. These changes are discussed in detail in Fenty (2010) along with application of the coupled ocean and sea ice MITgcm adjoint to estimating the state of the Labrador Sea during 1996–1997.

To conclude this section, we emphasize the coupled nature of the MITgcm ocean and sea ice adjoint. Figure 1 illustrates the relationship between control variables and the objective function J when using the tangent linear model (TLM, left diagram), or the adjoint model (ADM, right diagram). The control space consists of atmospheric perturbations (e.g., surface air temperature δT_a and precipitation δp), sea-ice perturbations (e.g., ice concentration δc and ice thickness δh), and oceanic perturbations (e.g., potential temperature $\delta\Theta$ and salinity δS). The left diagram depicts how each perturbation of an element of the control space leads to a perturbed objective function δJ via the TLM. In contrast, the right diagram shows the reverse propagation of *adjoint variables* or *sensitivities* labeled with an asterisk (*). The notation reflects the fact that adjoint variables are formally Lagrange multipliers or elements of the model's co-tangent space (as opposed to perturbations which are formally elements of the model's tangent space). For example, $\delta^* c$ refers to the gradient $\partial J/\partial c$. The aim of the diagram is to show (in a very simplified way) two things. First, it depicts how sensitivities of an objective function (e.g., sea ice export as will be defined later) to changes in, e.g., ice concentration $\partial J/\partial c$ is affected by changes in, e.g., ocean temperature via the chain rule $\partial J/\partial\Theta = \partial J/\partial c \cdot \partial c/\partial\Theta$. The adjoint model thus maps the adjoint objective function state to the adjoint sea-ice state, and from there to the coupled adjoint oceanic and surface atmospheric state. Second, it can be seen that the ADM maps from a 1-dimensional state ($\delta^* J$) to a multi-dimensional state ($\delta^* c, \delta^* h, \delta^* T_a, \delta^* p, \delta^* \Theta, \delta^* S$) whereas the TLM maps from a multi-dimensional state ($\delta c, \delta h, \delta T_a, \delta p, \delta \Theta, \delta S$) to a 1-dimensional state (δJ). This is the reason why only one adjoint integration is needed to assemble all the gradients of the objective function while one tangent linear integrations per dimension of the control space is needed to assemble the same gradient. Rigorous derivations can be found in, for example, Chapter 5 of the MITgcm documentation (Adcroft et al., 2002), in Wunsch (2006), or in Giering and Kaminski (1998).

3 A case study: Sensitivities of sea-ice export through Lancaster Sound

We demonstrate the power of the adjoint method in the context of investigating sea-ice export sensitivities through Lancaster Sound (LS). The rationale for this choice is to complement the analysis of sea-ice dynamics in the presence of narrow straits of Part 1. LS is one of the main paths of sea ice export through the Canadian Arctic Archipelago (CAA) (Melling, 2002, Prinsenberg and Hamilton, 2005, Michel et al., 2006, Münchow et al., 2006, Kwok, 2006). Figure 2 shows the intricate local geography of CAA straits, sounds, and islands. Export sensitivities reflect dominant pathways through the CAA, as resolved by the model. Sensitivity maps provide a very detailed view of various quantities affecting the sea-ice export (and thus the underlying propagation pathways). A caveat of this study is the limited resolution, which is not adequate to realistically simulate the CAA. For example, while the dominant

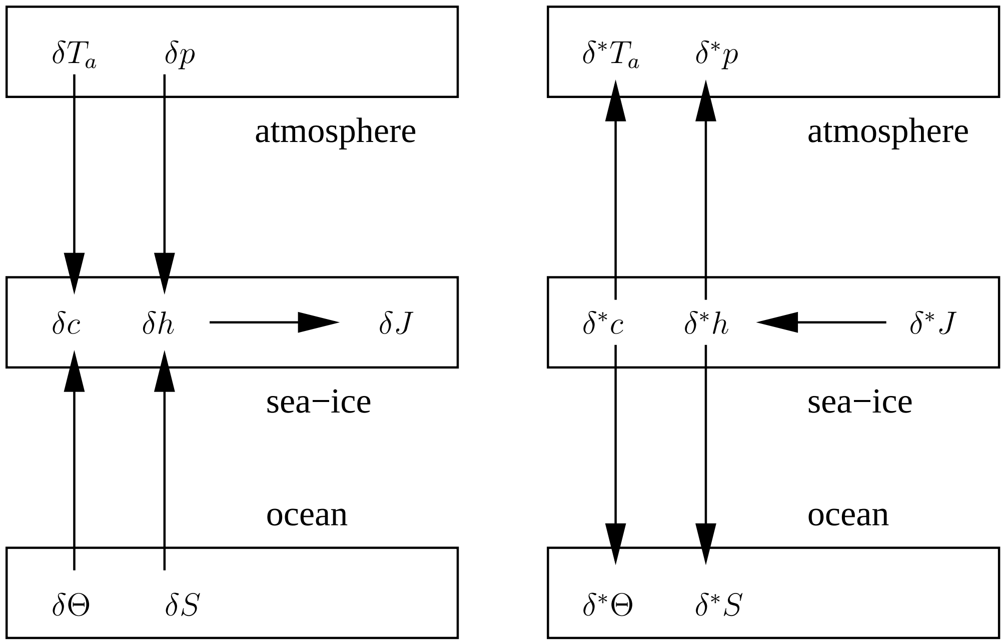

Fig. 1. This diagram illustrates how the tangent linear model (TLM, left panel) maps perturbations in the oceanic, atmospheric, or sea-ice state into a perturbation of the objective function δJ, whereas the adjoint model (ADM, right panel) maps the adjoint objective function $\delta^* J$ (seeded to unity) into the adjoint sea-ice state, which is a sensitivity or gradient, e.g., $\delta^* c = \partial J/\partial c$, and into the coupled ocean and atmospheric adjoint states. The TLM computes how a perturbation in *one* input affects *all* outputs whereas the adjoint model computes how *one* particular output is affected by *all* inputs.

circulation through LS is toward the East, there is a small Westward flow to the North, hugging the coast of Devon Island, which is not resolved in our simulation. Nevertheless, the focus here is on elucidating model sensitivities in a general way. For any given simulation, whether deemed "realistic" or not, the adjoint provides exact model sensitivities, which help inform whether hypothesized processes are actually borne out by the model dynamics. Note that the resolution used in this study is at least as good as or better than the resolution used for IPCC-type calculations.

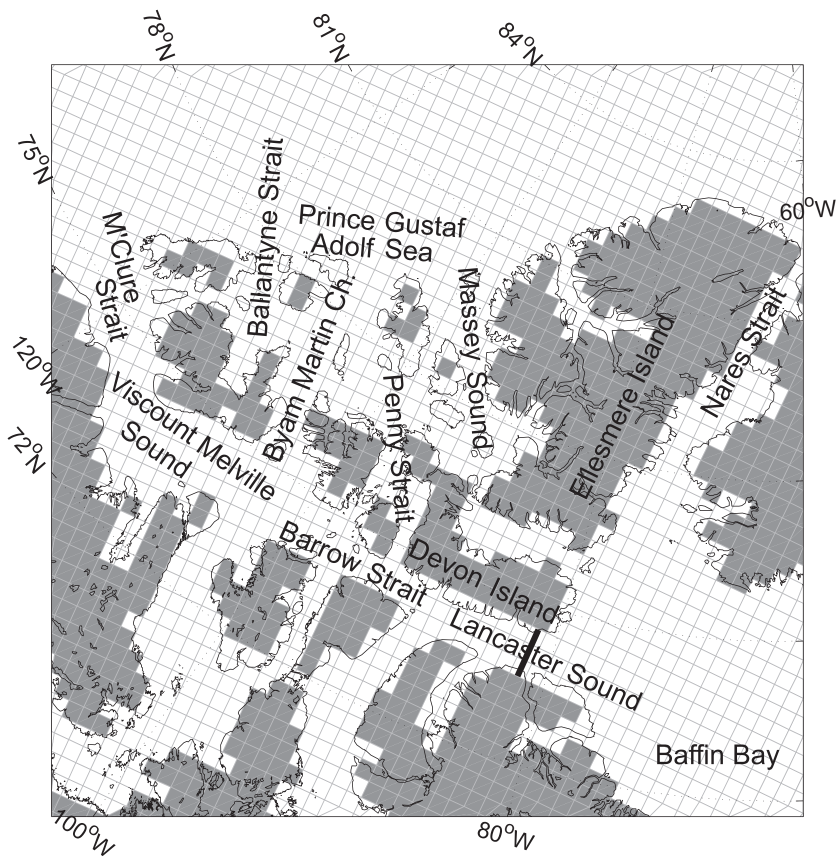

Fig. 2. Map of the Canadian Arctic Archipelago with model coastlines and grid (filled grey boxes are land). The black contours are the true coastlines as taken from the GSHHS data base (Wessel and Smith, 1996). The gate at $82°W$ across which the solid freshwater export is computed is indicated as black line.

3.1 The model configuration

The model domain is similar to the one described in Part 1. It is carved out from the Arctic face of a global, eddy-admitting, cubed-sphere simulation (Menemenlis et al., 2005) but with 36-km instead of 18-km grid cell width, i.e., coarsened horizontal resolution compared to the configuration described in Part 1. The vertical discretization is the same as in Part 1, i.e. the model has 50 vertical depth levels, which are unevenly spaced, ranging from 10 m layer thicknesses in the top 100 m to a maximum of 456 m layer thickness at depth. The adjoint model for this configuration runs efficiently on 80 processors, inferred from benchmarks on both an SGI Altix and on an IBM SP5 at NASA/ARC and at NCAR/CSL, respectively. Following a 4-year spinup (1985 to 1988), the model is integrated for an additional four years and nine months between January 1, 1989 and September 30, 1993. It is forced at the surface using realistic 6-hourly NCEP/NCAR atmospheric state variables. The objective function J is chosen as the "solid" freshwater export through LS, at approximately 74° N, 82° W in Fig. 2, integrated over the final 12-month period, i.e., October 1, 1992 to September 30, 1993. That is,

$$J = \frac{1}{\rho_{fresh}} \int_{\text{Oct 92}}^{\text{Sep 93}} \int_{\text{LS}} (\rho\, h\, c \,+\, \rho_s h_s c)\, u\, ds\, dt, \tag{1}$$

is the mass export of ice and snow converted to units of freshwater. Furthermore, for each grid cell (i,j) of the section, along which the integral $\int \ldots ds$ is taken, $c(i,j)$ is the fractional ice cover, $u(i,j)$ is the along-channel ice drift velocity, $h(i,j)$ and $h_s(i,j)$ are the ice and snow thicknesses, and ρ, ρ_s, and ρ_{fresh} are the ice, snow and freshwater densities, respectively. At the given resolution, the section amounts to three grid points. The forward trajectory of the model integration resembles broadly that of the model in Part 1 but some details are different due to the different resolution and integration period. For example, the differences in annual solid freshwater export through LS as defined in eqn. (1) are smaller between no-slip and free-slip lateral boundary conditions at higher resolution, as shown in Part 1, Section 4.3 ($91\pm85\,\text{km}^3\,\text{y}^{-1}$ and $77 \pm 110\,\text{km}^3\,\text{y}^{-1}$ for free-slip and no-slip, respectively, and for the C-grid LSR solver; $\pm$ values refer to standard deviations of the annual mean) than at lower resolution ($116 \pm 101\,\text{km}^3\,\text{y}^{-1}$ and $39 \pm 64\,\text{km}^3\,\text{y}^{-1}$ for free-slip and no-slip, respectively). The large range of these estimates emphasizes the need to better understand the model sensitivities to lateral boundary conditions and to different configuration details. We aim to explore these sensitivities across the entire model state space in a comprehensive manner by means of the adjoint model.

The adjoint model is the transpose of the tangent linear model operator. It thus runs backwards in time from September 1993 to January 1989. During this integration period, the Lagrange multipliers of the model subject to objective function (1) are accumulated. These Langrange multipliers are the sensitivities, or derivatives, of the objective function with respect to each control variable and to each element of the intermediate coupled ocean and sea ice model state variables. Thus, all sensitivity elements of the model state and of the surface atmospheric state are available for analysis of the transient sensitivity behavior. Over the open ocean, the adjoint of the Large and Yeager (2004) bulk formula scheme computes sensitivities to the time-varying atmospheric state. Specifically, ocean sensitivities propagate to air-sea flux sensitivities, which are mapped to atmospheric state sensitivities via the bulk formula adjoint. Similarly, over ice-covered areas, the sea-ice model adjoint (rather than the bulk formula adjoint) converts surface ocean sensitivities to atmospheric sensitivities.

3.2 Adjoint sensitivities

The most readily interpretable ice-export sensitivity is that to ice thickness, $\partial J/\partial(hc)$. Maps of transient sensitivities $\partial J/\partial(hc)$ are shown for free-slip (Fig. 3) and for no-slip (Fig. 4) boundary conditions. Each figure depicts four

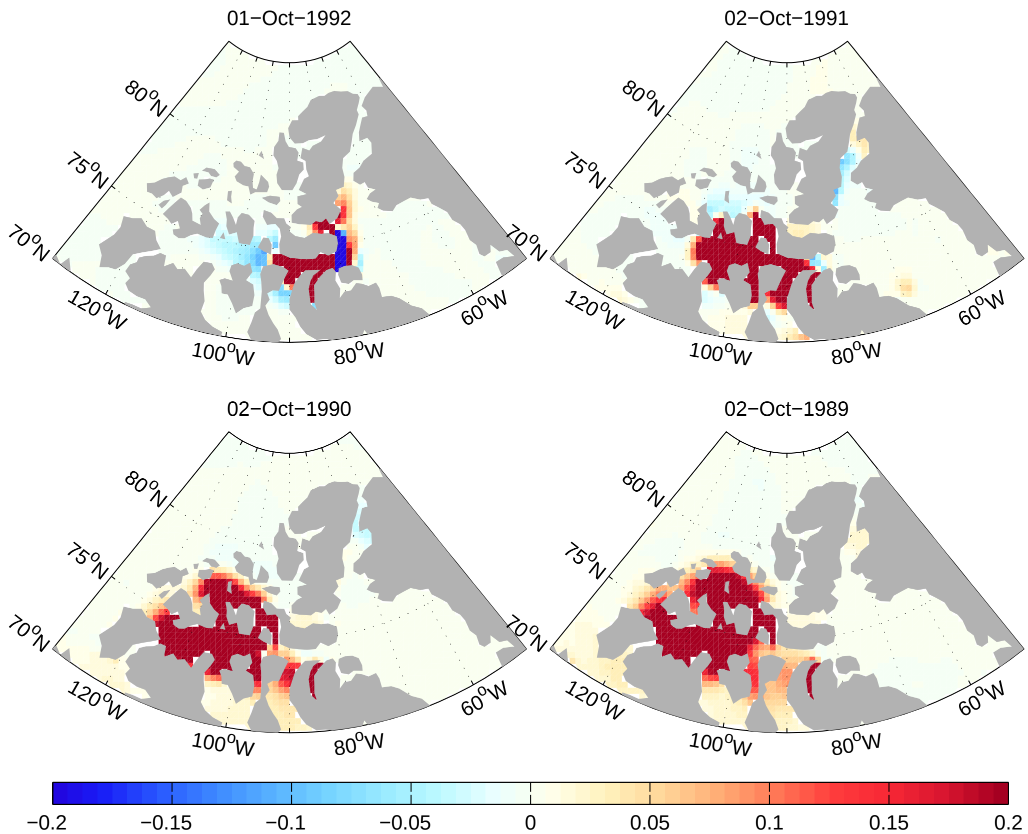

Fig. 3. Sensitivity $\partial J/\partial(hc)$ in $\mathrm{m}^3\,\mathrm{s}^{-1}/\mathrm{m}$ for four different times using free-slip lateral sea ice boundary conditions. The color scale is chosen to illustrate the patterns of the sensitivities. The objective function (1) was evaluated between October 1992 and September 1993. Sensitivity patterns extend backward in time upstream of the LS section.

sensitivity snapshots of the objective function J, starting October 1, 1992, i.e., at the beginning of the 12-month averaging period, and going back in time to October 2, 1989. As a reminder, the full period over which the adjoint sensitivities are calculated is (backward in time) between September 30, 1993 and January 1, 1989.

The sensitivity patterns for ice thickness are predominantly positive. The interpretation is that an increase in ice volume in most places west, i.e., "upstream", of LS increases the solid freshwater export at the exit section. The

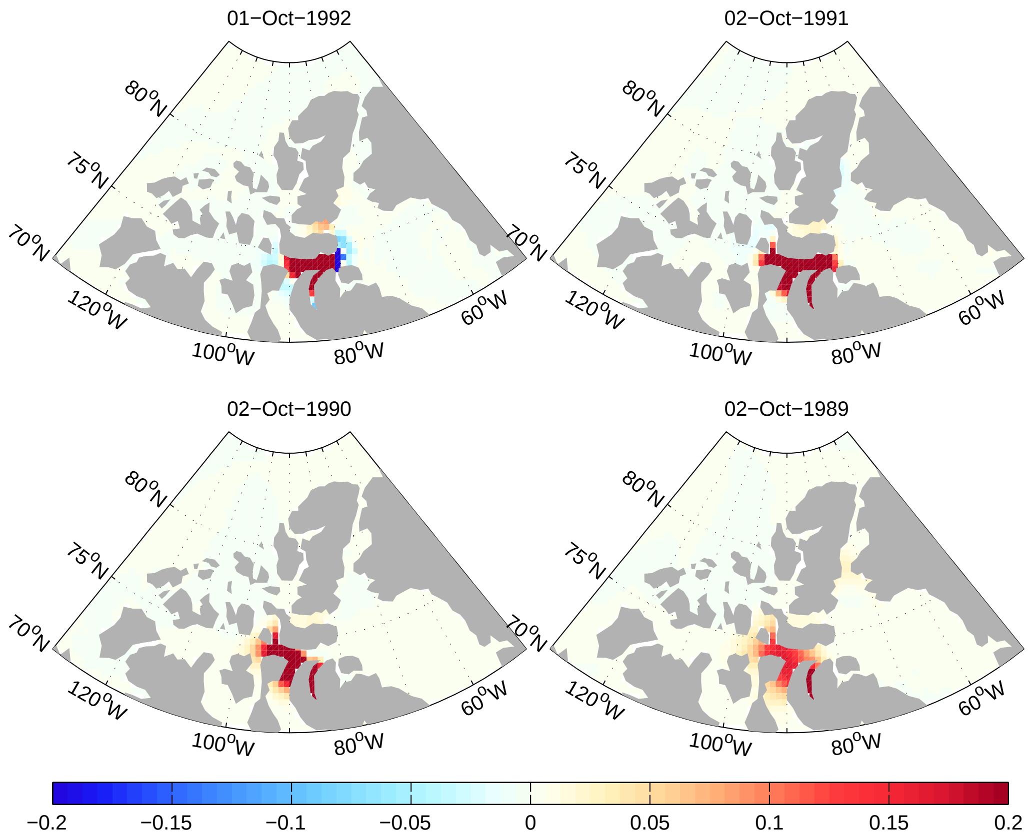

Fig. 4. Same as in Fig. 3 but for no-slip lateral sea ice boundary conditions.

transient nature of the sensitivity patterns is evident: the area upstream of LS that contributes to the export sensitivity is larger in the earlier snapshot. In the free-slip case, the sensivity follows (backwards in time) the dominant pathway through Barrow Strait into Viscount Melville Sound, and from there trough M'Clure Strait into the Arctic Ocean [2]. Secondary paths are northward from Viscount Melville Sound through Byam Martin Channel into Prince Gustav Adolf Sea and through Penny Strait into MacLean Strait.

There are large differences between the free-slip and no-slip solutions. By the end of the adjoint integration in January 1989, the no-slip sensitivities

[2] (the branch of the "Northwest Passage" apparently discovered by Robert McClure during his 1850 to 1854 expedition; McClure lost his vessel in the Viscount Melville Sound)

(Fig. 4) are generally weaker than the free slip sensitivities and hardly reach beyond the western end of Barrow Strait. In contrast, the free-slip sensitivities (Fig. 3) extend through most of the CAA and into the Arctic interior, both to the West (M'Clure Strait) and to the North (Ballantyne Strait, Prince Gustav Adolf Sea, Massey Sound). In this case the ice can drift more easily through narrow straits and a positive ice volume anomaly anywhere upstream in the CAA increases ice export through LS within the simulated 4-year period.

One peculiar feature in the October 1992 sensitivity maps are the negative sensivities to the East and, albeit much weaker, to the West of LS. The former can be explained by indirect effects: less ice eastward of LS results in less resistance to eastward drift and thus more export. A similar mechanism might account for the latter, albeit more speculative: less ice to the West means that more ice can be moved eastward from Barrow Strait into LS leading to more ice export.

The temporal evolution of several ice export sensitivities along a zonal axis through LS, Barrow Strait, and Melville Sound (115° W to 80° W, averaged across the passages) are depicted in Fig. 5 as Hovmoeller-type diagrams, that is, as two-dimensional plots of sensitivities as a function of longitude and time. Serving as examples for the ocean, sea-ice, and atmospheric forcing components of the model, we depict, from top to bottom, the sensitivities to ice thickness (hc), to ice and ocean surface temperature (SST), and to precipitation (p) for free-slip (left column) and for no-slip (right column) ice drift boundary conditions. The green line marks the starting time (1 Oct. 1992) of the 12-month ice export objective function integration (Eqn. 1). Also indicated are times when a perturbation in precipitation leads to a positive (Apr. 1991) or to a negative (Nov. 1991) ice export anomaly (see also Fig.

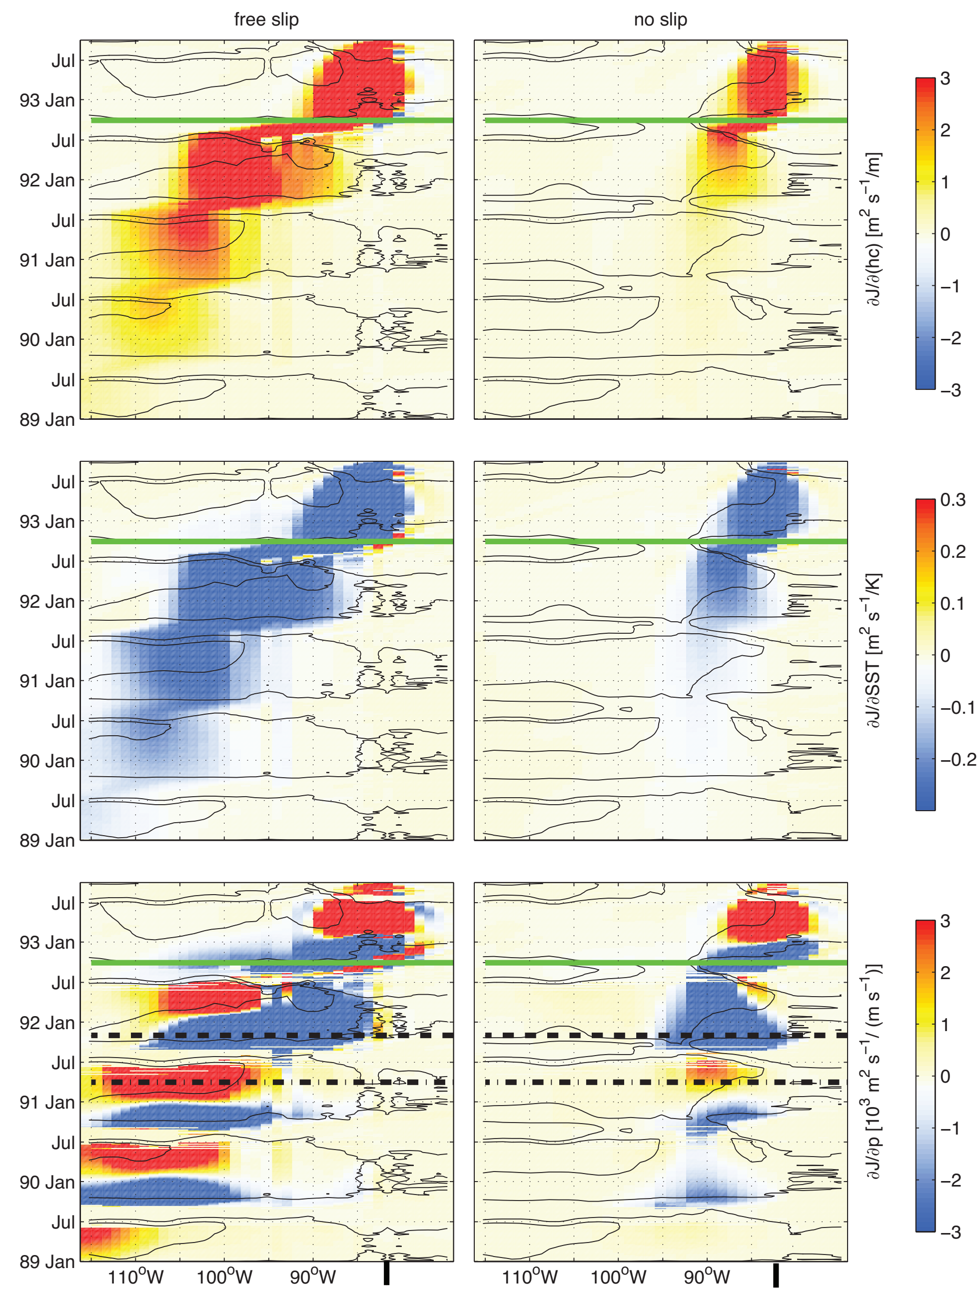

Fig. 5. Time vs. longitude diagrams along the axis of Viscount Melville Sound, Barrow Strait, and LS. The diagrams show the sensitivities (derivatives) of the solid freshwater export J through LS (Fig. 2) with respect to ice thickness (hc, top), to ice and ocean surface temperature (SST, middle), and to precipitation (p, bottom) for free-slip (left) and for no-slip (right) boundary conditions. J was integrated over the last year (period above green line). A precipitation perturbation during Apr. 1st. 1991 (dash-dottel line) or Nov. 1st 1991 (dashed line) leads to a positive or negative export anomaly, respectively. Contours are of the normalized ice strength P/P^{*}. Bars in the longitude axis indicates the flux gate at 82°W.

8). Each plot is overlaid with contours 1 and 3 of the normalized ice strength $P/P^* = (hc)\exp[-C\,(1-c)]$.

The Hovmoeller-type diagrams of ice thickness (top row) and SST (second row) sensitivities are coherent: more ice in LS leads to more export and one way to form more ice is by colder surface temperatures. In the free-slip case the sensitivities spread out in "pulses" following a seasonal cycle: ice can propagate eastward (forward in time) and thus sensitivities propagate westward (backwards in time) when the ice strength is low in late summer to early autumn (Fig. 6, bottom panels). In contrast, during winter, the sensitivities show little to no westward propagation as the ice is frozen solid and does not move. In the no-slip case the normalized ice strength does not fall below 1 during the winters of 1991 to 1993 (mainly because the ice concentrations remain near 100%, not shown). Ice is therefore blocked and cannot drift eastwards (forward in time) through the Viscount Melville Sound, Barrow Strait, and LS channel system. Consequently, the sensitivities do not propagate westward (backwards in time) and the export through LS is only affected by local ice formation and melting for the entire integration period.

It is worth contrasting the sensitivity diagrams of Fig. 5 with the Hovmoeller-type diagrams of the corresponding state variables (Figs. 6 and 7). The sensitivities show clear causal connections of ice motion over the years, that is, they expose the winter arrest and the summer evolution of the ice. These causal connections cannot easily be inferred from the Hovmoeller-type diagrams of ice and snow thickness. This example illustrates the usefulness and complementary nature of the adjoint variables for investigating dynamical linkages in the ocean/sea-ice system.

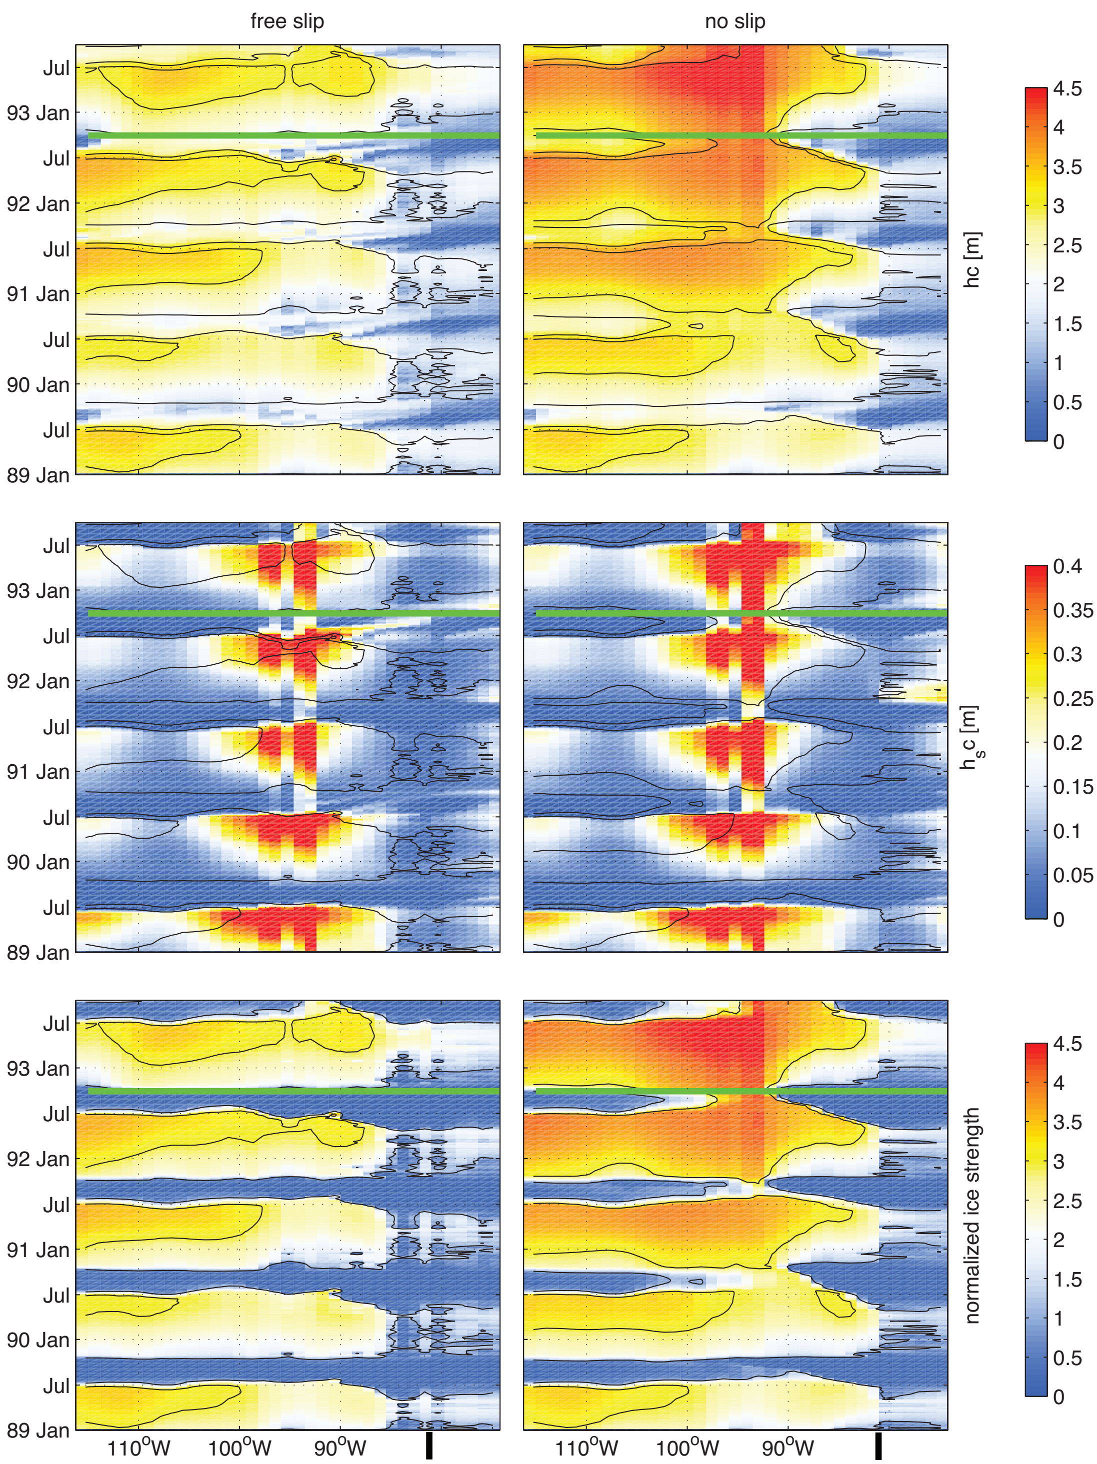

Fig. 6. Hovmoeller-type diagrams along the axis of Viscount Melville Sound, Barrow Strait, and LS. The diagrams show ice thickness (hc, top), snow thickness ($h_s c$, middle), and normalized ice strength (P/P^*, bottom) for free-slip (left) and for no-slip (right) sea ice boundary conditions. For orientation, each plot is overlaid with contours 1 and 3 of the normalized ice strength. Green line is as in Fig. 5.

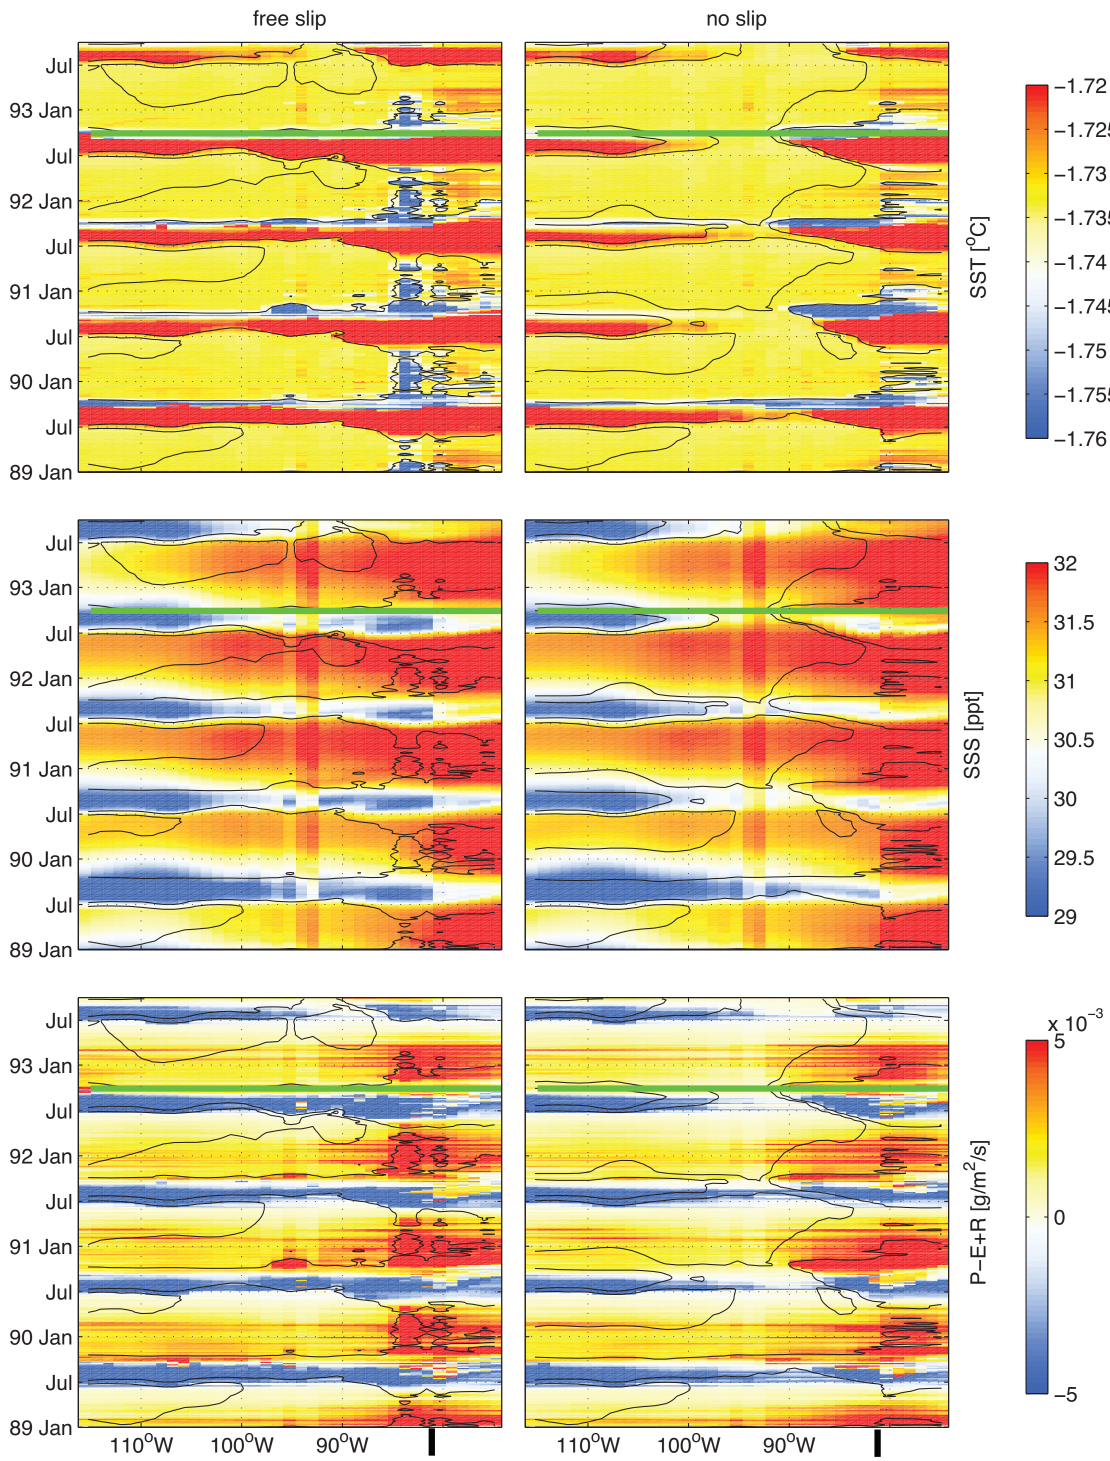

Fig. 7. Same as in Fig. 6 but for SST (top panels), SSS (middle panels), and precipitation minus evaporation plus runoff, $P - E + R$ (bottom panels).

The sensitivities to precipitation are more complex. To first order, they have an oscillatory pattern with negative sensitivity (more precipitation leads to less export) between roughly September and December and mostly positive sensitivity from January through June (sensitivities are negligible during the summer). Times of positive sensitivities coincide with times of normalized ice strengths exceeding values of 3. This pattern is broken only immediatly preceding the evaluation period of the ice export objective function in 1992. In contrast to previous years, the sensitivity is negative between January and August 1992 and east of 95° W.

We attempt to elucidate the mechanisms underlying these precipitation sensitivities in Section 3.4 in the context of forward perturbation experiments.

3.3 Forward perturbation experiments

Applying an automatically generated adjoint model under potentially highly nonlinear conditions incites the question to what extent the adjoint sensitivities are "reliable" in the sense of accurately representing forward model sensitivities. Adjoint sensitivities that are physically interpretable provide a partial answer but an independent, quantitative test is needed to gain confidence in the calculations. Such a verification can be achieved by comparing adjoint-derived gradients with ones obtained from finite-difference perturbation experiments. Specifically, for a control variable $\mathbf{u}$ of interest, we can readily calculate an expected change δJ in the objective function for an applied perturbation $\delta\mathbf{u}$ over domain A based on adjoint sensitivities $\partial J/\partial\mathbf{u}$:

$$\delta J = \int_A \frac{\partial J}{\partial \mathbf{u}} \, \delta \mathbf{u} \, dA \tag{2}$$

Alternatively, we can infer the magnitude of the objective perturbation δJ without use of the adjoint. Instead we apply the same perturbation $\delta \mathbf{u}$ to the control space over the same domain A and integrate the forward model. The perturbed objective function is

$$\delta J = J(\mathbf{u} + \delta \mathbf{u}) - J(\mathbf{u}). \tag{3}$$

The degree to which Eqns. (2) and (3) agree depends both on the magnitude of perturbation $\delta \mathbf{u}$ and on the length of the integration period.

We distinguish two types of adjoint-model tests. First there are finite difference tests performed over short time intervals, over which the assumption of linearity is expected to hold, and where individual elements of the control vector are perturbed. We refer to these tests as gradient checks. Gradient checks are performed on a routine, automated basis for various MITgcm verification setups, including verification setups that exercise coupled ocean and sea ice model configurations. These automated tests insure that updates to the MITgcm repository do not break the differentiability of the code.

A second type of adjoint-model tests is finite difference tests performed over longer time intervals and where a whole area is perturbed, guided by the adjoint sensitivity maps, in order to investigate physical mechanisms. The examples discussed herein and summarized in Table 2 are of this second type of sensitivity experiments. For nonlinear models, the deviations between Eqns. (2) and (3) are expected to increase both with perturbation magnitude as well as

Table 2

Summary of forward perturbation experiments and comparison of adjoint-based and finite-difference-based objective function sensitivities. All perturbations were applied to a region centered at 101.24°W, 75.76°N. The reference value for ice and snow export through LS is $J_0 = 69.6$ km$^3/yr$. For perturbations to the time-varying precipitation p the perturbation interval is indicated by Δt.

exp.	variable	time	Δt	$\delta \mathbf{u}$	$\frac{\delta J(adj.)}{km^3/yr}$	$\frac{\delta J(fwd.)}{km^3/yr}$	% diff.
ICE1	hc	1-Jan-89	init.	0.5 m	0.98	1.1	11
OCE1	SST	1-Jan-89	init.	0.5°C	-0.125	-0.108	16
ATM1	p	1-Apr-91	10 dy	$1.6 \cdot 10^{-7}$ m/s	0.185	0.191	3
ATM2	p	1-Nov-91	10 dy	$1.6 \cdot 10^{-7}$ m/s	-0.435	-1.016	57
ATM3	p	1-Apr-91	10 dy	$-1.6 \cdot 10^{-7}$ m/s	-0.185	-0.071	62
ATM4	p	1-Nov-91	10 dy	$-1.6 \cdot 10^{-7}$ m/s	0.435	0.259	40

with integration time.

Comparison between finite-difference and adjoint-derived ice-export perturbations show remarkable agreement for initial value perturbations of ice thickness (ICE1) or sea surface temperature (OCE1). Deviations between perturbed objective function values remain below 16% (see Table 2). Figure 8 depicts the temporal evolution of perturbed minus unperturbed monthly ice export through LS for initial ice thickness (top panel) and SST (middle panel) perturbations. In both cases, differences are confined to the melting season, during which the ice unlocks and which can lead to significant export. Large differences are seen during (but are not confined to) the period during which the ice export objective function J is integrated (grey box). As "predicted" by

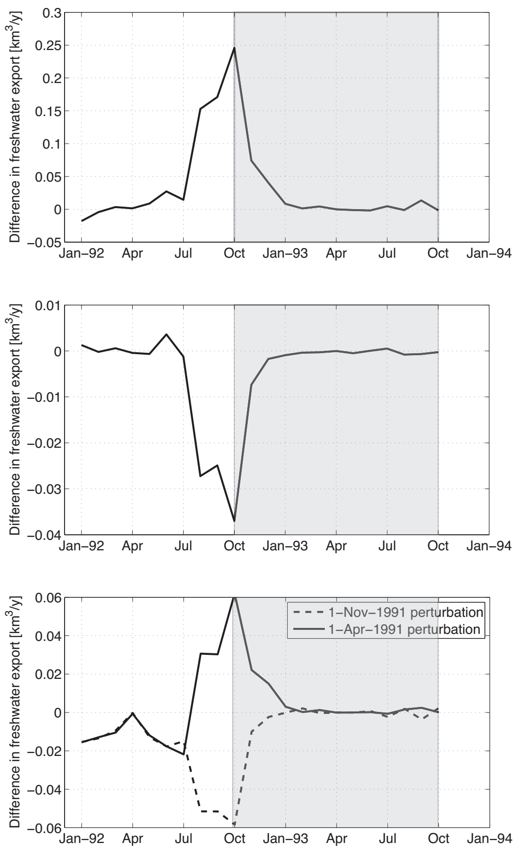

Fig. 8. Difference in monthly solid freshwater export at 82°W between perturbed and unperturbed forward integrations. From top to bottom, perturbations are initial ice thickness (ICE1 in Table 2), initial sea-surface temperature (OCE1), and precipitation (ATM1 and ATM2). The grey box indicates the period during which the ice export objective function J is integrated, and reflects the integrated anomalies in Table 2.

the adjoint, the two curves are of opposite sign and scales differ by almost an order of magnitude.

3.4 Sign change of precipitation sensitivities

Our next goal is to explain the sign and magnitude changes through time of the transient precipitation sensitivities. To investigate this, we have carried out the following two perturbation experiments: (i) an experiment labeled ATM1, in which we perturb precipitation over a 10-day period between April 1 and 10, 1991, coincident with a period of positive adjoint sensitivities, and (ii) an experiment labeled ATM2, in which we apply the same perturbation over a 10-day period between November 1 and 10, 1991, coincident with a period of negative adjoint sensitivities. The perturbation magnitude chosen is $\delta\mathbf{u} = 1.6 \times 10^{-7}$ m/s, which is of comparable magnitude with the standard deviation of precipitation. The perturbation experiments confirm the sign change when perturbing in different seasons. We observe good quantitative agreement for the April 1991 case and a 50% deviation for the November 1991 case. The discrepancy between the finite-difference and adjoint-based sensitivity estimates results from model nonlinearities and from the multi-year integration period. To support this statement, we repeated perturbation experiments ATM1 and ATM2 but applied a perturbation with opposite sign, i.e., $\delta\mathbf{u} = -1.6 \times 10^{-7}$ m/s (experiments ATM3 and ATM4 in Table 2). For negative $\delta\mathbf{u}$, both perturbation periods lead to about 50% discrepancies between finite-difference and adjoint-derived ice export sensitivities. The finite-difference export changes are different in amplitude for positive and for negative perturbations, confirming that model nonlinearities start to impact these calculations.

These experiments constitute severe tests of the adjoint model in the sense that they push the limit of the linearity assumption. Nevertheless, the results confirm that adjoint sensitivities provide useful qualitative, and, within certain limits, quantitative information of comprehensive model sensitivities that cannot realistically be computed otherwise.

To investigate in more detail the oscillatory behavior of precipitation sen-

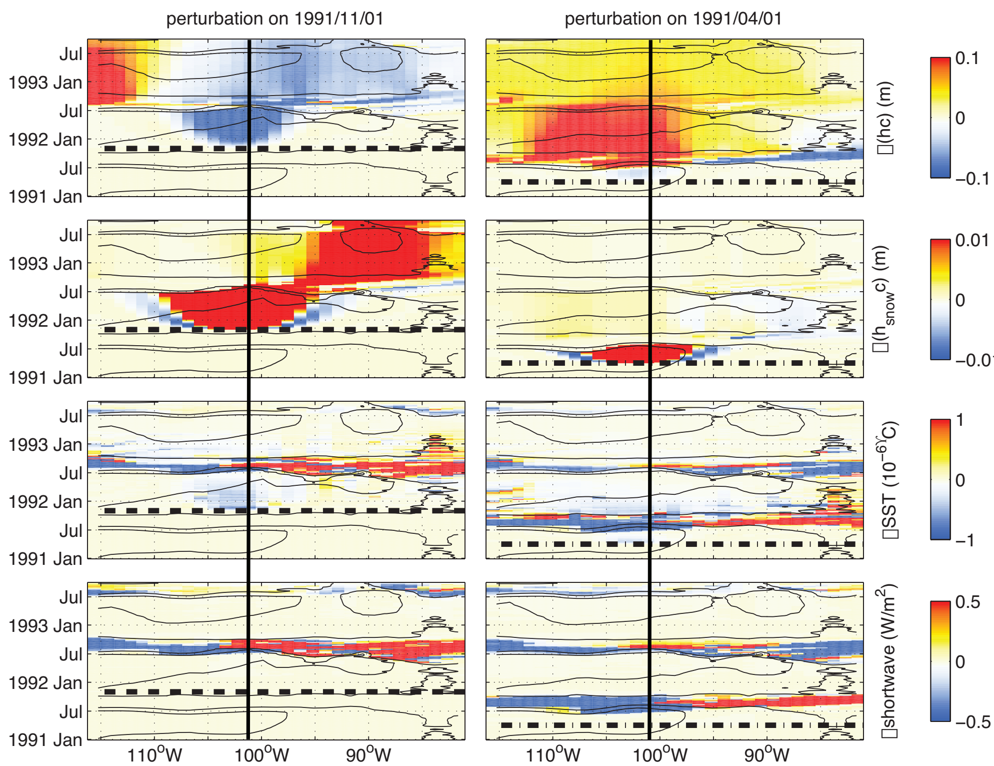

Fig. 9. Same as in Fig. 6 but restricted to the period 1991–1993 and for the differences in (from top to bottom) ice thickness (hc), snow thickness ($h_{\mathrm{snow}}c$), sea–surface temperature (SST), and shortwave radiation (for completeness) between a perturbed and unperturbed run in precipitation of $1.6 \times 10^{-1}\,\mathrm{m\,s^{-1}}$ on November 1, 1991 (left panels) and on April 1, 1991 (right panels). The vertical line marks the position where the perturbation was applied.

sitivities we have plotted differences in ice thickness, snow thicknesses, and SST, between perturbed and unperturbed simulations along the LS axis as a function of time. Figure 9 shows how the small localized perturbations of precipitation are propagated, depending on whether applied during *early* winter (November, left column) or *late* winter (April, right column). More precipation leads to more snow on the ice in all cases. However, the same perturbation in different seasons has an opposite effect on the solid freshwater export through LS. Both the adjoint and the perturbation results suggest the following mechanism to be at play:

- More snow in November (on thin ice) insulates the ice by reducing the effective conductivity and thus the heat flux through the ice. This insulating effect slows down the cooling of the surface water underneath the ice. In summary, more snow early in the winter limits the ice growth from above and below (negative sensitivity).
- More snow in April (on thick ice) insulates the ice against melting. Shortwave radiation cannot penetrate the snow cover and snow has a higher albedo than ice (0.85 for dry snow and 0.75 for dry ice in our simulations); thus it protects the ice against melting in the spring, more specifically, after January, and it may lead to more ice in the following growing season.

A secondary effect is the accumulation of snow, which increases the exported volume. The feedback from SST appears to be negligible because there is little connection of anomalies beyond a full seasonal cycle.

We note that the effect of snow vs rain seems to be irrelevant in explaining positive vs negative sensitivity patterns. In the current implementation, the model differentiates between snow and rain depending on the thermodynamic growth rate of sea ice; when it is cold enough for ice to grow, all precipitation is assumed to be snow. The surface atmospheric conditions most of the year in the Lancaster Sound region are such that almost all precipitation is treated as snow, except for a short period in July and August; even then, air temperatures are only slightly above freezing.

Finally, the negative sensitivities to precipitation between 95° W and 85° W during the spring of 1992, which break the oscillatory pattern, may also be explained by the presence of snow: in an area of large snow accumulation (almost 50 cm: see Fig. 6, middle panel), ice cannot melt and it tends to block the channel so that ice coming from the West cannot pass, thus leading to less ice export in the next season. The reason why this is true for the spring of 1992 but not for the spring of 1991 is that by then the high sensitivites have propagated westward out of the area of thick snow and ice around 90° W.

4 Discussion and conclusion

In this study we have extended the MITgcm adjoint modeling capabilities to a coupled ocean and sea-ice configuration. The key development is a dynamic and thermodynamic sea-ice model akin to most state-of-the-art models but that is amenable to efficient, exact, parallel adjoint code generation via automatic differentiation. At least two natural lines of applications are made possible by the availability of the adjoint model: (i) use of the coupled adjoint modeling capabilities for comprehensive sensitivity calculations of the ocean/sea-ice system at high Northern and Southern latitudes and (ii) extension of the ECCO state estimation infrastructure to derive estimates that are constrained both by ocean and by sea-ice observations.

The power of the adjoint method was demonstrated through a multi-year sensitivity calculation of solid freshwater (sea-ice and snow) export through Lancaster Sound in the Canadian Arctic Archipelago (CAA). The region was chosen so as to complement the forward-model study presented in Part 1, which examined the impact of rheology and dynamics on sea-ice drift through narrow straits. The transient adjoint sensitivities reveal dominant pathways of sea-ice propagation through the CAA. They clearly expose causal, time-lagged relationships between ice export and various ocean, sea-ice, and atmospheric variables of the coupled system. The computational cost of establishing all these relationships through pure forward calculations would be prohibitive. The sensitivity patterns (and thus causal relationships) differ substantially, depending on which lateral ice drift boundary condition (free-slip or no-slip) is imposed. Our results indicate that for the coarse-resolution configuration used here the free-slip boundary condition results in swifter ice movement and in a much larger region of influence than does the no-slip boundary condition. Note though that this statement may not hold for simulations at higher resolution.

The present calculations confirm some expected responses, for example, the increase in ice export with increasing ice thickness and the decrease in ice export with increasing sea surface temperature. They also reveal mechanisms which, although plausible, cannot be readily anticipated. As an example we presented precipitation sensitivities, which exhibit an annual oscillatory behavior, with negative sensitivities prevailing throughout the fall and early winter and positive sensitivities from late winter though spring. This behavior can be traced to the different impact of snow accumulation over ice, depending on the stage of ice evolution. For growing ice, snow accumulation suppresses ice growth (negative sensitivity) whereas for melting ice, snow accumulation suppresses ice melt (positive sensitivity). A secondary effect is the snow accumulation on downstream ice export (positive sensitivity). Differences between snow and rain seem negligible in our case study, since precipitation is in the form of snow for an overwhelming part of the year.

Given the automated nature of adjoint code generation and the nonlinearity of the problem when considered over sufficiently long time scales, independent tests are needed to gain confidence in the adjoint solutions. We have presented such tests in the form of finite difference experiments, guided by the adjoint solution, and we compared objective function differences inferred from forward perturbation experiments with differences inferred from adjoint sensitivity information. We found very good quantitative agreement for initial ice thickness and for sea surface temperature perturbations.

As described above, sensitivities to precipitation show an annual oscillatory behavior, which is confirmed by forward perturbation experiments. In terms of amplitude, precipitation shows a larger deviation (order of 50%) between adjoint-based and finite-difference-based estimates of ice and snow transport sensitivity through Lancaster Sound. Furthermore, finite difference perturbations exhibit an asymmetry between positive and negative perturbations of equal size. This points to the fact that, on multi-year time scales, nonlinear effects can no longer be ignored and it indicates a limit to the usefulness of the adjoint sensitivity information.

Given the urgency of understanding cryospheric changes, adjoint applications are emerging as powerful research tools, e.g., the study of Kauker et al. (2009) who attempt to isolate dominant mechanisms responsible for the 2007 Arctic sea-ice minimum, and the study of Heimbach and Bugnion (2009) who demonstrate how to infer Greenland ice sheet volume sensitivities from a large-scale ice sheet adjoint model. The results of the present study encourage application of the MITgcm coupled ocean/sea-ice adjoint system to a variety of sensitivity studies of Arctic and Southern Ocean climate variability. The system has matured to a stage where coupled ocean/sea-ice estimation becomes feasible. For the limited domain of the the Labrador Sea, single-year estimates have indeed successfully been produced by Fenty (2010) for the mid-1990s and mid-2000s, and will be reported elsewhere. Steps both toward a full Arctic and a global system are now within reach. The prospect of using observations of one component (e.g., daily sea-ice concentration) to constrain the other component (near-surface ocean properties) through the information propagation of the adjoint holds promise in deriving better, dynamically consistent estimates of the polar environments.

A Issues of AD-based adjoint code generation

TAF (Giering and Kaminski, 1998) and OpenAD (Utke et al., 2008) are source-to-source transformation tools, which take the Fortran source code of the nonlinear parent model (NLM) and generate Fortran code for the derivative model once the control space and objective function have been specified. The specification is an important step. It determines, in part, the structure of the TLM and ADM. For different control problems the TLM and ADM may be different, underlining the advantage of AD over hand-coding. At a basic level, the AD tool knows the derivative expression for all intrinsic Fortran functions (`+,-,*,/,SQRT,SIN,` etc.) and it readily produces line-by-line tangent linear code. The full tangent linear model is assembled by rigorous application of the chain rule (and the product rule) to the derivative line expressions. The adjoint code can be derived from the line-by-line TLM code, formulated in matrix form, by taking the matrix transpose and putting the resulting equations in code form.

An example Consider as a simple example the line of code for calculating the nonlinear bulk viscosity ζ from the shear viscosity η and from the ratio e

of the major to minor axis of the elliptical yield curve (Hibler, 1979):

$$\eta = \frac{\zeta}{e^2}. \tag{A.1}$$

The total derivative is

$$\begin{aligned} \delta\eta &= \frac{\partial\eta}{\partial\zeta}\,\delta\zeta + \frac{\partial\eta}{\partial e}\,\delta e \\ &= \frac{1}{e^2}\,\delta\zeta - \frac{2\zeta}{e^3}\,\delta e. \end{aligned} \tag{A.2}$$

The variables $\delta\eta$, δe, and $\delta\zeta$ are perturbations to the NLM state variables and may be viewed as elements of the TLM state space. Rewriting this in matrix form,

$$\begin{bmatrix} \delta\zeta \\ \delta e \\ \delta\eta \end{bmatrix}^{l+1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{1}{e^2} & \frac{-2\zeta}{e^3} & 0 \end{bmatrix} \begin{bmatrix} \delta\zeta \\ \delta e \\ \delta\eta \end{bmatrix}^{l}, \tag{A.3}$$

enables easy access to the transpose

$$\begin{bmatrix} \delta^*\zeta \\ \delta^* e \\ \delta^*\eta \end{bmatrix}^{l} = \begin{bmatrix} 1 & 0 & \frac{1}{e^2} \\ 0 & 1 & \frac{-2\zeta}{e^3} \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \delta^*\zeta \\ \delta^* e \\ \delta^*\eta \end{bmatrix}^{l+1}, \tag{A.4}$$

where $\delta^*\eta$, δ^*e, and $\delta^*\zeta$ are sensitivities, i.e., elements of the ADM state space or elements of the dual space to the TLM space. From this the adjoint code can easily be read-off as follows:

$$\begin{aligned} \delta^*\zeta &= \delta^*\zeta + \frac{1}{e^2}\delta^*\eta, \\ \delta^* e &= \delta^* e - \frac{2\zeta}{e^3}\delta^*\eta, \\ \delta^*\eta &= 0. \end{aligned} \tag{A.5}$$

Note that:

- the TLM propagates the impact of perturbing one input component ($\delta\eta$) on all output variables (a directional derivative), here just one scalar-valued objective function,

- the ADM accumulates the sensitivities of one output variable (here scalar-valued) to all input components (a gradient),

- the required variables are elements of the model state, which are needed to evaluate the derivative expression, including nonlinear functions and conditional statements, and for the ADM they need to be available in reverse order,

- Eqn. (A.5) states that the shear viscosity sensitivity $\delta^*\eta$ impacts the bulk viscosity sensitivity $\delta^*\zeta$ in a linear fashion, whereas it affects the ratio of the elliptic yield curve δ^*e nonlinearly.

Required variables and checkpointing An important issue is the evaluation of nonlinear or conditional expressions. In Eqn. (A.5) the values of e and ζ are required to evaluate the derivative. AD tools solve this problem for TLM generation by interlacing the TLM calculation with the NLM calculation. In this way, the state of e and ζ is known just when it is needed by the TLM. For the ADM the solution is significantly harder since the state of e and ζ are required in reverse order of the NLM execution. Overcoming this discrepancy is at the heart of implementing efficient adjoint code. The approach taken is a blend of two extremes, which are (i) recomputing the required state, or (ii) storing the whole state. For complex models, such as the MITgcm, neither of these in their pure form is feasible but an optimal blend, known as adjoint multi-level checkpointing, enables the generation of efficient and exact adjoint code. For TAF, which implements a recompute-all behavior as default, the task consists of targeting active variables in relevant, e.g., nonlinear or conditional, code expressions, whose storing will avoid excessive required recomputations. TAF directives enable the modeler to support TAF, alter its default behavior, and render the adjoint more efficient. A detailed description in the context of the MITgcm is given in Heimbach et al. (2005). Alternative approaches of store-all by default are implemented in other tools (e.g., OpenAD, see Utke et al., 2008).

Hand-coded adjoint models are sometimes considered as more efficient and faster in view of the ability of the code developer to explicitly optimize the code. This view needs to be formulated in more detail since it may be misleading in its general form. Significant code optimization can be obtained through relaxing the requirement of provision of the exact model forward state at the time of derivative evaluation. A code developer may decide that certain variables vary sufficiently slowly such that a time-mean (or, in certain applications, an equilibrium state) constitute an appropriate substitute. While this substitution leads indeed to significant adjoint model speed-up and/or memory reduction (omission of required recomputations) the comparison in performance is no longer warranted. This is because similar interventions are possible for AD generated code, in which recomputation or STORE/RESTORE operations may very well be replaced by similar approximations after the adjoint code has been generated. Code efficiently is thus not primarily an AD issue, but an issue of deciding which approximations to the exact linearizations are permissible. These decisions are either made at the outset (for hand-coding), or after the fact (for AD). Which of the routes of either simplifying an AD-generated adjoint or extending an approximate hand-coded adjoint is simpler and leads to more efficient adjoint models remains subject to research. Clearly, providing means (e.g. through directives) of prescribing approximation levels to AD tools would be an attractive feature of AD tools, and very useful for large-scale applications.

Retaining scalability of the coupled ocean/sea-ice adjoint Another aspect is ensuring scalability of the adjoint code on high performance computer systems. Here again, automatic differentiation provides adjoint code, which implements the same domain decomposition strategy adopted in the forward model. It thus inherits the same parallel modeling approach, and therefore essentially the same scalable code efficiency as the parent model. In terms of across-processor operations, such as exchanging information between processor tiles, global sums, etc., the same set of adjoint primitives can be used that have been developed for the MITgcm ocean component (Heimbach et al., 2005).

The main parallel operations are exchanges between processors (send/receive, gather/scatter), as well as global sums (reduce). All of these are linear operations in nature. Therefore there are no fundamental hurdles to parallel adjoint model execution. Adjoint primitives of the parallel support package have been written by hand since no adjoint support of the Message Passing Interface (MPI) is currently available (Heimbach et al., 2005). Nevertheless, efforts are currently under way to extend MPI libraries to include support for adjoint model generation (Utke et al., 2009).

Iterative solvers and their adjoint Next, we briefly describe the treatment of the sea-ice rheology solver. The solver used here is an adaptation of the line successive over-relaxation (LSOR) method of Zhang and Hibler III (1997) to an Arakawa C grid (see Part 1). At the heart of this method is an iterative approach used to solve the momentum equations for ice drift velocities, based on a tridiagonal matrix solver. A challenge is to generate the adjoint of the iterative procedure. A similar issue was encountered in the context of adding bottom topography as a control variable to the MITgcm, which breaks the self-adjoint property of the elliptic pressure solver and which required adjoint code generation for this routine (Losch and Heimbach, 2007). The approach taken here consists of invoking the implicit function theorem in order to simplify the reverse accumulation of sensitivities in terms of required variables during the (reverse) iteration, e.g., Christianson (1998) and Griewank and Walther (2008), chapter 15. Essentially this theorem states that only the variable *at the fixed point* is required, thus avoiding the potentially memory-intensive storing of the entire intermediate state of the iteration. TAF accommodates this feature via directives that identify a loop in the code as fixed-point iteration (Giering and Kaminski, 1998), and which we use here. We note that caveats exist between analytical derivation of the adjoint equations for implicit functions and its validity for numerical implementation (Giles, 2001). Deciding whether the generated code is reliable has to be based, somewhat heuristically, on detailed gradient checks, as was done in this study.

A note on recent developments in the use of fully implicit method in ocean, sea-ice and land-ice modeling seems warranted. Methods such as Jacobi-free Newton-Krylov (JFNK) methods enable very efficient model integrations using rather long time steps and showing very favorable convergence behavior.

Most implementations (in particular those aimed at scalable applications) take advantage of black-box solvers such as GMRES, Trilinos or PETSc. In such cases, differentiation through the solvers is either not possible (black-box) or very difficult and not recommended. Instead, use of the knowledge of the solver for the adjoint system of differential equations and implementation of the adjoint solver (usually part of the same black-box package) is preferable.

Approximating the adjoint of mixing parameterization schemes

Mixing schemes introduce additional nonlinear behavior on various time scales that may cause problems for the adjoint. Generating exact adjoint for most schemes does not per se present a fundamental problem. For example, Marotzke et al. (1999) describe in some detail the adjoint of the convective adjustment scheme. Ferreira et al. (2005) take advantage of the adjoint to estimate eddy-induced stresses in the ocean interior as a way to estimate parameters relevant for eddy-induced mixing.

However, with increasing time scales, resolution, nonlinearity of the scheme, or a combination thereof, the use of the adjoint will be prevented due to exponential growth of sensitivities. Approximating the adjoint under such circumstances has been found to be necessary to retain a stable solution. In the present calculation the approximation was made by excluding the adjoint of the non-local K-profile parameterization (KPP) scheme for vertical mixing (Large et al., 1994).

Some modifications have recently been made to the sea-ice thermodynamics, in particular to the treatment of sea-ice growth, in order to improve both certain forward model features as well as the adjoint model behavior. These changes will be discussed in detail elsewhere (Fenty, 2010).

Concluding remarks Many issues of generating efficient exact adjoint sea-ice code are similar to those for the ocean model's adjoint. Linearizing the model around the exact nonlinear model trajectory is a crucial aspect in the presence of different regimes. For example, is the thermodynamic growth term for sea-ice evaluated near or far away from the freezing point of the ocean surface? Adapting the (parent) model code to support the AD tool in providing exact and efficient adjoint code represents the main workload, initially. For legacy code, this task may become substantial but it is fairly straightforward when writing new code with an AD tool in mind. Once this initial task is completed, generating the adjoint code of a new model configuration takes about 10 minutes.

Acknowledgements

This work is a contribution to the ECCO2 project sponsored by the NASA Modeling Analysis and Prediction (MAP) program and to the ECCO-GODAE project sponsored by the National Oceanographic Partnership Program (NOPP). DM carried out this work at JPL/Caltech under contract with NASA. Computing resources were provided by NASA/ARC, NCAR/CSL, and JPL/SVF. Careful reviews by the two anonymous reviewers significantly improved the readability of the paper and are gratefully acknowledged.

References

Adcroft, A., Campin, J.-M., Heimbach, P., Hill, C., Marshall, J., 2002. Mitgcm release1 manual. (online documentation) , MIT / EAPS, Cambridge, MA 02139, USA, `http://mitgcm.org/sealion/online_documents/manual.html`.

Bugnion, V., Hill, C., Stone, P., 2006a. An adjoint analysis of the meridional overturning circulation in an ocean model. J. Clim. 19(15), 3732–3750.

Bugnion, V., Hill, C., Stone, P., 2006b. An adjoint analysis of the meridional overturning circulation in a hybrid coupled model. J. Clim. 19(15), 3751–3767.

Christianson, B., 1998. Reverse accumulation and implicit functions. Optimisation Methods and Software 9 (4), 307–322.

Duliere, V., Fichefet, T., 2007. On the assimilation of ice velocity and concentration data into large-scale sea ice models. Ocean Science 3, 321–335.

Fenty, I., 2010. State estimation of the Labrador Sea with a coupled ocean/sea-ice adjoint model. Ph.D. thesis, MIT, Program in Atmosphere, Ocean and Climate (PAOC), Cambridge (MA), USA.

Ferreira, D., Marshall, J., Heimbach, P., 2005. Estimating eddy stresses by fitting dynamics to observations using a residual-mean ocean circulation model and its adjoint. J. Phys. Oceanogr. 35(10), 1891–1910.

Galanti, E., Tziperman, E., 2003. A midlatitude–ENSO teleconnection mechanism via baroclinically unstable long Rossby waves. J. Phys. Oceanogr. 33, 1877–1887.

Galanti, E., Tziperman, E., Harrison, M., Rosati, A., Giering, R., Sirkes, Z., 2002. The equatorial thermocline outcropping - a seasonal control on the tropical Pacific ocean-atmosphere instability. J. Clim. 15, 2721–2739.

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

Giles, M., 2001. On the iterative solution of adjoint equations. In: Corliss, C., Faure, C., Griewank, A., Hascoet, L., Naumann, U. (Eds.), Automatic Differentiation: From Simulation to Optimization. Springer-Verlag, pp. 145–152.

Griewank, A., Walther, A., 2008. Evaluating Derivatives. Principles and Techniques of Algorithmic Differentiation, 2nd Edition. Vol. 19 of Frontiers in Applied Mathematics. SIAM, Philadelphia.

Heimbach, P., January 2008. The MITgcm/ECCO adjoint modelling infrastructure. CLIVAR Exchanges 44 (Volume 13, No. 1), 13–17.

Heimbach, P., Bugnion, V., 2009. Greenland ice sheet volume sensitivity to basal, surface, and initial conditions, derived from an adjoint model. Annals Glaciol. 52, in press.

Heimbach, P., Hill, C., Giering, R., 2005. An efficient exact adjoint of the parallel MIT general circulation model, generated via automatic differentiation. Future Generation Computer Systems 21(8), 1356–1371.

Heimbach, P., Wunsch, C., 2007. Estimating the Circulation and Climate of the Ocean – The ECCO Consortia. U.S. CLIVAR Variations V5N3, .

Hibler, III, W. D., 1979. A dynamic thermodynamic sea ice model. J. Phys. Oceanogr. 9 (4), 815–846.

Kauker, F., Kaminski, T., Karcher, M., Giering, R., Gerdes, R., Vosbeck, M., 2009. Adjoint analysis of the 2007 all time arctic sea-ice minimum. Geophys. Res. Lett. submitted.

Kim, J., Hunke, E., Lipscomb, W., 2006a. Sensitivity analysis and parameter tuning scheme for global sea-ice modeling. Ocean Modelling 14 (1-2), 61–80.

Kim, J., Hunke, E., Lipscomb, W., 2006b. A sensitivity-enhanced simulation approach for community climate system model. In: Alexandrov, V., van Albada, G., Sloot, P., Dongarra, J. (Eds.), Computational Science – ICCS 2006. Vol. 3994, Part IV of Lecture Notes in Computer Science. Springer Verlag, pp. 533–540.

Köhl, A., 2005. Anomalies of meridional overturning: Mechanisms in the North Atlantic. J. Phys. Oceanogr. 35(8), 1455–1472.

Kwok, R., 2006. Exchange of sea ice between the Arctic Ocean and the Canadian Arctic Archipelago. Geophys. Res. Lett. 33, L16501.

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

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

Lietaer, O., Fichefet, T., Legat, V., 2008. The effects of resolving the Canadian Arctic Archipelago in a finite element sea ice model. Ocean Modelling 24, 140–152.

Lisaæter, K., Evensen, G., Laxon, S., 2007. Assimilation of synthetic cryosat sea ice thickness in a coupled ice-ocean model. J. Geophys. Res. 112, C07023.

Losch, M., Heimbach, P., 2007. Adjoint sensitivity of an ocean general circulation model to bottom topography. J. Phys. Oceanogr. 37(2), 377–393.

Losch, M., Menemenlis, D., Campin, J., Heimbach, P., Hill, C., 2010. A dynamic-thermodynamic sea ice model for ocean climate modeling on an arakawa c-grid: Part 1: Forward model sensitivities. Ocean Modelling accepted, .

Marchuk, G., 1995. Adjoint equations and analysis of complex systems. Kluwer Academic Publishers.

Marotzke, J., Giering, R., Zhang, K., Stammer, D., Hill, C., Lee, T., 1999. Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport variability. J. Geophys. Res. 104, C12, 29,529–29,547.

Melling, H., 2002. Sea ice of the Northern Canadian Arctic Archipelago. J. Geophys. Res. 107 (C11), 3181.

Menemenlis, D., Hill, C., Adcroft, A., Campin, J.-M., Cheng, B., Ciotti, B., Fukumori, I., Koehl, A., Heimbach, P., Henze, C., Lee, T., Stammer, D., Taft, J., Zhang, J., 2005. NASA supercomputer improves prospects for ocean climate research. Eos Trans. AGU 86 (9), 89, 95–96.

Michel, C., Ingram, R., Harris, L., 2006. Variability in oceanographic and ecological processes in the Canadian Arctic Archipelago. Prog. Oceanogr. 71, 379–401.

Miller, P., Laxon, S., Feltham, D., Cresswell, D., 2006. Optimization of a sea ice model using basinwide observations of arctic sea ice thickness, extent, and velocity. J. Clim. 19 (1089-1108).

Moore, A., Arango, H., Di Lorenzo, E., Miller, A., Cornuelle, B., 2009. An adjoint sensitivity analysis of the Southern California Current circulation and ecosystem. J. Phys. Oceanogr. 39 (3), 702–720.

Münchow, A., Melling, H., Falkner, K., 2006. An observational estimate of volume and freshwater flux leaving the arctic ocean through Nares Strait. J. Phys. Oceanogr. 36 (11), 2025–2041.

Panteleev, G., Nechaev, D., Proshutinsky, A., Woodgate, R., Zhang, J., 2010. Reconstruction and analysis of the chukchi sea circulation in 1990-1991. J. Geophys. Res. accepted, .

Prinsenberg, S., Hamilton, J., 2005. Monitoring the volume, freshwater and heat fluxes passing though Lancaster Sound in the Canadian Arctic

Archipelago. Atmosphere-Ocean 43 (1), 1–22.

Stammer, D., Wunsch, C., Giering, R., Eckert, C., Heimbach, P., Marotzke, J., Adcroft, A., Hill, C., Marshall, J., 2002a. The global ocean circulation and transports during 1992 – 1997, estimated from ocean observations and a general circulation model. J. Geophys. Res. 107(C9), 3118.

Stark, J., Ridley, J., Martin, M., Hines, A., 05 2008. Sea ice concentration and motion assimilation in a sea ice-ocean model. J. Geophys. Res. 113.

Stoessel, A., 2008. Employing satellite-derived sea ice concentration to constrain upper-ocean temperature in a global ocean gcm. J. Clim. 21, 4498–4513.

Utke, J., Harscoet, L., Heimbach, P., Hill, C., Hovland, P., Naumann, U., 2009. Toward adjointable mpi. In: Proceedings of the 23rd IEEE International Parallel and Distributed Processing Symposium. Vol. in press. Rome, Italy, p. .

Utke, J., Naumann, U., Fagan, M., Tallent, N., Strout, M., Heimbach, P., Hill, C., Ozyurt, D., Wunsch, C., 2008. Openad/f: A modular open source tool for automatic differentiation of fortran codes. ACM Transactions on Mathematical Software 34(4), .

Veneziani, M., Edwards, C., Moore, A., 2009. A Central California modeling study. part ii: Adjoint sensitivities to local and remote driving mechanisms. J. Geophys. Res. in press, .

Wessel, P., Smith, W. H. F., 1996. A global self-consistent, hierarchical, high-resolution shoreline database. J. Geophys. Res. 101 (B4), 8741–8743.

Wunsch, C., 1996. The ocean circulation inverse problem. Cambridge University Press, Cambridge (UK).

Wunsch, C., 2006. Discrete Inverse and State Estimation Problems: With Geophysical Fluid Applications. Cambridge University Press.

Wunsch, C., Heimbach, P., 2007. Practical global oceanic state estimation. Physica D 230(1-2), 197–208.

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