UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS

School of Ocean and Earth Sciences

The Use of Adjoint Models for Determining the Sensitivity of Integral Quantities in an Eddy Resolving Ocean General Circulation Model

by

Fiona M^{c}Lay

Thesis for the degree of Doctor of Philosophy

September 2006

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF OCEAN AND EARTH SCIENCES

Doctor of Philosophy

THE USE OF ADJOINT MODELS FOR DETERMINING THE SENSITIVITY OF INTEGRAL QUANTITIES IN AN EDDY RESOLVING OCEAN GENERAL CIRCULATION MODEL

by Fiona M^{c}Lay

Adjoint models calculate the exact sensitivity of an output function of a model to infinitesimal perturbations in the forcing or initial conditions. In eddy resolving ocean models the presence of chaotic eddies is expected to lead to sensitivities to infinitesimal perturbations that are very different from the sensitivity to large perturbations and that no longer contain useful information. Previous studies disagree as to whether adjoint models can be used with eddy resolving ocean models on timescales longer than a few months.

Here the MIT ocean general circulation model and its adjoint are used to look at the sensitivity of the time mean heat content, kinetic energy, available potential energy and thermocline depth to the sea surface temperature, zonal wind stress, and vertical diffusivity in an eddy resolving model of a zonally reentrant channel. Using the tangent linear model the non linear timescale of the eddy resolving model is estimated at around 200 days. The adjoint model is integrated over 278 days and 690 days to see whether useful information remains in the sensitivities calculated by the adjoint model for longer than the non linear timescale of the system. The usefulness of the information in the sensitivities calculated by the adjoint model is assessed by comparison with integrations of the full non linear forward model with large spatial scale perturbations to the forcing, finite difference gradient checks, and sensitivities calculated by an adjoint model in a non eddy resolving channel where the adjoint method is known to provide useful information.

Finite difference gradients are found to be unsuitable for calculating sensitivities of time averaged climate quantities in an eddy resolving ocean model as they are also affected by chaos. Comparison of the sensitivities calculated by the adjoint model in the eddy and non eddy resolving models shows that information remains in the spatial structure of the adjoint model results in the eddy resolving model on a time scale of 278 days.

In the non eddy resolving case the adjoint model results agree well with the perturbed forward model experiments, and are clearly climatically relevant on a timescale of 690 days. Use of a parameterisation scheme that reduces the eddy kinetic energy gives adjoint sensitivities that agree with well the perturbed forward model experiments after 690 days, although there are areas of extremely high adjoint sensitivity that may not be physically realistic. Without this parameterisation scheme, adjoint sensitivities involving dynamic variables grow exponentially with time as expected in a chaotic system, but at the end of the integration time of 690 days there is some agreement between the adjoint and forward model results for sensitivities involving thermodynamic variables only.

These results show that even in the presence of chaotic eddies some useful information is retained in the adjoint model solution beyond the nonlinear timescale of the system.

Contents

List of Figures

List of Tables

DECLARATION OF AUTHORSHIP

I, Fiona McLay ....., [please print name]

declare that the thesis entitled [enter title]

The use of adjoint models for determining the sensitivity of integral quantities in an eddy resolving ocean general circulation model

and the work presented in it are my own. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- ~~none of this work has been published before submission; or~~ [delete as appropriate] parts of this work have been published as: [please list references]

McLay F and Marotzke J. Limitations on the use of adjoint models with eddy resolving ocean general circulation models. submitted to Ocean Modelling, 2006.

Signed: .....

Date:.....

Graduate School of the National Oceanography Centre, Southampton

This PhD dissertation by

Fiona M^{c}Lay

has been produced under the supervision of the following persons

Supervisor

Prof. Jochem Marotzke

Chair of Advisory Panel

Prof. Harry Bryden

Member of Advisory Panel

Dr. Neil Wells

Acknowledgements

I am grateful to many people for the help they gave me while I was working on this thesis:

Jochem Marotzke, my supervisor, provided advice and support throughout my time as a PhD student. Peter Korn and Luis Kornblueh gave lots of helpful suggestions, unfortunately not all of which were followed up. Patrick Heimbach, provided invaluable assistance in getting the adjoint and tangent linear models running and John Stark and Joël Hirschi taught me how to use a GCM.

Robin Smith and Dirk Notz provided lots of sage advice about how to do a PhD, having recently trodden down the path. Johanna Baehr and Kornelia Müller helped with negotiating German officialdom when I moved to Hamburg. Neil Wells and Harry Bryden were on my advisory panel at Southampton and made sure I stayed in the system when I moved to the MPI.

My friends in Hamburg and elsewhere made doing a PhD bearable, and sometimes even fun. In particular, Beena B S managed to make me laugh whatever mood I was in, Daniela Matei told me snap out of it when I was in a bad mood, and Clotilde Dubois put up with me for longest and did both at the same time. Finally my family were always there when I needed them.

Nomenclature and Symbols

ϵ Step size used in finite difference calculations.

ϵ_J Fractional accuracy with which the output function J is computed.

ϵ_m Machine accuracy.

δ Non dimensional strength of the thermohaline circulation in the Stommel (1961) box model.

η Elevation of the implicit free surface.

κ_h Horizontal diffusivity.

κ_v Vertical diffusivity.

μ Constants in the Stommel box model coupled to the Lorenz 1963 equations with a passive tracer.

Ξ Time mean advective transport across the mid basin line of a double gyre in Lea et al. (2002).

ρ_0 Reference density.

σ Potential density.

σ Independent variable in the Lorenz 1963 equations.

σ_{max} Maximum Eady growth rate.

τ Integration time.

τ_{CD} Coupling timescale in the C-D coupling scheme.

τ_T Restoring timescale for SST*.

τ_x Zonal windstress.

τ_{x_0} Maximum zonal windstress.

ϕ Latitude.

ϕ_m Latitude of the centre of the channel.

Ψ_n The numerical model operator at timestep n.

Ω Rate of rotation of the earth.

A Available potential energy.

A_h Horizontal viscosity.

A_v Vertical viscosity.

B Meridional gradient of SST* in the channel.

b Independent variable in the Lorenz 1963 equations.

C Concentration of the passive tracer in the Stommel box model coupled to the Lorenz 1963 equations with a passive tracer

$\overline{C}$ Percentage of time steps convection occurs between the top two model layers.

c Constant coupling parameter in coupled Stommel model.

D Thermocline depth.

E Non dimensional strength of the fresh water forcing in the Stommel (1961) box model.

F Fluxes in the Stommel box model coupled to the Lorenz 1963 equations with a passive tracer.

f Coriolis parameter.

g Acceleration due to gravity.

H Total heat content of the channel.

H Depth of the channel.

J Scalar output or cost function of the model output.

K Total kinetic energy of the channel.

Meridional Along a line of longitude.

N Buoyancy frequency.

P Pressure.

Q_T Surface heat flux.

Q_S Surface salinity flux.

S Salinity.

SSH Sea surface height.

SST Sea surface temperature.

SST* Apparent atmospheric temperature after effects of solar radiation and evaporation on the heat flux have been taken into account.

SST$_0^*$ Apparent atmospheric temperature in the centre of the channel.

T Potential temperature.

T_1 Potential temperature of the top model layer.

$T_{0_{\mathbf{maxsens}}}$ Potential temperature at the point of maximum sensitivity

TLM Tangent linear model.

t Time.

r Independent variable in the Lorenz 1963 equations.

U Mean zonal velocity.

u Zonal velocity.

v Meridional velocity.

$\mathbf{v_h}$ Horizontal velocity.

w Strength of the windstress curl in Lea et al. (2002).

$\mathbf{X_0}$ Vector of model inputs.

$\mathbf{X_n}$ State of the model at timestep n.

x Dependent variable in the Lorenz 1963 equations.

y Dependent variable in the Lorenz 1963 equations.

z Dependent variable in the Lorenz 1963 equations.

z Vertical coordinate in the channel model.

Zonal Along a line of latitude

Chapter 1

Introduction

1.1 Introduction

Almost all the questions that climate models are used to address are fundamentally ones of sensitivity. We wish to know how a change in the model inputs will affect the outputs. The output is the simulated climate of the model, or some function of it, while the inputs are not only the forcing fields used to drive the model but also the physics used to parameterise poorly understood or resolved processes. Adjoint models are an extremely powerful tool for sensitivity analysis of complex climate models, and they have already been widely used for sensitivity studies (e.g Marotzke et al. (1999); Junge and Haine (2001); Bugnion et al. (2006b)) and data assimilation (e.g Stammer et al. (2002); Moore et al. (2004)), both in oceanography and in other fields.

However previous studies have shown that there is a time limit to the applicability of the adjoint method in chaotic systems, including eddy resolving ocean models (Köhl and Willebrand, 2002; Lea et al., 2002). As computer power increases it is likely that eddy resolving ocean models will be used more frequently and guidance is needed about which problems the adjoint method can still provide insight. This chapter begins with an introduction to adjoint models and some of their uses in physical oceanography. The issues surrounding the use of adjoint models in chaotic systems are then discussed in the context of the Lorenz (1963) equations, with some extensions to the existing literature in sections 1.3.3

and 1.3.4. Previous work using adjoint models and eddy resolving ocean models is reviewed in section 1.3.5 and finally the questions addressed in this thesis are outlined in section 1.4.

1.2 A General Introduction to Adjoint Models

Two possible methods for calculating the sensitivity of a quantity are the direct method and the adjoint method. Before looking at the adjoint method in more detail, it is important to see how it compares to the direct method and the advantages that can be gained by using it.

The Direct Method

In the direct method a perturbation is added to the kth component of the vector of model inputs $X_{0_k} + \Delta X_{0_k}$, resulting in a new value of the output function $J + \Delta J$. The gradient of the output function with respect to this component of model input is then evaluated

$$\frac{\partial J}{\partial X_{0_k}} \simeq \frac{\Delta J}{\Delta X_{0_k}} \equiv \frac{J(X_{0_k} + \Delta X_{0_k}) - J(X_{0_k})}{\Delta X_{0_k}} \tag{1.1}$$

The sensitivity of all possible Js to $X_{0_k} + \Delta X_{0_k}$ can thus be calculated from two forward model integrations. However, we normally are interested in only a few J while $\mathbf{X_0}$ may have many thousands of components in a GCM. Ensembles of model runs with different perturbations must be used, and as physically meaningful changes in the parameters do not normally move the climatic mean of the model outside the range of fluctuations of the model's dependent variables, it is necessary to use integration times which are long compared to the natural timescales of the model, (Hall et al., 1982; Lea et al., 2000). The direct method is therefore prohibitively expensive in terms of computational time, although it is possible that distributed computing initiatives such as that of the climateprediction.net experiment, which aims to get the public to run climate models on their home PCs, may make it possible (Stainforth et al., 2005).

The Adjoint Method

In contrast to the direct method the sensitivity of a single J to perturbations in all possible components of $\mathbf{X}_0$, can be calculated by a single application of the adjoint model. The resulting computational time is only about 5 times greater than for the original forward integration (Griewank, 1992).

The general principle of adjoint and tangent linear models is outlined in several places, notably Errico (1997), Giering and Kaminski (1998) and Marotzke et al. (1999), but is included here for convenience. This approach follows that in Marotzke et al. (1999).

The state of the model at time-step n can be written $\mathbf{X}_n$, $0 \leq n \leq N$ where $\mathbf{X}_0$ is the vector of initial conditions, and $\mathbf{X}_N$ is the final state of the model. Any function of the model output $J = \sum_{n=1}^{N} f_n \circ \mathbf{X}_n$, can be computed from $\mathbf{X}_0$ by repeated application of the numerical model Ψ, where f_n gives the contribution to J at time step n only.

$$J = \sum_{n=1}^{N} f_n \circ \Psi_n \circ \Psi_{n-1} \circ \cdots \circ \Psi_2 \circ \Psi_1 \circ \mathbf{X}_0. \tag{1.2}$$

The sensitivity of J to $\mathbf{X}_0$ is then given by the chain rule.

$$\frac{\partial J}{\partial \mathbf{X}_0} = \sum_{n=1}^{N} 1 f_n' \left(\Psi_n' \left(\Psi_{n-1}' \left(\cdots \left(\Psi_2' \left(\Psi_1' \left(\mathbf{I} \right) \right) \right) \right) \right) \right) \tag{1.3}$$

where the prime denotes differentiation with respect to the argument, and each Ψ_n' is the Jacobian matrix of the model at time step n. Each Ψ_n' corresponds to a linearisation of the model around the model trajectory and is termed the tangent linear model (TLM). The TLM evolves an infinitesimal perturbation to the initial conditions forwards in time to give the perturbation at the final time, under the assumption of linear dynamics.

We can take the transpose of equation 1.3 to get

$$\left(\frac{\partial J}{\partial \mathbf{X}_0} \right)^T = \sum_{n=1}^{N} (\Psi_1')^T \left((\Psi_2')^T \left(\cdots \left((\Psi_{n-1}')^T \left((\Psi_n')^T \left((f_n')^T \left(1 \right) \right) \right) \right) \right) \right) \tag{1.4}$$

This is the adjoint model, which is often said to operate "backwards in time" as it calculates the sensitivity with respect to model input from a sensitivity with respect to model output. As a matrix has the same eigenvalues as its transpose, the adjoint and tangent linear models have the same eigenvalues, meaning that if perturbations in the TLM grow exponentially forwards in time, sensitivities in the adjoint model will grow exponentially backwards in time.

Although the above only considers initial condition sensitivity, it is possible to augment the state vector with model parameter and forcing values so that the adjoint model also calculates the sensitivity to these.

Due to the time taken to produce an adjoint model, many of the early adjoint models were approximate, and they did not become widely used until the development of software tools, such as TAMC, Tangent linear and Adjoint Model Compiler (Giering and Kaminski, 1998) which allow adjoint models to be developed automatically from the FORTRAN code for the forward model. Among others TAMC has been used for the MIT model (Marotzke et al., 1999) and the Hamburg Ocean Primitive Equation Model (HOPE) (Junge and Haine, 2001; van Oldenborgh et al., 1999).

1.2.1 Uses of Adjoint Models in Physical Oceanography

Adjoint models have been used for two major types of problem in oceanography, sensitivity analysis, and 4 dimensional variational data assimilation (4DVAR). In sensitivity analysis the output function is some quantity of interest in the forward model, such as a measure of the heat transport, and the sensitivity information from the adjoint model is used by itself to gain an understanding of physical processes. In data assimilation the output, or cost, function is a measure of the model-data misfit and the adjoint sensitivities are used as gradient information in descent algorithms.

Sensitivity Analysis

The first sensitivity studies using an adjoint of a OGCM were those of Marotzke et al. (1999) who examined the sensitivity of the heat transport across 29°N in the Atlantic, to temperature and salinity anomalies, and van Oldenborgh et al. (1999) who investigated some of the mechanisms behind ENSO. Since then adjoint models have been used for identifying an oscillatory mode in the THC (Sirkes and Tziperman, 2001), investigating the sensitivity of the meridional overturning circulation to global wind stress (Bugnion et al., 2006a,b), studying the heat uptake of the deep ocean (Huang et al., 2003a,b) and looking at the persistence of winter SST anomalies in the North Atlantic (Junge and Haine, 2001).

Data Assimilation

In meteorology 4DVAR is most commonly used to initialise forecast models, while in oceanography it is more commonly used to compensate for the sparsity of data. By bringing models into consistency with available data it is possible to study the general circulation and its variability. This approach has been used by Tziperman and Thacker (1989); Tziperman et al. (1992a,b), Marotzke and Wunsch (1993), and Stammer et al. (2002) and others.

1.3 Adjoint Models and Chaotic Systems

1.3.1 Chaos and Predictability

All the above examples of the uses of adjoint models used ocean models where the resolution was too low to resolve oceanic eddies. Increasing the resolution so that eddies are resolved leads to a more realistic representation of the previously parameterised small scales and of the large scale mean circulation (Smith et al., 2000; Hogan and Hurlburt, 2000). However, resolving eddies means that the flow may become chaotic.

In a chaotic system there is sensitive dependence on initial conditions so that the distance between two trajectories in phase space, $|\delta(t)|$, initially separated

by a distance $|\delta_0|$ grows exponentially so that $|\delta(t)| \sim |\delta_0|e^{\lambda t}$, where $\lambda > 0$. Predicting the state of the system to within an amount a of the true state is therefore only possible out to a time,

$$t_{horizon} \sim \left(\frac{1}{\lambda} \ln\left(\frac{a}{\delta_0}\right)\right) \tag{1.5}$$

where δ_0 is the difference between the true and measured states at the initial time. The extreme sensitivity to the initial conditions means that improving the measurement at the initial time does not significantly reduce $t_{horizon}$ due to logarithmic dependence on δ_0 (Strogatz, 1994). This results in a limit to what Lorenz (1975) termed *predictability of the first kind*, where the state of a system at a future time may be predicted given the initial conditions. Beyond this limit, the system may still have *predictability of the second kind* where the effect of a change in the forcing on the statistics of the system, or its climate, may still be predictable. It is predictability of the second kind that is normally of interest in climate modelling; we are not usually interested in instantaneous state of the turbulent ocean, but rather some time averaged quantity such as the mean heat transport. Lorenz (1975) argued that it is necessary for predictions of the second kind that integrations are extended beyond the limit of predictability of the first kind, or the climate will not be representative of the climate evaluated over a longer interval. If the integration time is long enough we might expect that the sensitivity of a time averaged climate quantity to the initial conditions calculated by the adjoint model is zero, and that the sensitivity to the forcing or parameter values approaches a constant value. In a non chaotic systems this indeed happens; Bugnion and Hill (2006) found that, in an ocean model that did not resolve chaotic eddies the sensitivity of the meridional overturning circulation to the initial sea surface temperature decayed to zero after around 400 years, while the sensitivity to a perturbation to the zonal windstress reached a constant value. However, studies have shown experimentally that, in a chaotic system, the sensitivity of a time averaged climate quantity to the initial conditions and forcing calculated by the adjoint method also grows exponentially backwards in time (Lea et al., 2000, 2002).

1.3.2 Studies of the Lorenz (1963) System

The Lorenz (1963) system of equations are a simplified description of convection between two parallel plates, which give rise to the well known butterfly shaped attractor. The system has been used by Stensrud and Bao (1992), Lea et al. (2000), and Köhl and Willebrand (2002) to study the limitations of the adjoint method in a chaotic system. The small number of parameters means that direct sensitivity calculations are computationally feasible, while for certain choices of parameter values the model exhibits chaotic behaviour, which is thought to be analogous to that in the atmosphere and oceans. The Lorenz (1963) equations are

$$\begin{aligned} \frac{\mathrm{d}x}{\mathrm{d}t} &= -\sigma x + \sigma y \\ \frac{\mathrm{d}y}{\mathrm{d}t} &= -xz + rx - y \\ \frac{\mathrm{d}z}{\mathrm{d}t} &= xy - bz \end{aligned} \tag{1.6}$$

Stensrud and Bao (1992) found that in the chaotic regime 4DVAR becomes less successful in recovering the initial conditions in the Lorenz (1963) model as the assimilation window is increased beyond a certain timescale, due to the development of local minima in the cost function. This is directly related to loss of predictability of the first kind. The exponential divergence of nearby trajectories means that the sensitivity to initial conditions grows exponentially backwards in time, and the output or cost function develops many local minima. The sensitivity calculated by the adjoint model is still exact but no longer contains useful information.

Lea et al. (2000) looked at the sensitivity of the time mean of z, $\overline{z}$, to r. Provided the integration time is much longer than the limit of predictability of the first kind $\overline{z}$ varies almost linearly with r over much of the range (figure 1.1A), and has predictability of the second kind. However, an infinitesimal change in a parameter value leads to exponential divergence of model trajectories similar to that seen for an infinitesimal change in the initial conditions. The climate

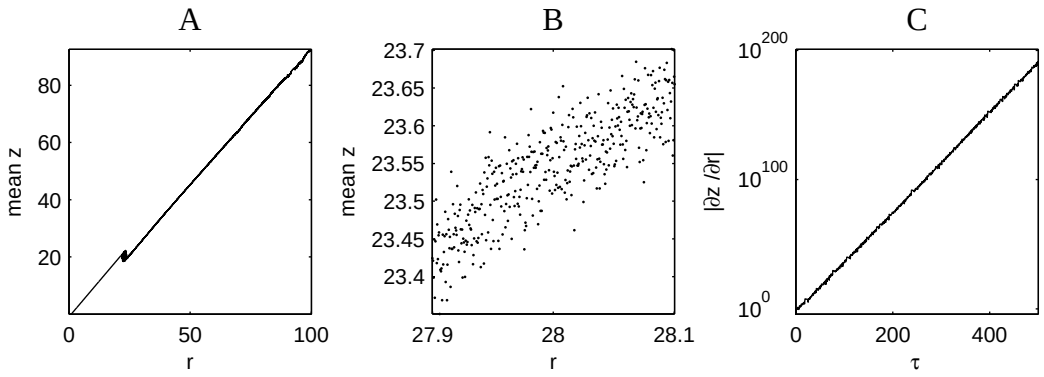

Figure 1.1: A) $\overline{z}$ as a function of r in the Lorenz (1963) equations, for $\sigma = 10$ and $b = 8/3$. B) A close up of $\overline{z}$ as a function of r in the Lorenz (1963) equations, for $\sigma = 10$ and $b = 8/3$. C) $d\overline{z}/dr$ calculated by adjoint model as a function of integration time, τ, for $r = 28$. After Lea et al. (2000).

of a chaotic system is not a smooth function of the parameter value but has a number of local maxima and minima (figure 1.1B) that increases exponentially with increasing integration time. The adjoint model calculates the sensitivity to an infinitesimal perturbation, which Lea et al. (2000) termed the microscopic sensitivity, as opposed to the macroscopic sensitivity which would be calculated for larger perturbations. The microscopic sensitivity grows exponentially backwards in time with the increasing number of local maxima and minima in the output function (figure 1.1C). Köhl and Willebrand (2002) looked at the similar problem of using the adjoint method for parameter value estimation in the Lorenz (1963) equations, and got similar results.

1.3.3 Finite Difference Gradient Checks in the Lorenz (1963) Equations

Lea et al. (2000) estimated the macroscopic sensitivity $\Delta\overline{z}/\Delta r$ from a linear fit to the graph of $\overline{z}$ as a function of r, where the time average had been evaluated over around 100 orbits of the chaotic attractor (figure 1.1A). This led them to conclude that the direct method can be used to calculate the macroscopic sensitivities in the Lorenz (1963) model. However, unlike the 2 integrations required for calculating sensitivities using the direct method in section 1.2, Lea et al. (2000)'s estimate of

$\frac{\Delta\overline{z}}{\Delta r}$ required $\overline{z}$ to be evaluated for a large range of r. This is feasible in a system with few parameters such as the Lorenz (1963) model, but not in a GCM where using 2 integrations per parameter is already computationally infeasible.

Lea et al. (2000) show experimentally that in the Lorenz (1963) equations the amplitude of the isolated extrema decreases with $1/\sqrt{\tau}$, where τ is the integration time, so that if the integration time is long enough, and the perturbation is large enough $\Delta\overline{z}/\Delta r$ calculated from a finite difference gradient should approach the macroscopic sensitivity calculated from a linear fit to the graph. However, it is not clear a priori what the correct integration time and perturbation size should be. For a non chaotic system the optimum step size for centred finite difference gradients scales as $\sqrt[3]{\epsilon_J}$ where ϵ_J is the fraction accuracy with which the output function J is computed (Press et al., 1992). In the best case scenario ϵ_J is equivalent to the machine accuracy ϵ_m which is around 2×10^{-16} for double precision floating point numbers, giving an optimal step size of 1×10^{-5}. In the non chaotic regime of the Lorenz (1963) equations finite difference gradients converge for $\epsilon > 1 \times 10^{-5}$ (figure 1.2A), while in the chaotic regime they only approach the larger scale sensitivity for $\epsilon = 1$ (figure 1.2B), even though in this case the integration time is longer than that used by Lea et al. (2000) for their estimate of the macroscopic sensitivity using a linear fit. Finite difference gradient checks are often used as a measure of the accuracy of sensitivity calculated by an adjoint model, the experience here with the Lorenz (1963) model shows that they need to be treated with caution in a chaotic system.

1.3.4 Lorenz (1963) System Coupled to a Non Chaotic System

Although the work of Lea et al. (2000) and Köhl and Willebrand (2002) showed that the adjoint method will not work with the Lorenz (1963) equations there are obvious ways in which the Lorenz (1963) equations are not a good analogy for the real climate. One of these is the absence of a much longer timescale, and it is interesting to see if these results hold when slower non chaotic elements are added

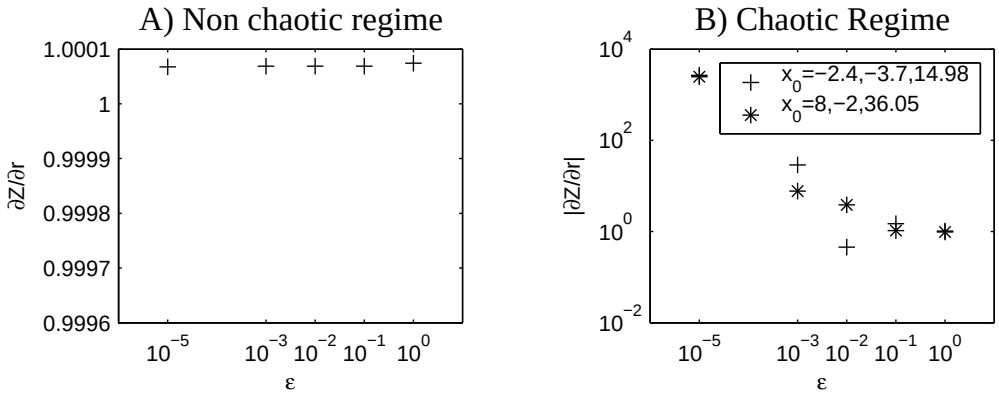

Figure 1.2: Finite difference gradient checks as a function of the step size ϵ for the Lorenz (1963) equations. A) Non chaotic regime $r = 10$, $b = 8/3$, $\sigma = 10$, $\tau = 500$ B) Chaotic regime $r = 28$, $b = 8/3$, $\sigma = 10$, for 2 sets of initial conditions x_0, integration time $\tau = 500$.

to the system. In particular do the adjoint sensitivities grow at the same rate as in the simple Lorenz (1963) model? The system considered here is the Stommel (1961) box model of the North Atlantic thermohaline circulation coupled to the Lorenz (1963) equations, and the same system with a passive tracer added.

Stommel (1961) Box Model Coupled to the Lorenz (1963) Equations

The Stommel (1961) box model can be fully described by the non dimensional equation,

$$\frac{\partial \delta}{\partial t} = E - |1 - \delta|\delta \tag{1.7}$$

where δ is the non dimensional strength of the thermohaline circulation and E is the non dimensionalised strength of the fresh water forcing (Marotzke, 2002). This is coupled to the Lorenz (1963) equations by writing

$$E = E_0 + cz \tag{1.8}$$

where E_0 and c are constants, and z is given by equation 1.6. The model can be considered roughly analogous to a chaotic atmosphere coupled to a non chaotic ocean. Provided the integration time is long enough the system reaches an equilibrium state where δ is determined by the mean value of E, $\bar{E}$. In this case the

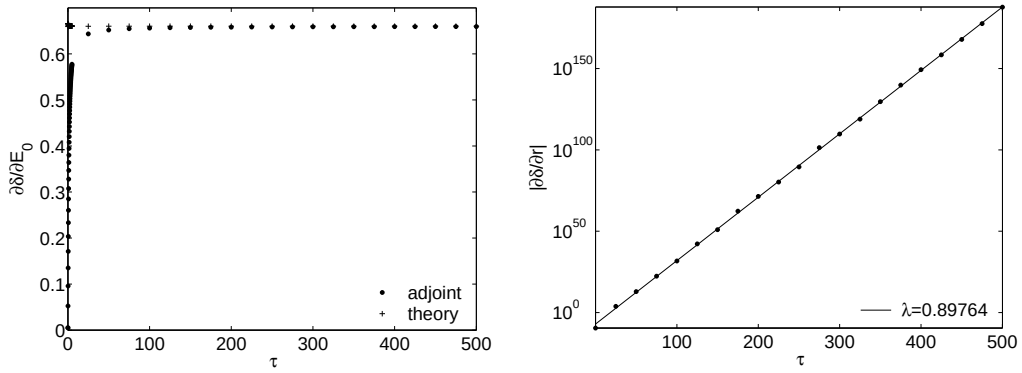

Figure 1.3: Sensitivities to parameters in the coupled Stommel (1961) model calculated by the adjoint model as a function of integration time τ. Left, sensitivity to a parameter in the Stommel (1961) model $\frac{\partial \delta}{\partial E_0}$. Right, sensitivity to a parameter in the Lorenz (1963) equations $\frac{\partial \delta}{\partial r}$.

sensitivities of δ to c and E_0 can be derived analytically to give

$$\frac{\partial \delta}{\partial c} = \bar{z} \left. \frac{\partial \delta}{\partial E} \right|_{\bar{E}} \tag{1.9}$$

$$\frac{\partial \delta}{\partial E_0} = \left. \frac{\partial \delta}{\partial E} \right|_{\bar{E}} \tag{1.10}$$

The equations are integrated numerically using a 4th order Runge-Kutta scheme (Press et al., 1992), with a timestep of 0.005, using the classical parameters for the Lorenz (1963) model $r = 28$, $\sigma = 10$, $b = 8/3$, and using $E_0 = 0.3$, $c = 0.001$, and $\delta_0 = 1.2$ in the Stommel (1961) model. The adjoint is generated with help of TAMC (Giering and Kaminski, 1998). The adjoint sensitivities $\frac{\partial \bar{\delta}}{\partial c}$ and $\frac{\partial \bar{\delta}}{\partial E_0}$ agree closely with the calculated values, however sensitivities to the parameters in the Lorenz (1963) model still grow exponentially backwards in time (figure 1.3), at the same rate as in the uncoupled Lorenz (1963) equations.

This result is not promising for the prospect of determining the sensitivity of climate quantities using an adjoint model. Adding a slower timescale to the system does not reduce the exponential growth of the sensitivity calculated by the adjoint method. However, it suggests that it may be possible to use an adjoint to determine the sensitivity of the concentration of a passive tracer to non dynamical parameters. This is explored in the next section.

Passive Tracers in a Stommel (1961) Box Model Coupled to the Lorenz (1963) Equations

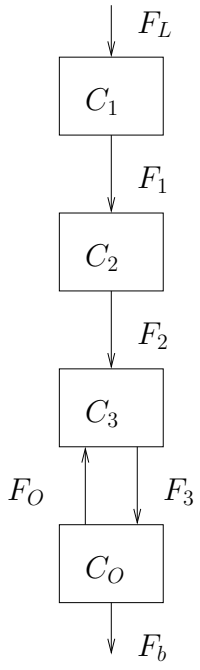

$$F_1 = \mu_1 \left|\frac{\partial x}{\partial t}\right| C_1 \qquad (1.11)$$

$$F_2 = \mu_2 \left|\frac{\partial y}{\partial t}\right| C_2 \qquad (1.12)$$

$$F_3 = \mu_3 \left|\frac{\partial z}{\partial t}\right| C_3 \qquad (1.13)$$

$$F_O = \mu_O C_O \qquad (1.14)$$

$$F_b = \mu_b \delta C_O \qquad (1.15)$$

Figure 1.4: Stommel (1961) box model coupled to the Lorenz (1963) equations with a passive tracer C.

The coupled Stommel (1961) model, section 1.3.4, is modified by the addition of a passive tracer C. For the tracer the model consists of 4 boxes, three 'chaotic atmosphere' boxes with tracer concentrations C_1, C_2, and C_3 and an 'ocean box' with tracer concentration C_O. The tracer is added to the system at a rate F_L, and then advected between the 'atmosphere' boxes according to equations 1.11-1.13, where μ_1, μ_2, and μ_3 are constants, $\frac{\delta x}{\delta t}$, $\frac{\delta y}{\delta t}$ and $\frac{\delta z}{\delta t}$, are given by equation 1.6, and F_1,F_2 and F_3 are the fluxes between the boxes. There are two fluxes out of the 'ocean' box, F_O, which is fed back into the 'atmosphere', equation 1.14, where μ_O is a constant, and F_b, equation 1.15, where δ is given by equation 1.7. F_b can be taken as representing the flux of the tracer to the deep ocean. The dynamical parameters have the same values as used in the coupled Stommel (1961) model, section 1.3.4, while the non dynamically active parameters have values $F_L = 0.5$, $\mu_1 = 0.0025$, $\mu_2 = \mu_3 = \mu_O = 0.001$ and $\mu_b = 0.1$, and are selected so that the system reaches equilibrium in a reasonable timescale.

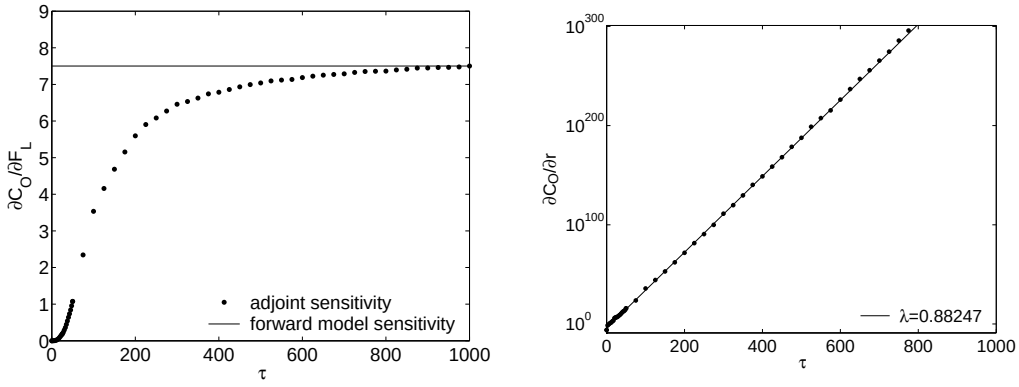

Figure 1.5: Sensitivities to parameters in the coupled Stommel (1961) model with passive tracers, calculated by the adjoint model, as a function of integration time τ. Left, sensitivity to a parameter in the tracer equations $\frac{\partial C_O}{\partial F_L}$. Right, sensitivity to a parameter in the Lorenz (1963) equations $\frac{\partial C_O}{\partial r}$.

As can be seen in figure 1.5, the sensitivities of C_O to the parameters in the Lorenz (1963) model grow exponentially backwards in time. However, adjoint sensitivities to F_L, μ_1, μ_2, μ_3, μ_O, and μ_b agree with those calculated from perturbed forward model runs. There are no problems associated with calculating the sensitivity of passive tracers to nondynamically active variables in a chaotic system using an adjoint model. This fact has been used by Fukumori et al. (2004) to look at the origin of water masses in a GCM.

1.3.5 Studies of Eddy Resolving Ocean Models with the Adjoint Method

Studies of the applicability of the adjoint method to eddy resolving, or eddy permitting models have shown it to have varying degrees of success, depending on the integration time, and type of model being used. For integration times shorter than the eddy timescale the flow still has predictability of the first kind, so the adjoint method is expected to provide useful information. Several studies have used adjoint models for data assimilation over this timescale. Schröter et al. (1993) successfully assimilated SSH measurements into an eddy resolving quasi-geostrophic model of the Gulf Stream Extension area over a time period of 34 days using the adjoint method, Morrow and De Mey (1995) successfully assimi-

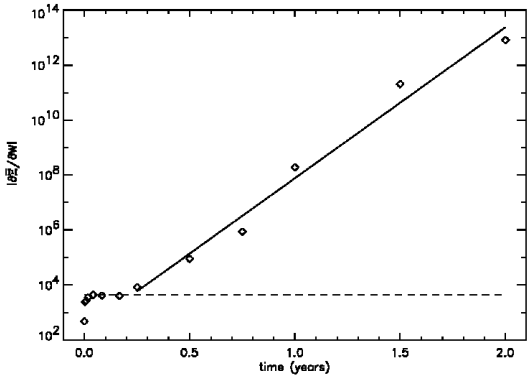

Figure 1.6: The exponential growth of adjoint gradients with increasing integration time in the shallow water model of Lea et al. (2002). The dashed line is the climate sensitivity calculated by the direct method Lea et al. (2002).

lated SSH and hydrographic data into an eddy resolving quasigeostrophic model of the Azores Current over 20 days, and Moore (1991) successfully assimilated altimeter data into an eddy resolving model of the Gulf Stream Region over 7 days.

Over longer timescales the situation is less clear, with the 3 existing studies providing different results. Lea et al. (2002) extend their analysis of the climate sensitivity of the Lorenz (1963) model to a shallow water model. They examine the sensitivity of the time mean advective transport of vorticity, $\overline{\Xi}$, across the mid-basin line of a double gyre configuration to the strength of the wind-stress curl, w, over climate timescales. As for $\overline{z}$ in the Lorenz (1963) model, $\overline{\Xi}$ clearly has predictability of the second kind and depends almost linearly on w for a wide range of values. Provided that the integration time is long enough the sensitivity $\frac{\Delta\overline{\Xi}}{\Delta w}$ can be estimated using forward model experiments, by fitting a line to the graph of $\overline{\Xi}$ against w. Lea et al. (2002) identified three regions in the adjoint sensitivities $\frac{\partial\overline{\Xi}}{\partial w}$; between 0 and 0.05 years the adjoint method underestimates the climate sensitivity as this is a shorter timescale than that needed for the model to spin up, from 0.05 to 0.20 years there is a plateau in the adjoint sensitivities where the adjoint method provides a reasonable estimate of the climate sensitivity, and for times greater than 0.2 years there is an exponential growth in the adjoint gradients (figure 1.6).

The quasigeostrophic model of Lea et al. (2002) only contains fast dynamical processes, and not slower thermodynamic processes. There are two instances where the adjoint method has been used with eddy resolving general circulation models for timescales that are long compared to the eddy timescale of the system, and these give different results. Gebbie (2004) used an adjoint model in the same form as was used for a non eddy resolving model, to assimilate data into an eddy resolving model of the eastern subtropical North Atlantic (40°W to 8°W, 12°N to 40°N). He found that the adjoint gradients appeared to be stable for an integration time of 2 years. The eddy timescale in his system is of the order of 3 to 4 months so it would be expected that the exponentially increasing adjoint gradients associated with chaotic error growth would be apparent on this timescale (G. Gebbie, personal communication). However, Köhl and Willebrand (2003) found that sensitivities calculated by the adjoint method did grow exponentially in an eddy resolving North Atlantic GCM. The reasons for the differing success of the adjoint method in these two cases is not clear, but Gebbie (2004) suggests that the dynamics in his model are only weakly nonlinear.

1.3.6 Possible Solutions

Three modifications of the adjoint method have been proposed to deal with the problem of exponentially growing sensitivities in a chaotic system, the ensemble adjoint technique of Lea et al. (2000), the coarse resolution adjoint of Köhl and Willebrand (2002), and the Fokker-Planck adjoint approach of Thuburn (2005).

Ensemble Averaging Approach of Lea et al. (2000)

Lea et al. (2000) suggested that sensitivity estimates could be improved by averaging over an ensemble of shorter integrations of the adjoint model, the timescale being set by when the plateau in the adjoint sensitivity occurs (figure 1.6). They tested this approach in the Lorenz (1963) model (Lea et al., 2000), and in a shallow water model in a double gyre configuration (Lea et al., 2002). For the double gyre configuration the ensemble adjoint approach calculated the sensitivity of the

vorticity transport across the mid-basin line to wind stress curl successfully, but it failed when applied to the variance of the vorticity transport Lea et al. (2002). This method has potential problems if it is to be used with a GCM as the shorter integrations may be too short to account for the non-local aspect of sensitivity (Thuburn, 2005). Also Eyink et al. (2004) found experimentally that at least in the Lorenz (1963) equations, the ensemble average sensitivity only converges slowly towards the finite difference sensitivity.

Coarse Resolution Adjoint Approach of Köhl and Willebrand (2002)

Köhl and Willebrand (2002) suggested a solution suitable for data assimilation. For the forward run statistical moments are calculated by averaging the solution of the high resolution model, while for the adjoint run a coarse resolution model twin linearised about the mean of the high resolution model is used. By parameterising rms variability of sea surface height in terms of density gradients, Köhl and Willebrand (2002) were able to assimilate sea surface height variability into a $1/3^\circ$ eddy permitting model of the North Atlantic, using an adjoint constructed from a 1° model. This improved the state estimation by steepening frontal structures, particularly the Azores front and the Gulf Stream, and consequently increasing variability so levels are similar to those observed (Köhl and Willebrand, 2003). However, due to the parameterisations involved in constructing the lower resolution adjoint it is unlikely that this method would provide useful information for sensitivity studies.

Fokker-Planck Adjoint Approach of Thuburn (2005)

The Thuburn (2005) solution is more complicated as it modifies the way in which the adjoint is derived from the forward model. By adding stochastic forcing to the forward model equations, it is possible to describe the evolution of the forward model probability density function using the Fokker-Planck equation, and the adjoint of this equation is then used to calculate sensitivities. The method is successful at calculating the climate sensitivity of the Lorenz (1963) model, but is not computationally feasible for a GCM.

1.4 This Thesis

Previous studies with the Lorenz (1963) equations have demonstrated that the adjoint method cannot provide useful information on climate sensitivities in simple chaotic systems (section 1.3.2). However, it is not clear to what extent these results apply to eddy resolving ocean models as the three existing studies give differing answers about the extent to which adjoint models can be used (section 1.3.5). Amongst the many unanswered questions are; over what timescale can an adjoint model can be expected to provide useful information? is this timescale the same for all quantities? and how can we tell if the sensitivity information provided by an adjoint model is useful or not?

Here we aim to address some of these questions by looking at the behaviour of sensitivities calculated by an adjoint model in an eddy resolving general circulation model of a zonally reentrant channel. Results are compared with those from a non eddy resolving channel and a model where the eddies are damped by the use of a parameterisation scheme. The idealised configuration makes interpretation of the adjoint sensitivities easier, and the small domain size allows the adjoint model to be run many times with different output functions and makes comparison with perturbed forward model experiments practical.

The model setup and the circulation in the forward model are described in chapter 2. Chapter 3 looks at the timescale over which we can still extract useful information from the adjoint model results, and how we can see if there is useful information contained in the adjoint sensitivity. Chapter 4 looks at the physical interpretation of the sensitivities calculated by the adjoint model and their spatial structure, and conclusions and outlook are given in chapter 5.

1.5 Summary

- Adjoint models are a powerful and efficient tool for sensitivity analysis in coarse resolution ocean models.
- Studies with the Lorenz (1963) equations have shown that sensitivities cal-

culated by an adjoint model only contain useful information on a finite timescale in a chaotic system. This also applies to the sensitivity of time averaged climate quantities, which have predictability of the second kind.

- Previous studies using adjoint models with eddy resolving ocean models have been inconclusive about the timescale over which adjoint models can be used successfully, and to what extent the results from the Lorenz (1963) model are relevant to an ocean GCM.

Chapter 2

Description of the model

2.1 Introduction

The model used for this study is the Massachusetts Institute of Technology general circulation model (MITgcm) and its adjoint. The MIT model was chosen for this study primarily because it has a flexible adjoint maintained alongside the forward model code. The model code is freely available through the MITgcm website http://mitgcm.org/, and is described in detail in Marshall et al. (1997a,b), and in the manual that accompanies the code. This chapter begins with a summary of the MITgcm and its adjoint, before describing the exact model configurations used in this study and the circulation in the forward models, with a particular emphasis on the levels of eddy kinetic energy.

2.2 MITgcm

The MITgcm is a primitive equation ocean general circulation model (OGCM), which solves the equations for the ocean circulation under the Boussinesq approximation, where the density is assumed to be constant unless density variations give rise to buoyancy forces. The MITgcm is novel in that it can be run in both hydrostatic and nonhydrostatic modes. The hydrostatic mode is used here so

that the equations solved by the model are,

$$\frac{D\mathbf{v_h}}{Dt} + \frac{1}{\rho_0}\nabla_h p + 2\Omega \times \mathbf{v_h} - A_h \nabla_h^4 \mathbf{v_h} - A_v \frac{\partial^2 v_h}{\partial z^2} = \frac{1}{\rho_0}\nabla \cdot \tau \quad (2.1)$$

$$\frac{\partial \eta}{\partial t} - \nabla \cdot \int_{-H}^{0} \mathbf{v_h}\, dz = 0 \quad (2.2)$$

$$\frac{dp}{dz} = -g\rho \quad (2.3)$$

$$\rho = \rho(T, S, z) \quad (2.4)$$

$$\frac{DT}{Dt} - \kappa_h \nabla_h^4 T - \kappa_z \frac{\partial^2 T}{\partial z^2} = \nabla \cdot Q_T \quad (2.5)$$

$$\frac{DS}{Dt} - \kappa_h \nabla_h^4 S - \kappa_z \frac{\partial^2 S}{\partial z^2} = \nabla \cdot Q_S \quad (2.6)$$

$$(2.7)$$

where $\mathbf{v_h}$ is the horizontal velocity, p is the deviation in pressure from that of an ocean at rest with density ρ_0, Ω is the rotation of the earth, A_h is the (biharmonic) horizontal viscosity, A_v is the vertical viscosity,τ is the surface windstress, η is the elevation of the implicit free surface, T is the potential temperature, S is the salinity, H is the depth of the ocean, κ_h is the horizontal diffusivity, κ_v is the vertical diffusivity, and Q_T is the surface heat flux.

The equations are discretised using a z-level coordinate system in the vertical, where the vertical coordinate is the distance from the ocean surface, and an Arakawa C-grid, in the horizontal (figure 2.1 left). A quasi second order Adams-Bashforth scheme is used for the time discretisation. Any choice of coordinate system has its advantages and limitations, a z-level coordinate system allows easy representation of the horizontal pressure gradient and equation of state in a Boussinesq fluid, but it is difficult to represent the bottom topography and isopycnal diffusivity (Griffies et al., 2000). Arakawa C-grids are usually the preferred horizontal grid in models that resolve the Rossby radius as they represent inertia-gravity waves well. However, the u and v velocity points are staggered in space on a C-grid so that calculating the Coriolis terms involves spatial averaging. If the resolution of the grid is low compared to the Rossby radius, this spatial averaging leads to false minima in the dispersion relation and regions where the group velocity has the wrong sign. Energy is thus fed into short scale perturba-

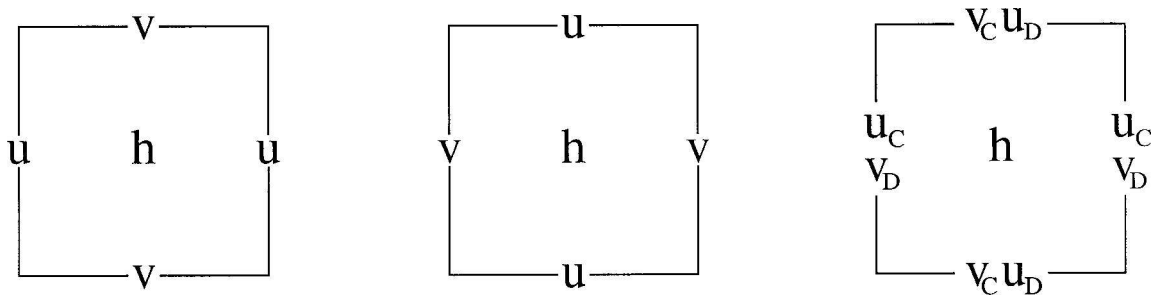

Figure 2.1: Arrangement of the model variables on the C grid (left) the D grid (centre0 and the C-D grid (right). From Adcroft et al. (1999).

tions and there is a checkerboard pattern of computational noise (Adcroft et al., 1999). For this reason it has been traditional to run models that do not resolve the Rossby radius on a grid where the velocity points coincide.

In the MITgcm this problem is solved by using a C-D coupling scheme. A D-grid is overlaid on the C-grid so that the u velocities on the C-grid coincide with the v velocities on the D grid (figure 2.1) right. The velocities on the C-grid are used for calculating the Coriolis terms on the D-grid and vice versa, thus removing the need for the spatial averaging (Adcroft et al., 1999). The D-grid velocities are relaxed back to the C-grid velocities with a timescale τ_{CD}. The C-D scheme effectively doubles the resolution for the velocity, but not for the tracers, so that half the model kinetic energy is associated with the D-grid velocities (Adcroft et al., 1999). This scheme is not required when the Rossby radius is resolved, so it is not normally used at this resolution, although the dispersion relation on the C-D grid is similar to that on the C grid and it should give similar results.

When using the hydrostatic mode a convective adjustment scheme is necessary to maintain static stability in the model. This scheme checks adjacent model levels for static stability and if they are found to be statically unstable it mixes them so that they have the same salinity and potential temperature.

2.3 MITgcm Configuration

2.3.1 Model Domain and Resolution

The basic configuration is a flat bottomed zonally reentrant channel, 9.6° by 9.6° by 5000m, with the southern wall at 45°N. The model is run at 2 resolutions, $0.1° \times 0.1°$, corresponding to a maximum grid spacing of 10km at 45°N, and $0.48° \times 0.48°$, corresponding to a maximum grid spacing of 50km. In order to allow chaotic eddies to form the first baroclinic Rossby radius, $a = c/f$, where f is the Coriolis parameter and c is the Rossby wave speed, must be resolved. In the 0.1° model the 1st baroclinic Rossby radius ranges from 32km at 45°N to 20km at 54.5°N, and is resolved everywhere so that the model is termed eddy resolving. The 0.48° model is non eddy resolving, the first baroclinic Rossby radius ranges from 40km at 45°N to 20km at 54°N, and is only resolved north of 50°N. At both resolutions there are 19 levels in the vertical, ranging in thickness from 30m at the surface to 550m at depth.

Although 0.1° resolution is sufficient for eddies the form and the model to exhibit the chaotic behaviour of interest in this study (Hogan and Hurlburt, 2000) found that levels of kinetic energy increased in their model of the Japan sea, untie the resolution was 1/32°. It is thus likely that increasing the resolution beyond 0.1° would also increase the eddy kinetic energy in the high resolution model used here. However, this would significantly increase the integration time making many the integrations of the forward and adjoint models, which are a requirement of this study computationally impractical.

Biharmonic horizontal viscosity and diffusion are used for the lateral dissipation of momentum and lateral mixing of temperature, respectively as they are more scale selective than Laplacian friction or mixing and cause less damping of mesoscale eddies for the same damping of computational noise (Semtner and Mintz, 1977). To ensure that the integrations at different resolutions have equivalent values of mixing and viscosity, the parameters are scaled so that $A_{0.48} = A_{0.1} \left(\frac{\Delta x_{0.48}}{\Delta x_{0.1}}\right)^4 \frac{\Delta t_{0.1}}{\Delta t_{0.48}}$ where $A_{0.48}$ is the mixing/viscosity in the 0.48° model, $A_{0.1}$ is the mixing/viscosity in the 0.1° model, Δx is the grid spacing and Δt is the

Parameter	Coarse Resolution	High Resolution
Resolution	$0.48^\circ \times 0.48^\circ$	$0.1^\circ \times 0.1^\circ$
Grid	$20 \times 20 \times 19$	$96 \times 96 \times 19$
Time step	2400s	500s
Biharmonic horizontal viscosity	$5.5 \times 10^{12}\mathrm{m}^4\mathrm{s}^{-1}$	$5 \times 10^{10}\mathrm{m}^4\mathrm{s}^{-1}$
Biharmonic horizontal diffusivity of heat and salt	$1.1 \times 10^{12}\mathrm{m}^4\mathrm{s}^{-1}$	$1 \times 10^{10}\mathrm{m}^4\mathrm{s}^{-1}$
Vertical viscosity	$2 \times 10^{-3}\mathrm{m}^2\mathrm{s}{-1}$	$2 \times 10^{-3}\mathrm{m}^2\mathrm{s}^{-1}$
Vertical diffusivity of heat and salt	$1 \times 10^{-4}\mathrm{m}^2\mathrm{s}^{-1}$	$1 \times 10^{-4}\mathrm{m}^2\mathrm{s}{-1}$
Quadratic bottom friction	0.001	0.001
CD coupling time scale (damped model only)	172800s	86500s
SST restoring time scale	86500s	86500s

Table 2.1: Model parameters.

timestep. As the vertical resolution is the same at both resolutions the vertical viscosity and vertical diffusivity are kept constant. More sophisticated schemes for isopycnal mixing are not used as they tend to be highly nonlinear and were found to cause the adjoint model to become unstable after a few timesteps.

Quadratic bottom drag and no slip boundary conditions are applied at the bottom boundary. Bottom friction is necessary for the barotropic energy to reach equilibrium, and because with low bottom friction large scale meanders are favoured over mesoscale eddies (Rivière et al., 2004). Free slip boundary conditions are applied at the side boundaries.

The C-D coupling scheme described in section 2.2 is used in all the experiments with the 0.48° model. Although it is not necessary to include this scheme in an eddy resolving model it was suggested that the C-D scheme may lead to smoother velocity fields and a better behaved adjoint solution in an eddy resolving model (Patrick Heimbach, personal communication), so for some of the 0.1° models runs the CD scheme is included. Eddies in these runs are less vigorous, so this model is referred to as the damped 0.1° model, while the model without the CD scheme is referred to as the undamped 0.1° model.

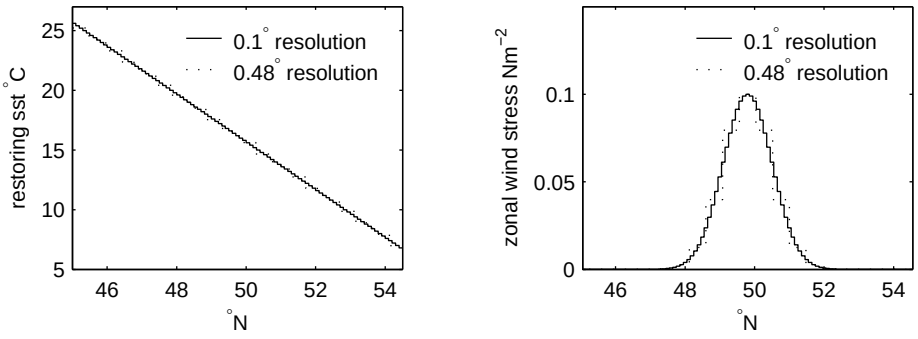

Figure 2.2: Model forcing fields. Left, SST*. Right, zonal windstress τ_x.

2.3.2 Model Surface Forcings

The model is forced at the surface with a temporally and zonally constant zonal windstress, which is sharply peaked in the centre of the channel,

$$\tau_x(\phi) = \tau_{x_0} \exp\left(-(\phi - \phi_m)^2\right) \tag{2.8}$$

where τ_{x_0} is the maximum windstress, ϕ is the latitude and ϕ_m is the latitude at the centre of the channel (figure 2.2). Initially τ_{x_0} is set to 0.1Nm^{-2}, which is close to the winter mean for 55°N (Peixoto and Oort, 1992).

For the thermal forcing, restoring boundary conditions are used, where the heat flux into the ocean is given by,

$$Q_T = (\text{SST}^* - T_1)/\tau_T \tag{2.9}$$

where SST* is the apparent atmospheric temperature seen by the ocean surface, after the effects of solar radiation and evaporation on the heat flux have been taken into account (Haney, 1971), τ_T is the restoring timescale, and T_1 is the temperature of the top model layer. The heat content of the model is thus free to vary over time as the heat flux into the model depends on the state of the surface layer.

In the current case SST* is zonally constant and varies linearly with latitude (figure 2.2),

$$\text{SST}^*(\phi) = \text{SST}^*_0 + (\phi - \phi_m)\, B \tag{2.10}$$

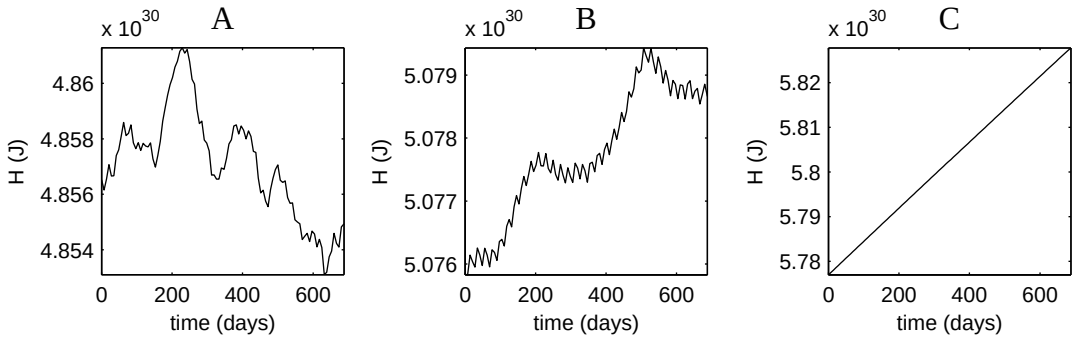

Figure 2.3: Heat content, after the 20 year spin up, as a function of time. A) the 0.1° model B) the damped 0.1° model and C) the 0.48° model. Note the different scales.

where SST_0^* is the apparent atmospheric temperature in the centre of the channel, and B is the meridional gradient of SST^*. The restoring timescale of 1 day is relatively short, and causes significant damping of temperature anomalies at the surface (section 2.3.4). A short τ_T is necessary to maintain the meridional temperature gradients in the model. Other studies using GCMs to study the flow in a channel have either restored to climatological fields over a longer timescale but over the entire interior of the model domain (Alves and de Verdière, 1999), or have applied sponge boundary conditions at the side boundaries. Restoring over the entire domain is inappropriate here as we are interested in the sensitivity of climate quantities to changes in the forcing, while sponge boundary conditions would make the results harder to interpret. Initially $\mathrm{SST}_0^* = 16°\mathrm{C}$ and $B = -2°\mathrm{C}°N^{-1}$. Both these values are rather higher than observed in the real ocean where the zonal mean temperature is 8°C at 50°N and B is $-0.4°\mathrm{C}°N^{-1}$ (Peixoto and Oort, 1992), but were chosen so that the model would be highly unstable. To simplify the analysis, salinity is kept constant at 35 p.s.u.

2.3.3 Model Initialisation

A perturbation is added to the initial conditions to allow eddies to develop more quickly, and both models are spun up for 20 years prior to the start of the adjoint integrations. This is not long enough for the model to achieve full thermodynamic

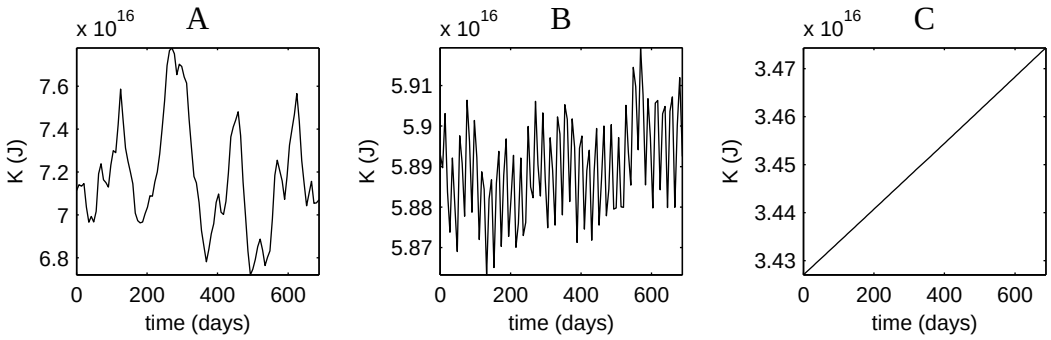

Figure 2.4: Kinetic energy, after the 20 year spin up, as a function of time. A) the 0.1° model B) the damped 0.1° model and C) the 0.48° model. Note the different scales.

equilibrium, but it should be long enough for the velocity field to adjust to the density structure. In the 0.48° model there is an obvious drift in the heat content (figure 2.3C) and kinetic energy (figure 2.4C), but the drift in both the damped and undamped 0.1° models is much smaller (figure 2.3A,B and figure 2.4A,B) .

As the aim of this thesis is to look at the usefulness of a particular method for calculating climate sensitivities and how this relates to the predictability of the system, and not to examine its equilibrium behaviour, the drift of the model is not considered a problem.

2.3.4 Forward Model State

Due to the short timescale used for restoring to SST*, deviations of the temperature from SST* are strongly damped in the surface layer in all models (figure 2.5). Away from the surface layer in both the undamped and damped 0.1° model there is a sharp front, and associated zonal jet, centred on 49.5°N. Across the front the temperature falls by around 5°C over 1° latitude between 50 and 500m (figure 2.5). The front meanders between with a wavelength equal to twice the length of the channel (figure 2.6). In the 0.48° model the front is broader due to the lower resolution, and the temperature change across the front is only 8°C over 3° latitude, between 50 and 500m (figure 2.5 right).

The kinetic energy in the channel is sharply peaked at the position of the

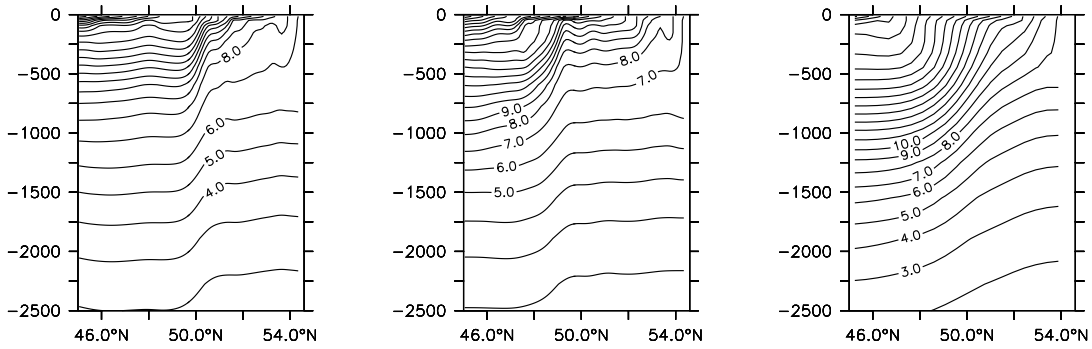

Figure 2.5: Instantaneous temperature at 690 days after the 20 year spin up. Left: undamped 0.1° model at 4.75°E; centre: damped 0.1° model at 4.75°E; right: 0.48° model at 4.56°E. Contour interval is 1°C.

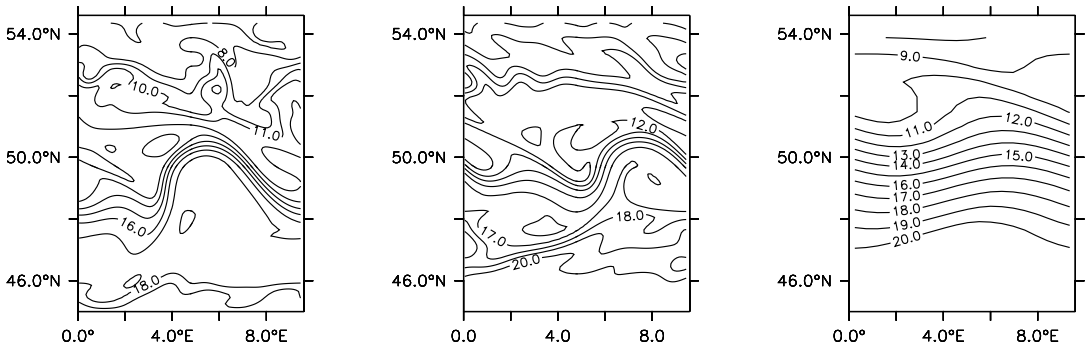

Figure 2.6: Instantaneous temperature at 95m, at 690 days after the 20 year spin up. Left: undamped 0.1° model; centre: damped 0.1° model; right: 0.48° model. Contour interval is 1°C

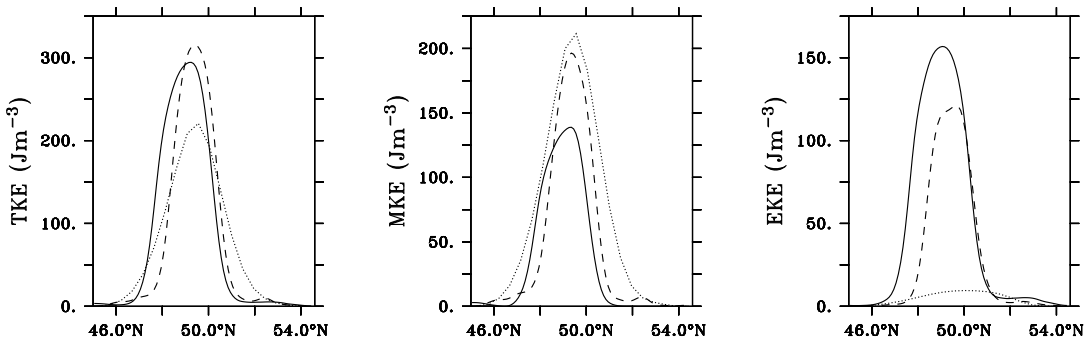

Figure 2.7: Zonally averaged kinetic energy per unit volume at 95m. Left the time mean total kinetic energy TKE = $\overline{uu}+\overline{vv}$. Centre, kinetic energy due to the mean flow MKE = $\bar{u}\bar{u} + \bar{v}\bar{v}$. Right, eddy kinetic energy EKE = TKE − MKE. Solid line: undamped 0.1° model; dashed line: damped 0.1° model; dotted line: 0.48° model. All time averages are calculated over the 690 period shown in figure 2.4.

front where there is a zonal jet, 49° in the undamped 0.1° model and 49.5° in the damped 0.1° and 0.48° models (figure 2.7 left). The peak is broader in the 0.48° model, due to the less steep front, and kinetic energy decreases more rapidly with depth (not shown), so that the total kinetic energy in the 0.48° model is half that of the 0.1° models (figure 2.4). Above 500m all three models have similar levels of kinetic energy although there are major differences in the way it is partitioned between the mean and time varying parts of the flow. In the damped 0.1° model more of the kinetic energy is contained in mean flow (figure 2.7 centre) and less in the eddying field (figure 2.7 right) than in the undamped 0.1° model. Over all depths there is a 50% reduction in the eddy kinetic energy in the damped 0.1° model, compared to the undamped 0.1° model. This is seen in the snapshots of the temperature field at 95m, where there are more eddies present in the undamped 0.1° model than in the damped 0.1° model (figure 2.6). In the 0.48° model the grid resolution is too coarse to permit eddies to form; almost all the kinetic energy is due to the time mean flow (figure 2.7 left and right) and no eddies are visible in the snapshots of the temperature field (figure 2.6 right).

At 52°N there is a local maximum in the kinetic energy due to the mean flow in the damped 0.1° model (figure 2.7 centre). This is due to a much weaker secondary

zonal jet at this latitude, which can also be seen in the plots of instantaneous temperature (figure 2.6). This peak is not present in the undamped 0.1° model, but there is slower decay of eddy kinetic energy with latitude to the north of the front than south of it (figure 2.7). In the undamped 0.1° model, most of the eddies are formed north of the front (figure 2.6 left). Alves and de Verdière (1999) found a similar asymmetry in their model of the Azores Front. A water parcel of a given volume, between two isopycnals south of the front has higher relative vorticity and is thicker than one north of it, as, away from the front, the isopycnals in the surface layer in the basic state are nearly horizontal (figure 2.8). A water parcel moving northwards gains negative relative vorticity, and decreases in thickness, while one moving south gains positive relative vorticity and increases in thickness. This leads to larger and weaker anticyclones north of the jet, and smaller stronger cyclonic filaments south of it. The cyclonic filaments south of the jet are quickly dissipated by the subgridscale dissipation in the model, while the anticyclonic filaments north of it persist long enough to form pinched off eddies (Alves and de Verdière, 1999).

The varying levels of eddy activity can also be seen in the fluctuation in time of the kinetic energy and heat content of the channel. In the undamped 0.1° model the kinetic energy fluctuates about the mean by around 10% compared to 1% in the undamped model (figure 2.4) while in the heat content fluctuates by around 0.2% in both the damped and undamped 0.1° models. The 0.48° model has clearly failed to reach equilibrium, but even after detrending there are no obvious fluctuations of the kinetic energy or heat content with time.

Thus although the general circulation in the three models is similar they have varying amounts of eddy activity and allow the performance of the adjoint method to be studied in systems that are expected to be highly chaotic, the undamped 0.1° model, less chaotic the damped 0.1° model, and non chaotic, the 0.48° model.

An estimate of the speed at which disturbances in the flow grow is given by the Eady model of baroclinic instability. The growth rate of the most unstable mode is

$$\sigma_{max} = 0.3098 \frac{f}{N} \frac{\mathrm{d}U}{\mathrm{d}z} \tag{2.11}$$

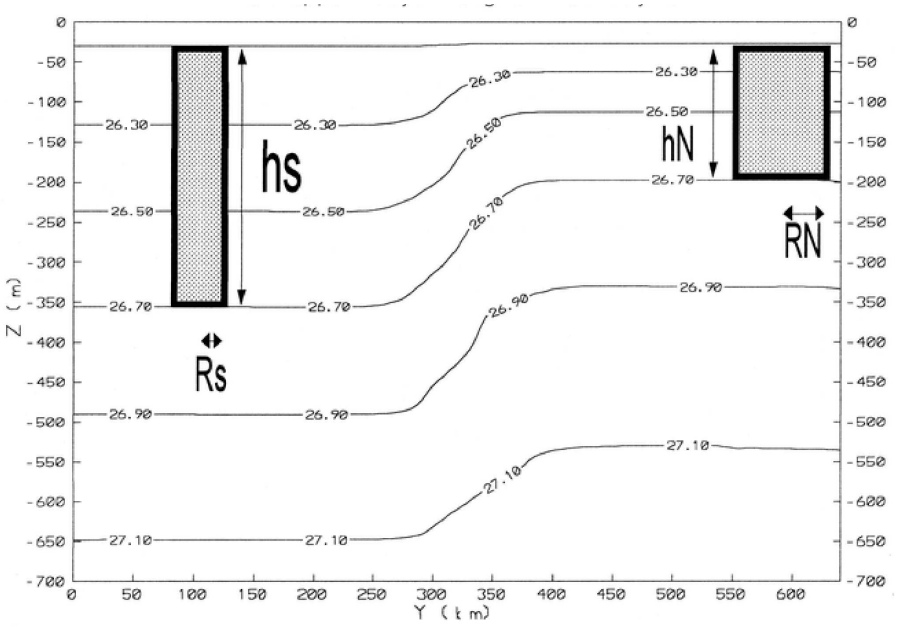

Figure 2.8: The distribution of isopycnals across the front in Alves and de Verdière (1999)'s model of the Azores Front. Near the surface the isopycnals are nearly horizontal, so that a water parcel south of the front has a greater thickness, and lower volume than one north of it. From Alves and de Verdière (1999).

where f is the Coriolis parameter, $N = \frac{-g}{\rho_0}\frac{d\rho}{\mathrm{d}z}$ is the buoyancy frequency, and U is the mean velocity (Gill, 1982, p557). In all models the core of the jet is the most baroclinically unstable part of the flow, with an e-folding time for the instabilities, $1/\sigma_{max}$, of around 6-days (figure 2.9). Along the boundaries and at depth this increases to around 5000 days.

2.4 The MITgcm Adjoint and Tangent Linear Models

2.4.1 Generation of the Adjoint and Tangent Linear Models

There are two possible methods for generating an adjoint or tangent linear model, *linearise then discretise* where the adjoint and tangent linear equations are first derived from the continuous formulation of the nonlinear forward model and then discretised, or *discretise then linearise* where the nonlinear forward model is first

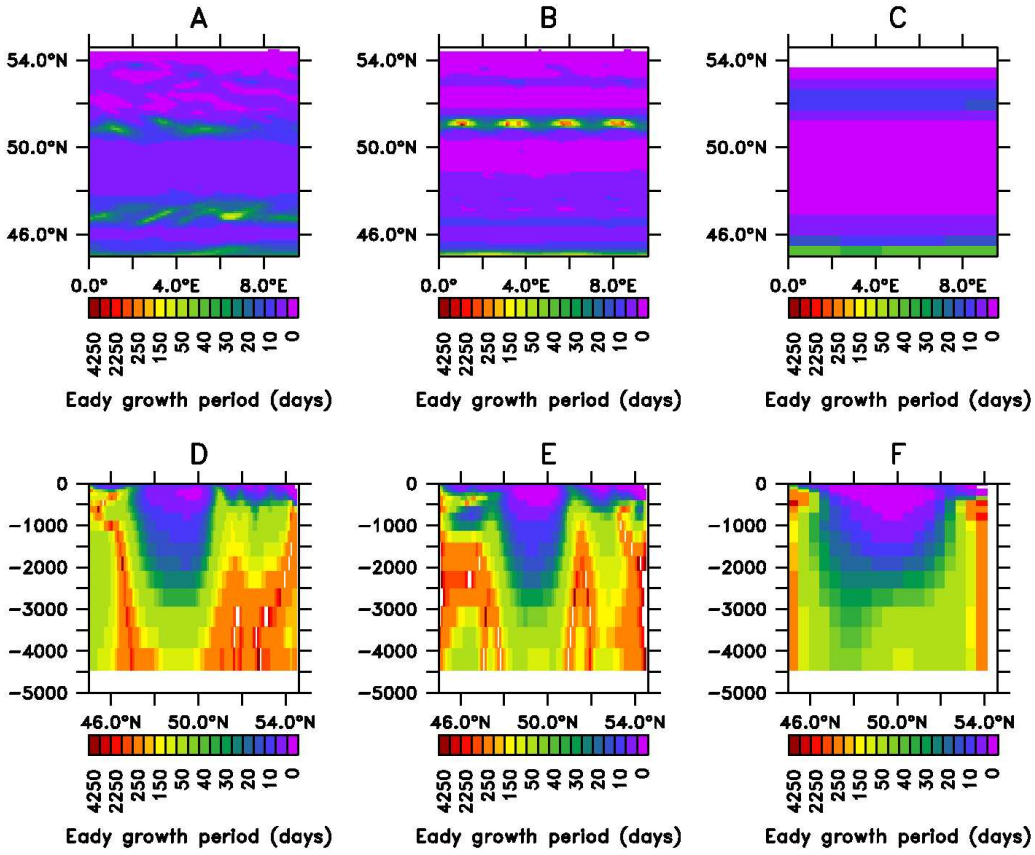

Figure 2.9: Eady growth period in days calculated from the time averaged velocity and temperature fields, for the 690 days after the 20 year spin up. Top row (A-C): Eady growth period at 50m. Bottom row (D-F): Eady growth period at 4.6°E. Left column (A,D) shows the undamped 0.1° model, middle column (B,E) shows the damped 0.1° model, right column (C,F) shows the 0.48° model. Note that the colour scale is not linear.

discretised and the adjoint and tangent linear equations are either derived from the discretised equations or from the computational code. The MITgcm uses the discretise linearise approach. The discretise then linearise approach gives the exact adjoint to the model code, resulting in a sensitivity that is correct to within computational round off error, while the adjoint-discretise approach is only accurate to within the accuracy of the finite difference approximations (Sirkes and Tziperman, 1997). However, in some instances the linearise then discretise method may be preferable. Sirkes and Tziperman (1997) observed an unphysical computational mode due to the leapfrog time differencing scheme in the solution of an adjoint model generated using the discretise then linearise method but not in that generated by the linearise then discretise method. The leapfrog scheme is not used in the MITgcm and so this is not a problem for this study. Zhu and Kamachi (2000) showed that timestepping schemes that are unconditionally stable in a nonlinear model may be only conditionally stable in a TLM requiring a much shorter timestep to be used in a TLM; this is a particular problem with some parameterisations for vertical diffusivity, and is the reason why a constant vertical diffusivity is used in this study. The linearise then discretise method allows the use of an approximate adjoint model, such as that of Schiller and Willebrand (1995), which considerably reduces the computational cost of running the adjoint model.

The main advantage of the discretise then linearise approach is that it enables automatic differentiation software to be used to generate the adjoint code, without which developing an adjoint model is extremely time consuming, and changes to the forward model code may require the adjoint model to be completely rewritten. The MITgcm adjoint and tangent linear models are generated using the free access automatic differentiation software TAMC (Giering, 1999; Giering and Kaminski, 1998), and later its commercial successor TAF (Giering, 2005) which allows the adjoint model to be maintained alongside the forward model code (Marshall et al., 2004). The principal advantage of TAF over TAMC is speed. Generating and compiling the adjoint model takes around 1 hour using TAMC but only 10 minutes using TAF. The adjoint code generated using TAF

also allows the adjoint integration to be split so that it can fit into a computer queue, while the TAMC generated code does not (Marshall et al., 2004). There is also a bug in the TAMC generated tangent linear model code that causes the forward model to crash due to counters for an iterative loop in the routine used to solve for the pressure field not being reset correctly. This is corrected by hand (Patrick Heimbach personal communication). The code generated by TAF and TAMC is otherwise identical. Both TAMC and TAF deal with nondifferentiable conditional statements in the forward model by requiring the adjoint model to take the same branch of the code as the forward model (Giering and Kaminski, 1998). As the convective adjustment scheme contains conditional statements this means that if convection occurs in the forward model, it also occurs in the adjoint model, so the adjoint model does not calculate the part of sensitivity due to a possible increase or decrease in convection.

2.4.2 Checkpointing Scheme and Computational Requirements

The adjoint model requires the state of the forward model in the reverse order to which it is computed. Storing the forward model state at every timestep would require prohibitively large amounts of memory, while recomputing at every timestep would require prohibitively large amounts of computer time. The MITgcm adjoint employs a checkpointing scheme with three loops, which reduces the amount of storage required without involving excessive recomputations (figure 2.10). The resulting computational cost of the adjoint model is only around 5 times that of a single forward model integration (Griewank, 1992).

Using $n_1 = n_2 = 65$, $n_3 = 30$, a 690 day integration of the 0.1° adjoint model uses 4Gb of disk space and 2.5Gb of RAM. The adjoint model was mainly run on 4 processors of the Sun SMP server at the MPI for Meteorology in Hamburg, where a 690 day integration of the adjoint model took about 7 days. A centred difference gradient check of the adjoint model results requires 2 further integrations of the forward model. This significantly increases the run time, if several gradient

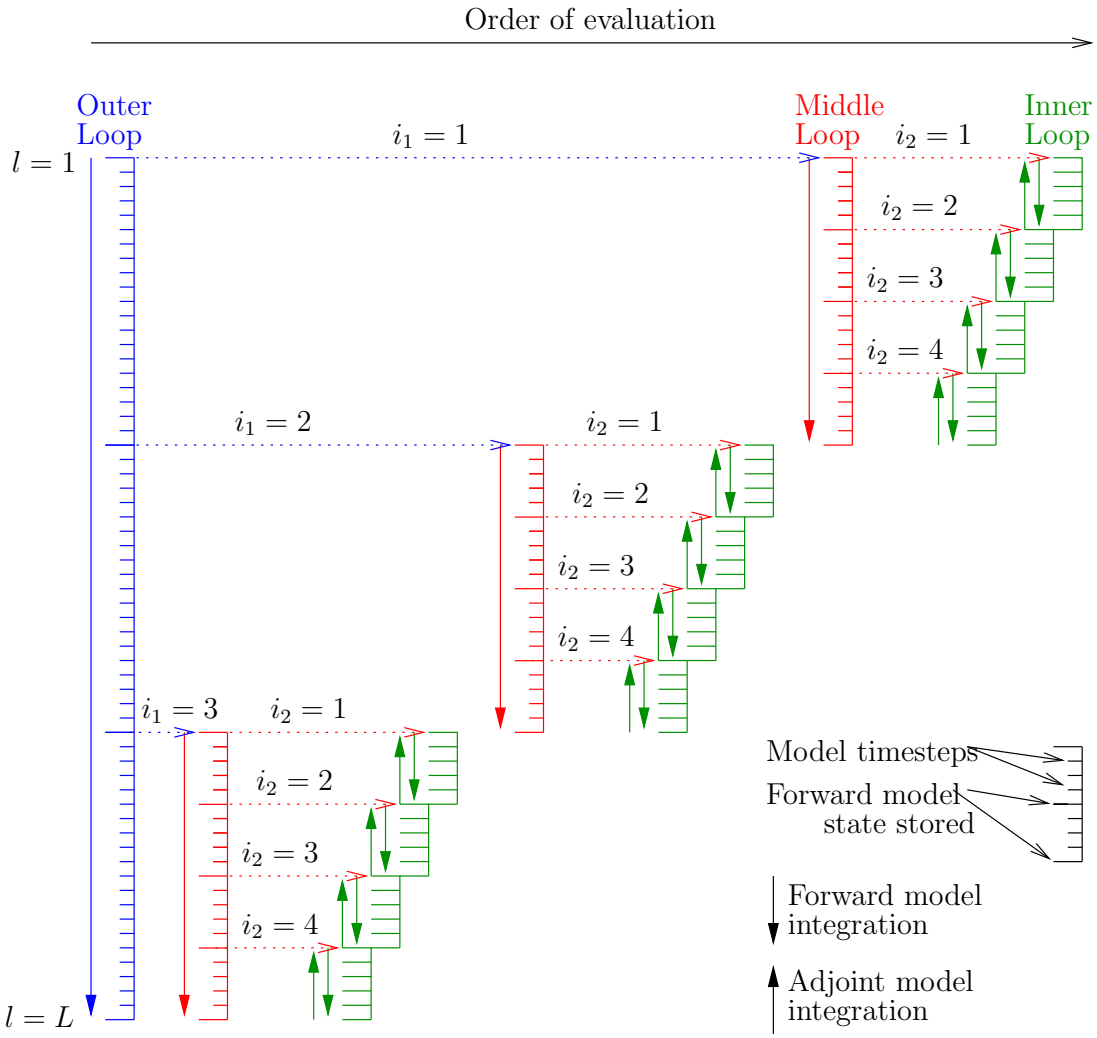

Figure 2.10: Schematic of the checkpointing scheme in the MIT model. The outer loop is divided into n_1 segments and runs over the entire integration time of L timesteps. The model state is written to disk when $l = (i_1 - 1)L/n_1$, where $i_1 = 1, 2, \ldots, n_1$. The middle loop is divided into n_2 segments and runs between timesteps $l = (i_1 - 1)L/n_1$ and $l = i_1 L/n_1$, with the first integration of the middle loop starting from $l = (n_1 - 1)L/n_1$. The model state is written to disk when $l = (i_1 - 1)L/n_1 + (i_2 - 1)L/(n_1 n_2)$, where $i_2 = 1, 2, \ldots, n_2$. The inner loop contains n_3 timesteps and the model state is stored in memory at every timestep. The adjoint model is then integrated backwards in time over the n_3 timesteps of the inner loop (Marotzke et al., 1999; Marshall et al., 2004). For illustration the simple case $n_1 = 3$, $n_2 = 4$, $n_3 = 5$ is shown here. Adapted from Marshall et al. (2004, p.248)

checks are required. Most of the forward model experiments were computed at the Rechenzentrum Garching, but the difficulty of splitting the adjoint model integrations using the TAMC generated code, meant that this was impractical for the longer adjoint model runs.

2.5 Summary

- The basic model is a flat bottomed zonally reentrant channel, 9.6° by 9.6° with the southern boundary at 45°N. The model is forced at the surface by restoring boundary conditions for temperature, and by zonal windstress.
- The model is run at 2 different resolutions 0.1° and 0.48°. Two versions of the 0.1° model are run, damped and undamped. These differ only in the inclusion of the CD coupling scheme in the damped model.
- The different levels of eddy activity in the undamped 0.1°, damped 0.1° and 0.48° models allow the usefulness of the adjoint method to be studied in systems that are highly chaotic, slightly chaotic and non chaotic respectively.

Chapter 3

Time Limits on the Applicability of the Adjoint Model

3.1 Introduction

The aim of this chapter is to establish a timescale over which useful information remains in the sensitivities calculated by the adjoint model, and to look at methods of assessing whether sensitivities calculated by the adjoint model contain useful information. We begin by looking at the results from integrations of the tangent linear model, which allow us to define an approximate maximum timescale over which we expect the linearisation to be valid in the most chaotic undamped 0.1° model. The adjoint model is then run over this timescale (around 280 days), and again over a much longer timescale of around 690 days, on which timescale the linear approximation in the undamped 0.1° model is clearly no longer valid. The usefulness of the information in the sensitivities calculated by the adjoint model is then assessed by comparing the adjoint model results with finite difference gradient checks, and forward model experiments where a larger spatial scale perturbation to the forcing has been made. The time evolution of the adjoint sensitivities is also examined.

The two output functions considered in this chapter are the time mean total heat content $\overline{H}$ and the time mean total kinetic energy $\overline{K}$. Time averaged inte-

gral quantities are chosen as we wish to look at the upper limit at which useful information is retained in sensitivities calculated by the adjoint method. It is thought that this upper limit may be longer for integral quantities that clearly have predictability of the 2nd kind (Lorenz, 1975), even on a relatively short timescale of a couple of years. Use of $\overline{H}$ and $\overline{K}$ allows a comparison between the behaviour of dynamic and more slowly varying thermodynamic variables. The majority of this work is covered in McLay and Marotzke (2006).

3.2 Validity of the Linearisation

3.2.1 Tangent Linear Model Results

Although it is known that there is a limit of applicability of the adjoint method in any chaotic system, it is not clear what this is in any given model. The solution of the TLM can be compared with the difference between 2 nonlinear model integrations to give an indication of the timescale over which the linearisation required for the adjoint model is valid (Errico and Vukicevic, 1992; Kleist and Morgan, 2005). Following Kleist and Morgan (2005) a 1°C perturbation was applied to the initial temperature field at the grid cell that showed the highest adjoint sensitivity of $\overline{H}$ to the temperature at the initial time (4.56°E, 50.52°N, 2703m for the 0.48° model, 5.35°E, 52.55°N, 324m for the damped 0.1° model, and 2.8°E, 51.9°N, 324m for the undamped 0.1° model). The location of the maximum sensitivity of $\overline{K}$ to the temperature at the initial time is in approximately the same location.

The trajectories in the two perturbed non linear model integrations diverge in both the undamped 0.1° model and the damped 0.1° model, but the difference between the trajectories saturates at a higher value in the undamped 0.1° model due to the more vigorous eddies. Exponential divergence of nearby trajectories is a characteristic of chaotic dynamics (Strogatz, 1994), but cannot be seen between the two perturbed non linear integrations in either model due to saturation. The tangent linear solution shows clear exponential growth in the undamped 0.1°

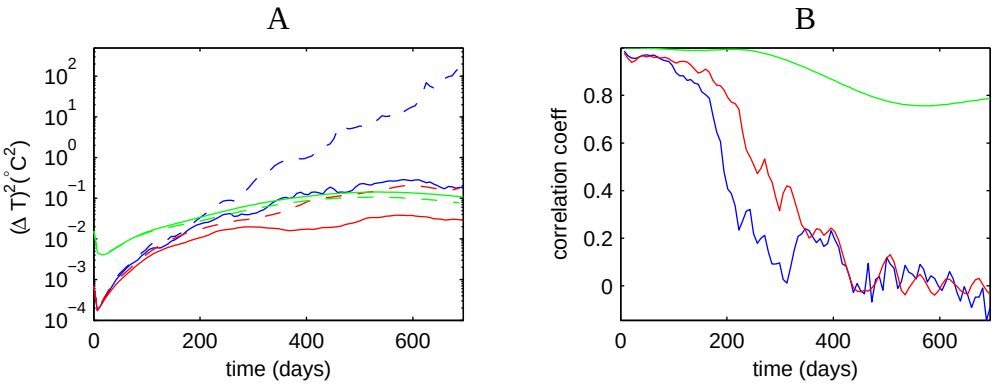

Figure 3.1: A) Divergence of the potential temperature calculated from two integrations of the nonlinear model (solid lines) and calculated by the tangent linear model (dotted lines). B) Correlation between the perturbation potential temperature in the difference between 2 integrations of the non linear forward model and the TLM, at 50m. Undamped 0.1° (blue) damped 0.1° model (red) and 0.48° model (green).

The divergence is defined as, $\frac{\sum_{ijk}\left(V_{ijk}\Delta T_{ijk}\right)^2}{\sum_{ijk} V_{ijk}}$ where ΔT_{ijk} is the perturbation temperature of the grid cell with index ijk and V_{ijk} is its volume. The correlation is defined as $\frac{\sum_{ij} \Delta T_{TLM_{ij}} \Delta T_{NLM_{ij}}}{\sqrt{\left(\sum_{ij} \Delta T_{TLM_{ij}} \Delta T_{TLM_{ij}} \sum_{ij} \Delta T_{NLM_{ij}} \Delta T_{NLM_{ij}}\right)}}$, where ΔT_{TLM} is the perturbation calculated by the TLM, and ΔT_{NLM} is the difference between the two nonlinear model integrations. The perturbed non linear model and the TLM are initialised with a 1°C perturbation to the initial potential temperature at the grid cell that shows the highest sensitivity in the adjoint model runs, see text. Other state variables behave similarly.

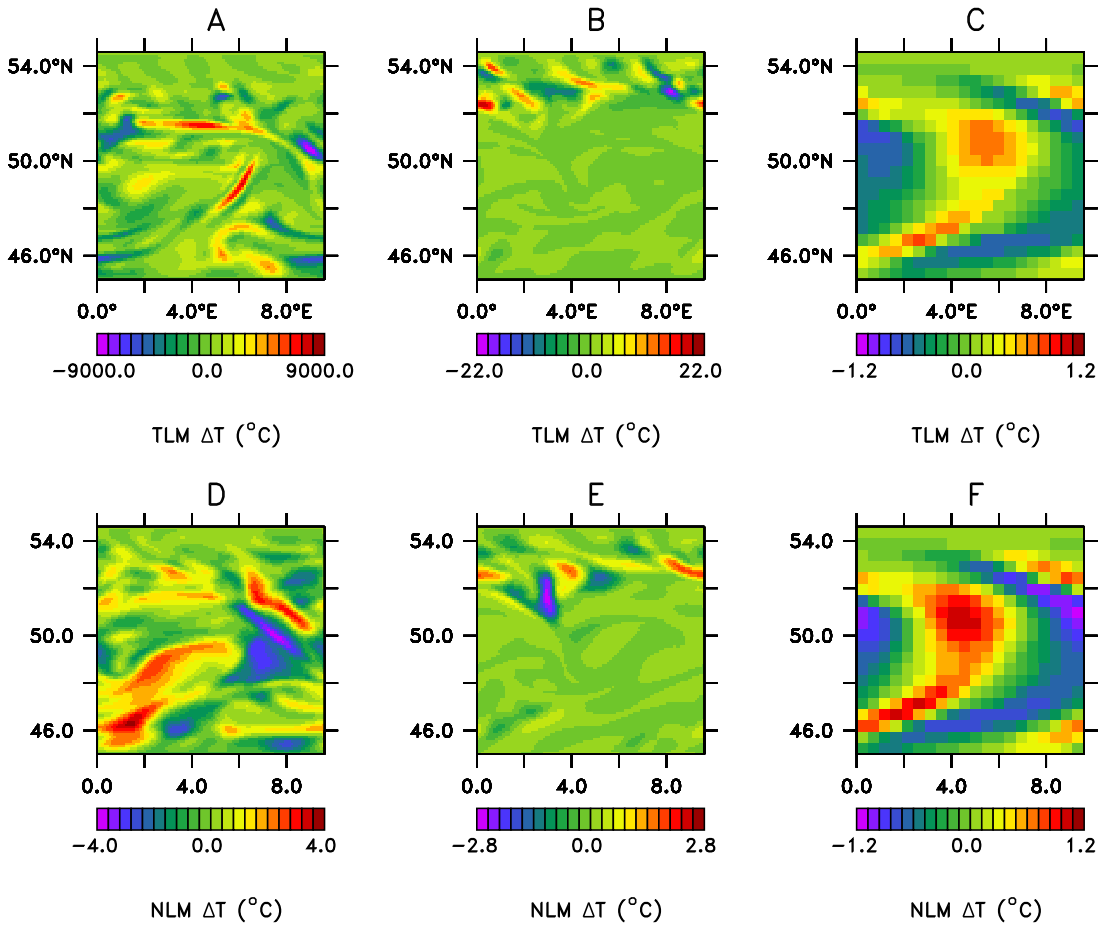

Figure 3.2: Perturbation temperature at 50m after 694 days. The top row (panels A-C) shows the results from the TLM, the bottom row (D-E) shows the difference between 2 integrations of the nonlinear forward model. Left column (A,D) shows the undamped 0.1° model, middle column (B,E) the damped 0.1° model, and the right column (C,F) the 0.48° model. Note the difference in scales.

model and grows more slowly in the damped 0.1° model (figure 3.1A). In the damped 0.1° model the growth of the tangent linear solution is much slower. However, after 690 days, the perturbation variables in the tangent linear solution are still an order of magnitude greater than in the difference between the nonlinear forward models. The integration time is insufficient to determine whether the TLM solution is growing exponentially in the damped 0.1° model, which would be indicative of chaos, or linearly. There is only slight divergence between the 2 integrations of the 0.48° nonlinear model, and the perturbation quantities are the same size in the tangent linear and non linear model integrations (figure 3.1A). The initial difference between the trajectories is larger than in the 0.1° models, as the initial perturbation has been applied to a grid cell with a larger volume.

Although trajectories diverge much more quickly in the undamped 0.1° model than in the damped 0.1° model, the linearisation is valid for a similar time interval, with the correlation between the TLM and the difference between two non linear forward model integrations falling below 0.8 after 200 days in the undamped 0.1° model and 250 days in the damped 0.1° model (figure 3.1B). There are therefore similar differences between the spatial structure of the perturbations calculated by the perturbed non linear model integrations and the TLM in both the undamped and damped 0.1° models, and the main difference between the two models is the much larger local growth of perturbations in the undamped 0.1° model. In contrast the correlation coefficient does not fall below 0.8 in the 0.48° model.

In the undamped 0.1° model there is a large difference in the eddy field between the two perturbed forward model integrations. This leads to high perturbation temperatures distributed throughout the channel, in both the TLM (figure 3.2A) and non linear model (figure 3.2D). In the damped 0.1° model the perturbations are concentrated to the north of the central jet in both the difference between the two non linear models (figure 3.2B) and in the TLM (figure 3.2E).

In the 0.48° model the perturbation to the non linear model shifts the phase of the wave travelling through the channel slightly, resulting in alternating crescents of positive and negative perturbation temperature in the centre of the channel (figure 3.2F). While the TLM shows a slightly smaller phase shift the pattern is

similar (figure 3.2C).

The validity of the linearisation obviously depends on the perturbation made, but the TLM results suggest that the adjoint method is likely to provide reliable information over a timescale of 200 days for the undamped 0.1° model and 250 days for the damped 0.1° model. For the 0.48° model there is no obvious time limit on its applicability.

3.2.2 Finite Difference Gradient Checks

The second common check on the usefulness of the adjoint gradients is to compare the sensitivities calculated by the adjoint model with those from finite difference gradients (section 1.2). Two different sets of finite difference gradient checks are carried out, different magnitude perturbations to the temperature at the point of maximum sensitivity at the initial time as calculated from the 690 day adjoint model integration, $T_{0_{\text{maxsens}}}$, and perturbing the surface wind stress at a number of grid cells.

As discussed in section 1.3.1 the macroscopic sensitivity of a time averaged quantity to the initial conditions is expected to approach 0 as the integration time increases. A maximum integration time of 2000 days is too short for this to happen for the sensitivity of $\overline{H}$ to $T_{0_{\text{maxsens}}}$ for any of the finite difference sensitivities calculated (figure 3.3D-F). In the damped and undamped 0.1° models the magnitude of the finite difference sensitivity increases with increasing integration time and with decreasing $T_{0_{\text{maxsens}}}$ (figure 3.3D-E). The increasing sensitivity with decreasing $T_{0_{\text{maxsens}}}$ is similar to that seen in the chaotic regime of the Lorenz (1963) equations (section 1.3.3). In the 0.48° model there is closer agreement between the finite difference sensitivities calculated for different $T_{0_{\text{maxsens}}}$ but, unlike the 0.1° models, an increase in $T_{0_{\text{maxsens}}}$ does not always lead to a decrease in $\frac{\Delta\overline{H}}{\Delta T_{0_{\text{maxsens}}}}$ (figure 3.3F).

The finite difference sensitivities agree well with the sensitivities calculated by the adjoint model at least out to 690 days in the damped 0.1° model and the 0.48° model for all $\Delta T_{0_{\text{maxsens}}}$ except $\Delta T_{0_{\text{maxsens}}} = 1$ (figure 3.3B-C). In the

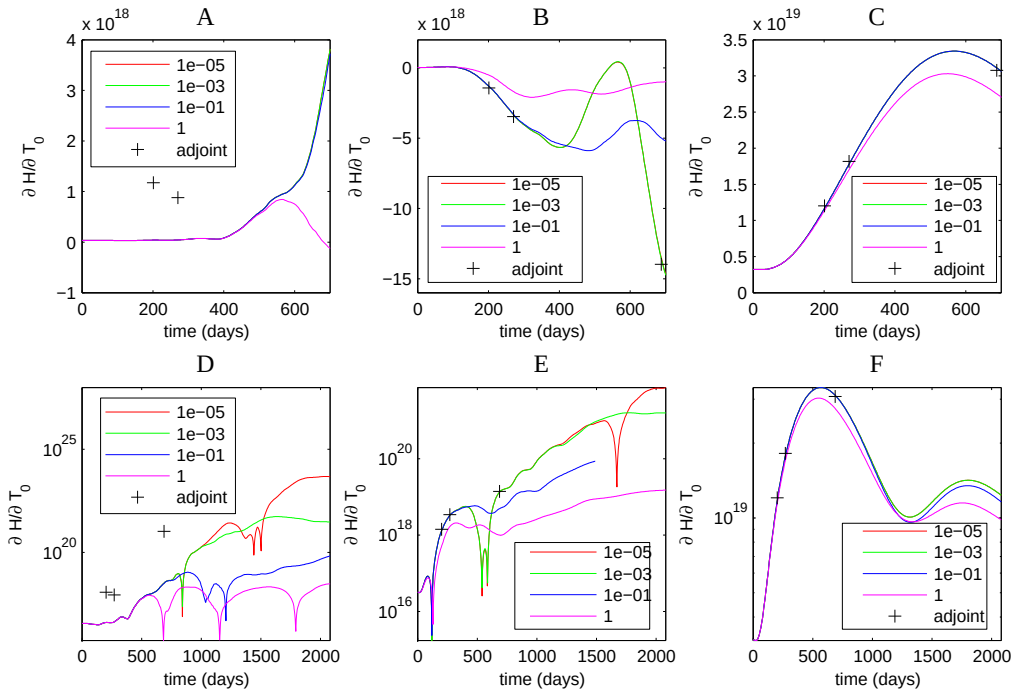

Figure 3.3: Centred difference finite difference gradients $\Delta\overline{H}/\Delta T_{0_{\mathrm{maxsens}}}$ for different $\Delta T_{0_{\mathrm{maxsens}}}$ (colour coding). Top row (A-C): Linear plot of $\Delta\overline{H}/\Delta T_{0_{\mathrm{maxsens}}}$ over the integration time used in the adjoint model experiments, A) undamped 0.1° model, B) damped 0.1° model, and C) 0.48° model. Bottom row (D-F): Logarithmic plot of $\Delta\overline{H}/\Delta T_{0_{\mathrm{maxsens}}}$ over over a period of 2000 days, A) undamped 0.1° model, B) damped 0.1° model, and C): 0.48° model. The results from the different length adjoint model integrations in section 3.3 are also shown. Note the different scales, and that the red and green curves are totally obscured by the blue curve in A and C, and the red curve is totally obscured by the green curve in B and F.

undamped 0.1° model the sensitivities calculated by the adjoint model are much larger than those calculated using finite difference gradients (figure 3.3A). In the undamped 0.1° model the sensitivity calculated by the adjoint model to $\Delta T_{0_{\text{maxsens}}}$ does not contain useful information, but it is not certain to what extent the finite difference sensitivities do. Results for the sensitivity of the $\overline{K}$ to $\Delta T_{0_{\text{maxsens}}}$ are similar and are not shown here.

A perturbation of $1 \times 10^{-5}\text{Nm}^{-2}$ is also made to τ_x at 5 different grid cells. After 690 days adjoint and finite difference gradients for $\overline{H}$ agree to within 2% in the 0.48° model, to within 12% in the damped 0.1° model, and to within 31% in the undamped 0.1° model. The finite difference sensitivities in the 0.48° model are an order of magnitude larger than in the damped 0.1° model due to the perturbation being made over a larger grid cell. However, the finite difference sensitivities calculated in the undamped 0.1° model are also an order of magnitude bigger than in the damped 0.1° model, for the same $\Delta\tau_X$ (table 3.1). Results from the experiments described in section 3.3 show that the sensitivity of $\overline{H}$ to the magnitude of the surface windstress τ_{X_0} is not an order of magnitude larger in the undamped 0.1° model than in the damped 0.1° model and 0.48° model (figure 3.4). For small ΔX_{0_i}, finite difference sensitivities do not contain any useful information due to the same local maxima and minima in the output function that affect the adjoint model results. Although it may be possible to choose a larger ΔX_{0_i} that is not effected by these local maxima and minima, larger scale curvature in the output function may mean that the resulting sensitivities are also of limited use.

3.3 Comparison with Perturbed Forward Model Experiments

Tanguay et al. (1995) demonstrated that the linearisation becomes inaccurate first at the smallest length scales. It is therefore reasonable to assume that information may be retained in the adjoint model solution for longer at larger spatial

Undamped 0.1° model

Adjoint (1×10^{20} JN^{-1}m^2)	Finite Difference (1×10^{20} JN^{-1}m^2)	1-$\frac{\text{finite difference}}{\text{adjoint}}$ (%)
2.0	1.5	25
2.3	1.6	27
2.4	1.7	29
2.6	1.8	31
2.7	1.9	31

Damped 0.1° model

Adjoint (1×10^{18} JN^{-1}m^2)	Finite Difference (1×10^{18} JN^{-1}m^2)	1-$\frac{\text{finite difference}}{\text{adjoint}}$(%)
1.8	1.9	-8.2
1.7	1.8	-9.2
1.5	1.7	-10
1.4	1.6	-11
1.3	1.5	-12

0.48° model

Adjoint (1×10^{19} JN^{-1}m^2)	Finite Difference (1×10^{19} JN^{-1}m^2)	1-$\frac{\text{finite difference}}{\text{adjoint}}$ (%)
1.1	1.1	-0.10
4.8	4.8	0.72
4.2	4.2	1.7
5.8	5.7	0.34
3.1	3.1	0.45

Table 3.1: Finite difference gradients $\frac{\Delta\overline{H}}{\Delta\tau_{X_0}}$ and adjoint gradients $\frac{\partial\overline{H}}{\partial\tau_{X_0}}$ for a 690 day adjoint model integration for different grid cells (rows).

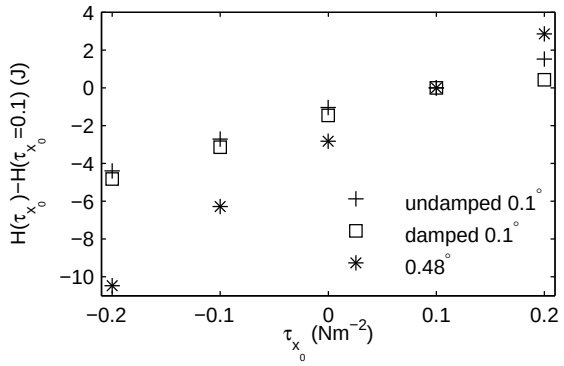

Figure 3.4: $\overline{H}$ as a function of τ_{X_0} for a 690 day integration. To allow all 3 models to be plotted on the same axis $\overline{H}$ from the unperturbed experiment τ_{X_0} has been subtracted.

scales. Instead of comparing the adjoint solution with point-wise gradient checks, the effect of larger-scale changes in the forcing or parameter value $\Delta\mathbf{X_0}$ can be computed from the sensitivities calculated by the adjoint model by $\sum_i \Delta\mathbf{X_{0_i}} \frac{\partial J}{\partial \mathbf{X_{0_i}}}$. This can then be directly compared with sensitivities calculated from perturbed forward model experiments. The adjoint model results can be said to agree well with the forward model results if a change in $\overline{H}$ or $\overline{K}$ calculated by the adjoint model is tangent to a curve fitted to the forward model results.

A series of 4 perturbed forward model experiments are used; 1) varying the vertical diffusivity, κ_v, over the entire model domain 2) increasing or decreasing the restoring SST* by a constant factor throughout the entire model domain 3) increasing or decreasing the meridional restoring SST* gradient while keeping the mean restoring SST* constant 4) varying the strength of the maximum zonal windstress, τ_{x_0}. All other forcings and parameters are held fixed at the same values as used in the adjoint run (table 2.1). The experiments are started from the same initial state as the forward model integration about which the adjoint model is linearised, and integrated for 690 days with the perturbed forcing or parameter value. The sensitivity to the perturbation calculated from the adjoint model results, cannot easily be compared with the sensitivity calculated from the forward model experiments at different times. A single integration of the adjoint model calculates the sensitivity of the output function evaluated at the final time to a perturbation in the forcing made at progressively earlier time steps, while in the perturbed forward model experiments the perturbation to the forcing is made at the initial time and the output function is evaluated at progressively later time steps. For a comparison at different times the adjoint model must either be started from different points along the forward model integration, or the forward model experiments must be repeated with the perturbation to the forcing made at different times. The former approach is used here, but due to the computational expense of running the adjoint model the comparison is only made at two times, 278 days and 690 days which are chosen based on the TLM results (section 3.2.1). After 278 days the response to a perturbation calculated by the TLM is an order of magnitude larger than that calculated by perturbed

nonlinear model integrations, in the undamped 0.1° model; after 690 days, the same applies for the response to a perturbation in the damped 0.1° model.

The dependence of $\overline{H}$ on all forcing is qualitatively similar between the models indicating similar physical processes are responsible, so only the undamped 0.1° model results are shown here. 278 days after the perturbation to the forcing is made $\overline{H}$ is a monotonic smoothly varying function of the mean SST* (crosses in figure 3.5A), τ_{x_0} (crosses in figure 3.5C), and κ_v (crosses in figure 3.5D). $\overline{H}$ is also a monotonic smoothly varying function of the SST* gradient in the undamped 0.1° model (crosses in figure 3.5B), and damped 0.1° model (not shown), but not in the 0.48° model where both an increase and a decrease in the SST* gradient cause $\overline{H}$ to decrease (not shown), and the response is clearly different to what would be expected in equilibrium. Nevertheless, the smooth response of $\overline{H}$ to changes in the forcing, suggests that $\overline{H}$ is clearly predictable after 278 days. 278 days is too short a timescale for the model to fully adjust to a change in the forcing, so the response after 690 days is larger, but has a similar form (stars in figure 3.5).

$\overline{H}$ depends almost linearly on τ_x and κ_v within the range of values used, and there is good agreement between a linear fit to the forward model results and the change in $\overline{H}$ calculated from the adjoint model results, in all models after 278 days (solid and dotted lines in figure 3.5 C and D). After 690 days the agreement between a change in $\overline{H}$ due to a change in τ_{x_0} or a change in κ_v calculated by the adjoint model is still good in the damped 0.1° model and the 0.48° model. In the undamped 0.1° model the agreement between the change in $\overline{H}$ due to a change in κ_v calculated using the forward and adjoint models has deteriorated after 690 days, but adjoint model results still have the correct sign and order of magnitude (dashed and dot-dashed lines in figure 3.5D). In contrast the sensitivity to τ_x calculated by the adjoint model is 2 orders of magnitude larger than that calculated from the forward model experiments and has the wrong sign, and there is no information retained in the adjoint solution even on this spatial scale (dashed and dot-dashed lines in figure 3.5C).

$\overline{H}$ does not depend linearly on either the mean SST*, or on the SST* gradient

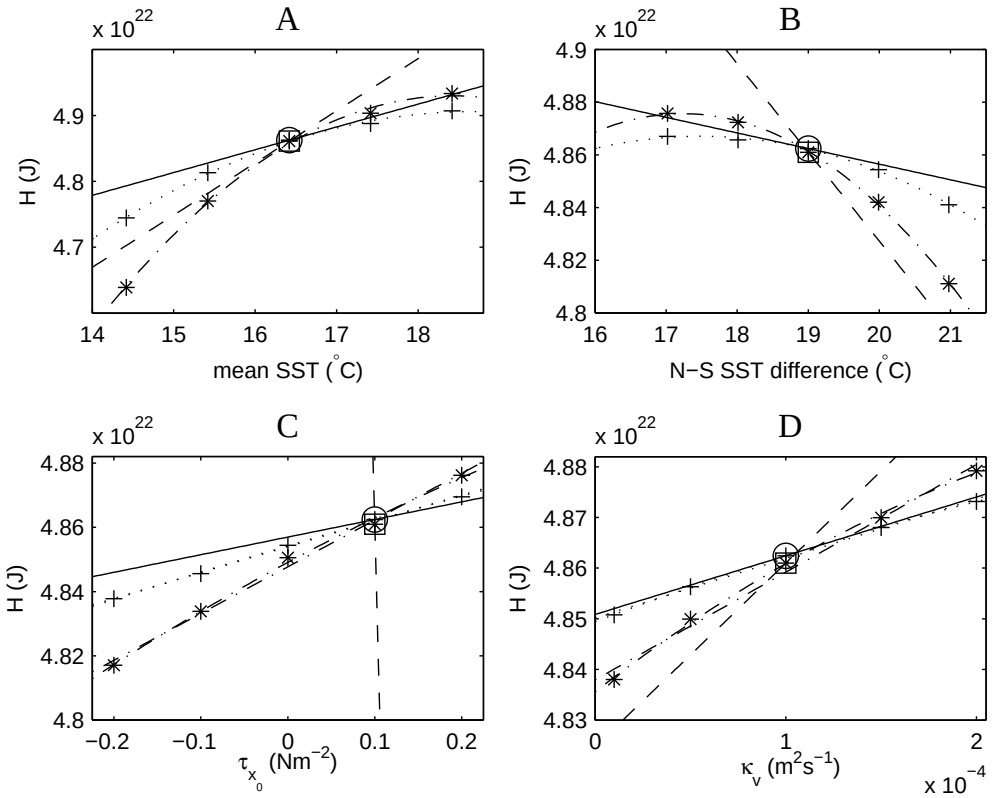

Figure 3.5: $\overline{H}$ in the perturbed forward model experiments, and as predicted from the adjoint model results for the undamped 0.1° model. A) $\overline{H}$ as a function of mean SST* B) $\overline{H}$ as a function of the North South SST* difference across the channel C) $\overline{H}$ as a function of the maximum zonal windstress τ_{x_0} D) $\overline{H}$ as a function of κ_v. Crosses, results from the perturbed forward model experiments after 278 days. Dotted line, quadratic fit to perturbed forward model experiments after 278 days. Solid line, response to perturbation predicted from the adjoint model results after 278 days. Circle, forward model state about which the adjoint model is linearised in 278 day integration. Stars, results from the perturbed forward model experiments after 690 days. Dash-dot line, quadratic fit to perturbed forward model experiments after 690 days. Dashed line, response to perturbation predicted from the adjoint model results after 690 days. Square, forward model state about which the adjoint model is linearised in 690 day integration.

as a cooling at the surface leads to convection, while a warming does not. In this case a quadratic is fitted to the data, as it is the lowest order polynomial which can approximate the data. The adjoint model is constructed so that conditional statements always take the same value as in the forward model (Giering, 1999), so it is unable to calculate the sensitivity associated with an increase in convection correctly. However, in the region of the unperturbed forward model experiments the adjoint model results still agree well with the perturbed forward model experiments in all models after 278 days (solid and dotted lines in figure 3.5A and B). After 690 days there is still good agreement between the change in $\overline{H}$ calculated from the adjoint model results, and from the perturbed forward model experiments, in the damped 0.1° and 0.48° model. Again the agreement in the undamped 0.1° model has deteriorated, but the response calculated by the adjoint model, still has the correct sign and order of magnitude (dashed and dot-dashed lines in figure 3.5A and B).

In the forward model experiments the dependence of $\overline{K}$ on the SST* gradient, τ_{x_0} and κ_v is qualitatively the same after 278 days in all models, so again only the undamped 0.1° model results are shown here (crosses in figure 3.6B-D). As for $\overline{H}$ the response after 690 days is larger, but has a similar form (stars in figure 3.6B-D). $\overline{K}$ does not depend linearly on either the SST* gradient, τ_{x_0} and κ_v so to enable comparison with the adjoint model results, a quadratic is fitted to the forward model results in all experiments. After 278 days there is good agreement between the adjoint and forward model results in all experiments in all 3 models (solid lines and dotted lines in figure 3.6B-D). After 690 days there is still good agreement with the adjoint and forward model results in the damped 0.1° model and 0.48° model, but no agreement at all in the undamped 0.1° model (dashed and dot-dashed lines in figure 3.6B-D).

The dependence of $\overline{K}$ on the mean SST* is more complicated, in the undamped 0.1° model a decrease in the mean SST* leads to an increase in $\overline{K}$ (crosses in figure 3.6A) while in the 0.48° model it leads to a decrease (not shown). After 278 days the adjoint and forward model results agree well in the undamped 0.1° model (solid and dotted lines figure 3.6A), and in the damped 0.1° model and 0.48° model

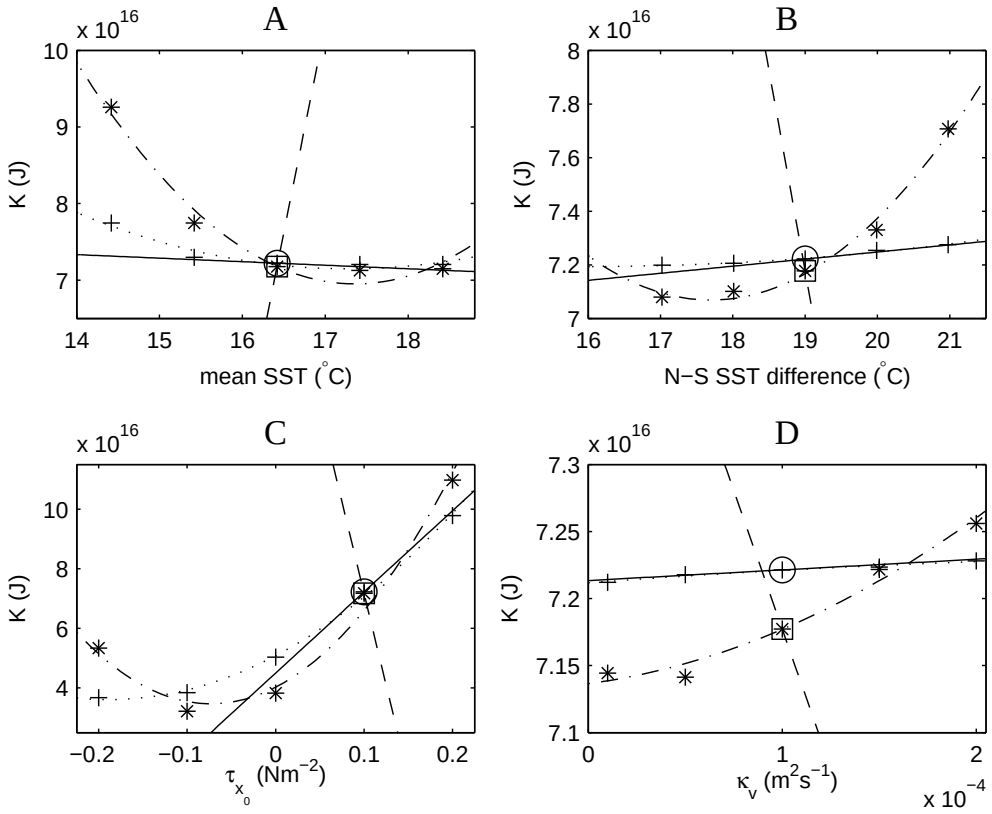

Figure 3.6: $\overline{K}$ in the perturbed forward model experiments, and as predicted from the adjoint model results for the undamped 0.1° model. A) $\overline{K}$ as a function of mean SST* B) $\overline{K}$ as a function of the North South SST* difference across the channel C) $\overline{K}$ as a function of the maximum zonal windstress τ_{x_0} D) $\overline{K}$ as a function of κ_v. Crosses, results from the perturbed forward model experiments after 278 days. Dotted line, quadratic fit to perturbed forward model experiments after 278 days. Solid line, response to perturbation predicted from the adjoint model results after 278 days. Circle, forward model state about which the adjoint model is linearised in 278 day integration. Stars, results from the perturbed forward model experiments after 690 days. Dash-dot line, quadratic fit to perturbed forward model experiments after 690 days. Dashed line, response to perturbation predicted from the adjoint model results after 690 days. Square, forward model state about which the adjoint model is linearised in 690 day integration.

(not shown). After 690 days there is still good agreement between the adjoint and forward model results in the damped 0.1° model, but response calculated by the adjoint model is too large and has the wrong sign in the undamped 0.1° model (dashed and dot-dashed lines in figure 3.6A), again there is no information remaining in the sensitivities calculated by the adjoint model even at the largest spatial scale. Although in both 0.1° models there is a greater response to the change in SST* in the forward model after 690 days than after 278 days, in the 0.48° model the response is reduced. In 0.48° model the adjoint method also gives the wrong sign for the sensitivity of $\overline{K}$ to the mean SST* after 690 days.

The greater success of the adjoint model at calculating the sensitivity of the heat content to SST* and κ_v can also be seen in the rate of growth of adjoint sensitivities. The sensitivity of $\overline{H}$ and $\overline{K}$ to SST*, κ_v and τ_x grows throughout the integration time in all the models (figure 3.7). If the integration time is shorter than the time needed for the model to reach equilibrium, a perturbation applied over a longer time period will have a greater effect on the time averaged climate quantities. However, although the adjoint sensitivities in the 0.48° and damped 0.1° model grow at similar rates, the sensitivities in the undamped 0.1° model, show the broadly exponential growth expected in a chaotic system (figure 3.7). The growth is not a pure exponential, there are times when the sensitivity grows less rapidly, or even decays. However the sensitivities of $\overline{H}$ to SST* (figure 3.7A), and $\overline{H}$ to κ_v (figure 3.7B), in the undamped 0.1° model remain the same order of magnitude as in the damped 0.1° model and 0.48° model for a longer time than the sensitivity of $\overline{H}$ to τ_{x_0} (figure 3.7C) and the sensitivity of $\overline{K}$ to SST*, κ_v, and τ_{x_0} (figure 3.7D-F). This suggests that sensitivities calculated by the adjoint model involving thermodynamic variables only may contain useful information on longer time scales than those that also involve dynamic variables as it takes longer for the exponential growth to dominate in this case. This may also be affected by the strength of restoring boundary conditions at the surface.

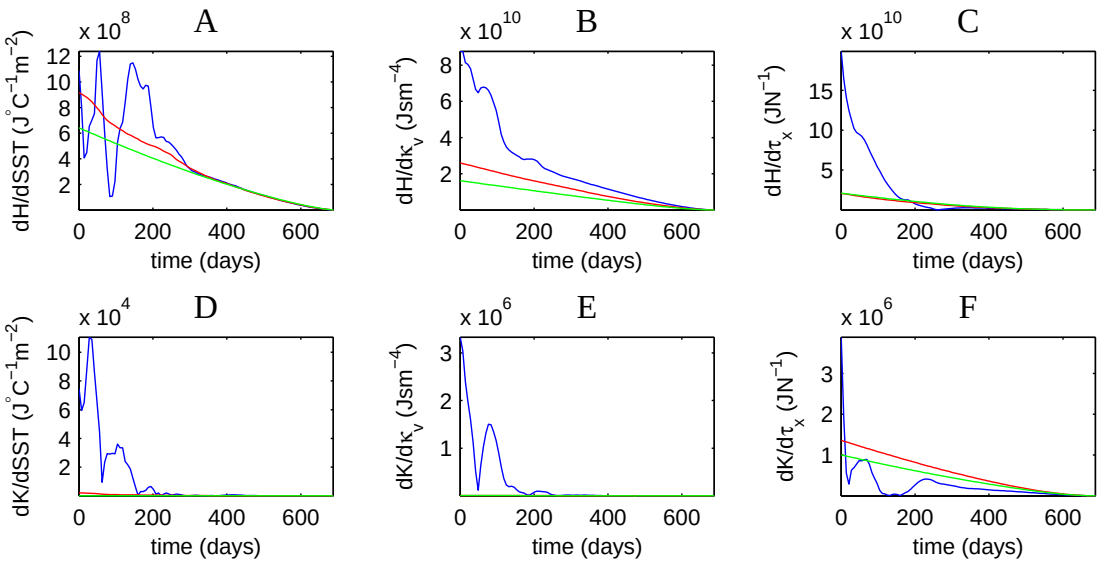

Figure 3.7: Growth of the sensitivity calculated by the adjoint model in the undamped 0.1° model, blue, the damped 0.1°model, red, and the 0.48° model, green. A) mean sensitivity per unit area of $\overline{H}$ to SST* B) mean sensitivity per unit volume of $\overline{H}$ to κ_v at 50m C) mean sensitivity per unit area of $\overline{H}$ to τ_x D) mean sensitivity per unit area of $\overline{K}$ to SST* E) mean sensitivity per unit volume of $\overline{K}$ to κ_v at 50m F) mean sensitivity per unit area of $\overline{K}$ to τ_x. Note the adjoint model starts from the end of the forward model integration and runs backwards in time, right to left here.

3.4 Discussion

The results presented here agree with those of Lea et al. (2002) and Köhl and Willebrand (2002), that chaos in eddy resolving ocean models does provide a time limit beyond which the adjoint model ceases to give useful information, which is related to the limit of predictability of the first kind. After 278 days when the comparison of the nonlinear forward model and the tangent linear model results suggest that the linear approximation assumed in the adjoint model is likely to be valid, the adjoint model is able to give the correct sensitivity of $\overline{H}$ and $\overline{K}$ to κ_v, τ_{x_0}, the mean SST* and the meridional SST* gradient in all models. After 690 days, when comparison between the nonlinear forward model and tangent linear model integrations suggest that the adjoint model is unlikely to provide useful information in the undamped 0.1° model the adjoint model is no longer able to give the correct sensitivity of $\overline{K}$ to any quantity or of $\overline{H}$ to τ_{x_0}, however, the adjoint model is still able to give the correct sensitivity of $\overline{H}$ to the mean SST*, SST* gradient and κ_v. This indicates that some information remains in the adjoint sensitivities for far longer than the non linear timescale of the system.

The amount of useful information in the adjoint method also depends on the quantity being studied. The sensitivity of the heat content to thermodynamic variables, SST*, and κ_v contains useful information over longer timescales than the sensitivity of the heat content to the dynamic variable τ_x, and the sensitivity of the kinetic energy to all variables. This indicates that the adjoint method is more stable with regards to the more slowly varying thermodynamic quantities, than to dynamic quantities, but caution is needed in this interpretation as the strong restoring to SST* damps eddies in the surface layer of the model and sensitivities calculated by the adjoint model may grow more rapidly in a model with weaker restoring.

Finite difference gradient checks fail to provide an indication of when an adjoint model can give us climatologically relevant sensitivity information. Lea et al. (2000) introduced the concepts of macroscopic and microscopic sensitivity of time averaged climate quantities, where the macroscopic sensitivity describes the re-

sponse of the climate to large scale changes in the forcings and parameter values, and the microscopic sensitivity describes the fluctuations in the response due to changes in the model trajectory caused by the change in forcing. The adjoint model calculates the microscopic sensitivity, and it is dominated by chaos. For short time scales finite difference methods also give the microscopic sensitivity. It is possible that on longer timescales finite difference gradients may give the macroscopic sensitivity, but this is not true of the adjoint method. In contrast experiments where a larger spatial scale perturbation has been made to the forcing can be used to show if the information obtained from the adjoint model is climatically relevant. However, they do not show if information is also retained at small spatial scales and this is the subject of the next chapter.

3.5 Summary

- Information remains in the large scale structure of sensitivities calculated by the adjoint model for far longer than the nonlinear timescale of the system. There is information remaining in the sensitivities calculated by the adjoint model on a timescale of at least 690 days, in the chaotic undamped 0.1° model.

- The adjoint method is able to calculate sensitivities only involving thermodynamic variables over a longer timescale than sensitivities that involve dynamic variables.

- Finite difference gradient checks do not provide a good indication of whether useful information remains in gradients calculated by an adjoint model in eddy resolving ocean models. Perturbed forward model experiments provide an alternative check on the usefulness of adjoint gradients.

Chapter 4

Spatial Structure of Sensitivities Calculated by the Adjoint Model

4.1 Introduction

The results of the previous section show that information is retained on the largest spatial scales, beyond the time at which the system can be expected to behave linearly. However, this does not automatically imply that the adjoint model still provides information efficiently. Running the adjoint model is around 5 times as expensive as running the forward model, so the sensitivity information provided by the adjoint model has to be in greater detail than could easily be obtained from perturbed forward model experiments. In this chapter the spatial structures in sensitivities calculated by the adjoint model are examined to see if they make physical sense, and if they provide information at a resolution that could not easily have been provided by perturbed forward model experiments.

In the previous chapter only two output functions, the heat content and the total kinetic energy were used. Two further output functions are introduced here, the available potential energy and the thermocline depth. As in the previous chapter sensitivities are analysed after 278 days when we expect the adjoint model to give good results, and after 690 days when we do not.

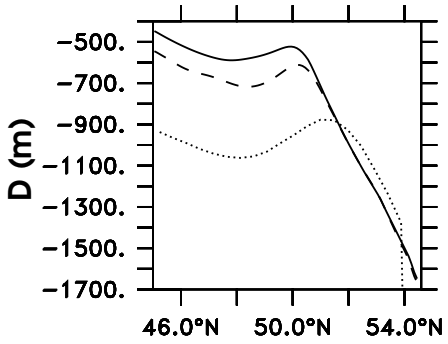

Figure 4.1: Zonally averaged thermocline depth. Solid line: undamped 0.1° model; dashed line: damped 0.1° model; dotted line: 0.48° model. All averages are calculated over the 690 day period used for the long adjoint model runs.

4.1.1 Thermocline Depth

Several definitions of the thermocline depth, D, exist. Here we use the e-folding scale for the potential density, σ, meaning that density of the thermocline is defined by

$$\sigma(D) = \sigma(-H) + \frac{\sigma(0) - \sigma(-H)}{e} \tag{4.1}$$

where H is the depth of the channel, $\sigma(-H)$ is the potential density at the bottom, and $\sigma(0)$ is the potential density at the surface. The depth of the thermocline is then found by searching for the first model layer k which is denser than $\sigma(D)$ and then linearly interpolating between the depths z_k and z_{k-1}.

$$D = (z_k - z_{k-1})\frac{\overline{\sigma(D)} - \overline{\sigma(z_{k-1})}}{\overline{\sigma(z_k)} - \overline{\sigma(z_{k-1})}} + z_{k-1} \tag{4.2}$$

The thermocline depth is similar in the damped and undamped 0.1° models. There is a local minimum at 49°N where there is convection associated with upwelling south of the front, and a sharp decrease in thermocline depth north of 50°N which is also associated with convection. In the 0.48° model the thermocline is deeper, but has a similar shape (figure 4.1).

The actual output function used here is the spatial and temporal mean thermocline depth $\overline{D}$, where the overbar denotes that the thermocline depth has been calculated from the spatial and temporal mean temperature and salinity. Defin-

ing the thermocline depth in the way used here involves conditional statements, which mean the sensitivity is poorly defined if $\sigma(z_k) = \sigma(D)$. By using a definition for the mean thermocline depth based on time and spatially averaged quantities, $\overline{D} = D(\sigma(\overline{T_z}))$, rather than the time averaged thermocline depth calculated from instantaneous quantities, this becomes less likely and $\sigma(z_k)$ was not equal to $\sigma(\overline{D})$ for any of the adjoint integrations used here.

4.1.2 Available Potential Energy

The kinetic energy of the channel can be changed, either by direct forcing through increased windstress, or indirectly as a response to a change in the density structure, and thus the potential energy, of the channel. The available potential energy (APE) provides a measure of how much of the potential energy in the channel is available for conversion to kinetic energy, and looking at the sensitivity of the APE should thus allow us to distinguish between sensitivity to direct and indirect forcing.

The total potential energy contained in the worlds oceans,

$$\text{TPE} = \iiint \rho g z \, \mathrm{d}x \, \mathrm{d}y \, \mathrm{d}z \tag{4.3}$$

is seven orders of magnitude greater than the total kinetic energy (Oort et al., 1989). However, not all of the potential energy is available for transformation into kinetic energy. Lorenz (1955) defined the concept of available potential energy, APE, as the difference between the total potential energy and the potential energy of a reference state that has the minimum total potential energy which could result from an adiabatic redistribution of mass. The reference state is stably stratified and has horizontal isopycnals, and the sum of available potential energy and kinetic energy is conserved under adiabatic flow. For the global ocean the available potential energy is just 0.001% of the total potential energy (Oort et al., 1989).

Several methods of calculating the oceanic APE, or more specifically the energy of the reference state, have been proposed. Here we use the method of

Oort et al. (1989), which is a modified version of Lorenz (1955)'s method for the atmosphere. The APE is defined as,

$$\mathrm{A} = -\frac{1}{2} g \iiint \frac{(\rho - \tilde{\rho})^2}{\delta\tilde{\sigma}/\mathrm{d}z} \,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z \tag{4.4}$$

where $\tilde{\rho}$ is the global mean density at the constant height z and $\delta\tilde{\sigma}/\mathrm{d}z$ is the gradient of potential density. This definition assumes that horizontal gradients in density are much smaller than vertical gradients, which is normally the case in the ocean. It is easily seen from equation 4.4 that A decreases if the isopycnals become more horizontal, or the stratification becomes more stable.

4.2 Sensitivity after 278 Days

4.2.1 Sensitivity to SST*

In all models regions of high sensitivity to SST* are associated with convection in the forward model, as perturbations to SST* in these regions are mixed to a greater depth and effect a larger volume of water. In all models convection is most frequent and deepest along the northern boundary of the channel (figure 4.2D-F). There are also shallower bands of convection associated with upwelling south of the front between 49°N and 50°N in the 0.1° models and between 47° and 50°N in the 0.48° model (figure 4.2D-F). In the undamped 0.1° model there is also a band of convection at 52°N (figure 4.2E) this is associated with upwelling south of the secondary zonal jet (section 2.3.4) .

Sensitivity of $\overline{H}$ to SST*

The sensitivity of $\overline{H}$ to SST*, in the 0.48° model is highest along the northern boundary, and south of the front at 48°N where there is convection in the forward model (figure 4.3C). After 278 days both the damped and undamped 0.1° models have regions of high sensitivity of $\overline{H}$ to SST*, at the northern boundary and at 50°N, also associated with convection (figure 4.3A and B). Although convection is stronger at 52°N than at 50°N in the undamped 0.1° model there is no clearly

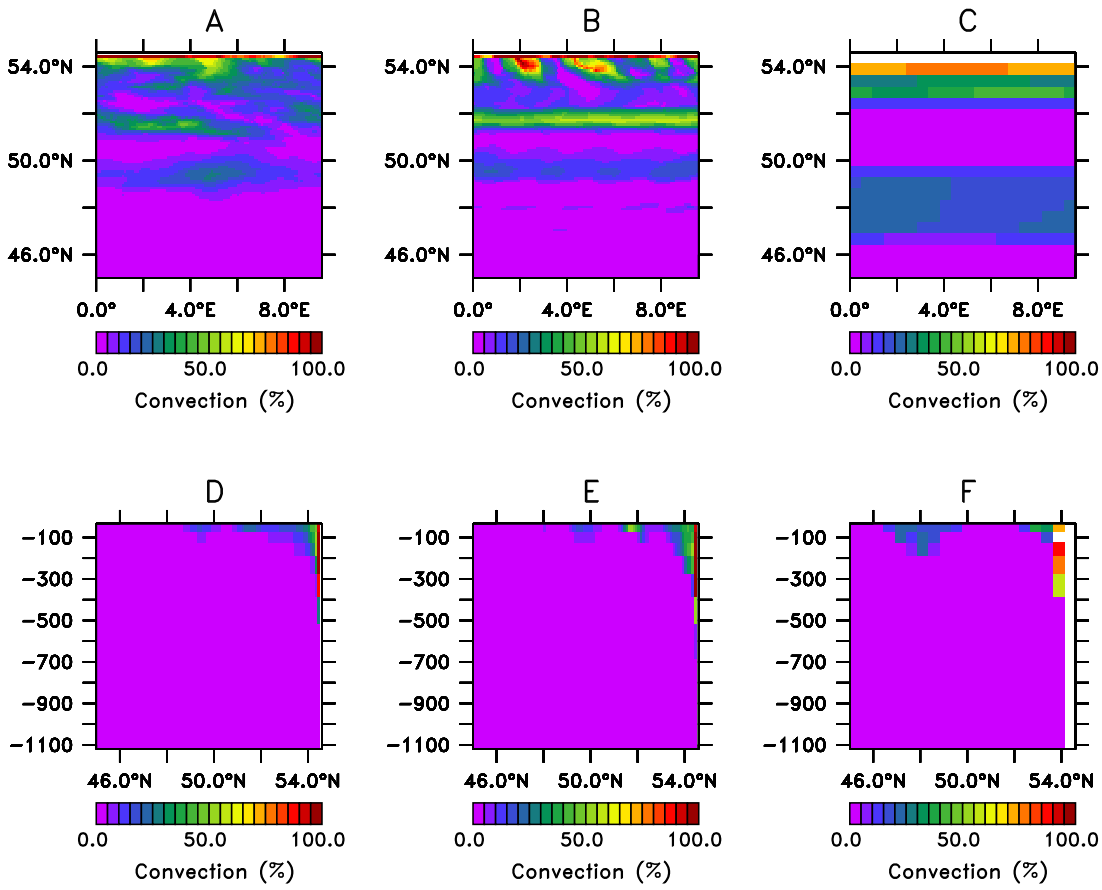

Figure 4.2: Number of timesteps convection occurs averaged over 278 days. Top row (A-C) convection between the top two model layers. Bottom row (D-F) zonally averaged convection. Left (A,D): undamped 0.1° model, centre (B,E): damped 0.1°, right (C,F): 0.48° model.

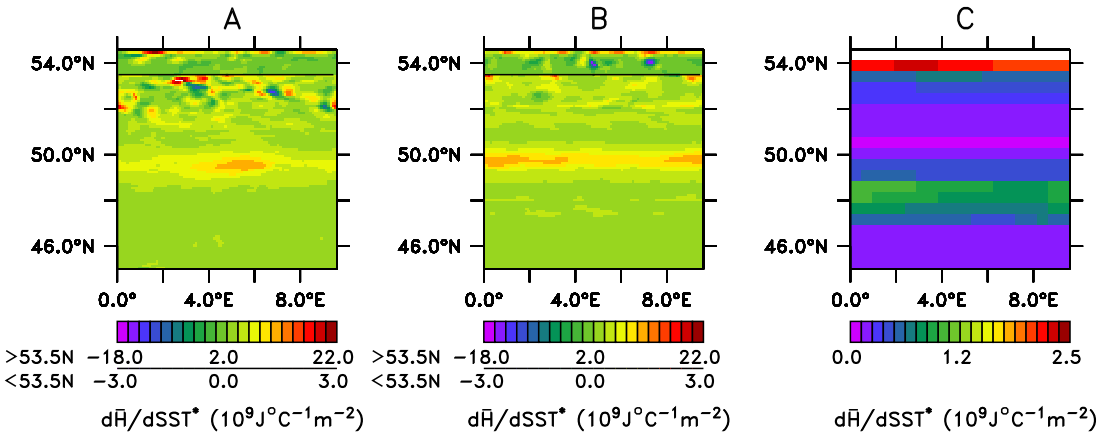

Figure 4.3: Sensitivity of $\overline{H}$ to the restoring SST* as calculated by the adjoint model at T=0 in a 278 day integration. A: undamped 0.1° model B: damped 0.1° C: 0.48° model. Due to the extremely high sensitivity to SST* along the northern boundary in both 0.1° models, a different scale is used north of 53.5°N (black line in panels A and B, and 2 sets of numbers along the colour key). Also note the difference in scale between the panels.

discernible band of high sensitivity at 52°N (figure 4.3B). At 52°N the convection is confined to the second model layer (figure 4.2E), so that a SST* perturbation in this region effects the shallow surface layers only .

Sensitivity of $\overline{D}$ to SST*

In the damped 0.1° model there is a band of high negative sensitivity at the northern boundary and a band of positive sensitivity at 48°N (figure 4.4B). In the undamped model the thermocline depth is 738m and is similar to the maximum depth of convection at the northern boundary (figure 4.2E). An increase in SST* at the northern boundary is thus able to reach the depth of the thermocline in the damped 0.1° model, causing it to deepen ($\overline{D}$ becomes more negative). The band of convection at 50°N is only 150m deep (figure 4.2E), and the integration time is insufficient for an increase in SST* in this region to diffuse down to the thermocline, and the resulting warming is all above the level of the thermocline. The sensitivity of $\overline{D}$ to SST* is positive throughout the channel in the 0.48° model, including the northern boundary (figure 4.4C). In the 0.48° model the thermocline is deeper, 1013m, and the maximum depth of convection is shallower

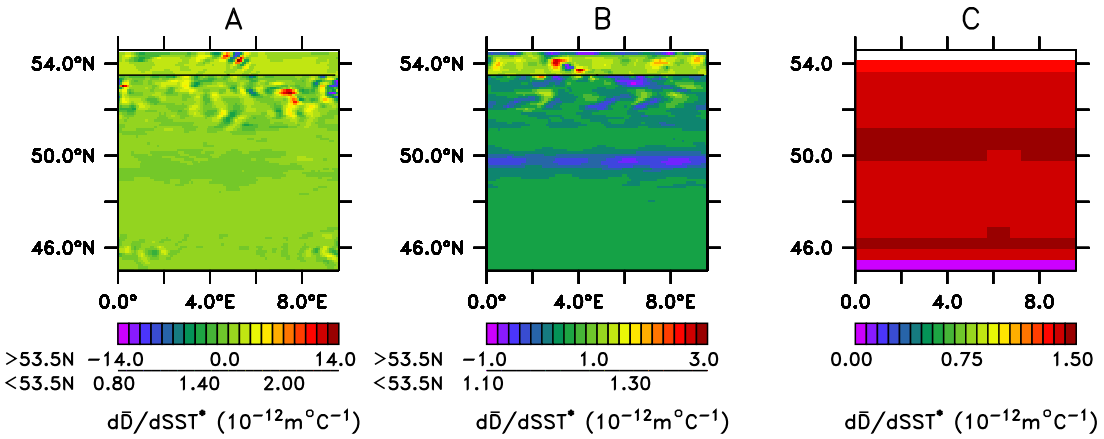

Figure 4.4: Sensitivity of $\overline{D}$ to SST* as calculated by the adjoint model after 278 days of a 690 day integration. A: the undamped 0.1° model; B: the damped 0.1° model; C: the 0.48° model. Due to the extremely high sensitivity to SST* along the northern boundary in both 0.1° models, a different scale is used north of 53.5°N (black line in panels A and B and 2 sets of numbers along the colour key). Also note the difference in scale between the panels.

(figure 4.2F), so that throughout the entire channel the increase in SST* all the heat uptake due to an increase in SST* is above the thermocline. In the undamped 0.1° model there is a band of positive sensitivity at 50°N but it is less distinct than in the damped 0.1° model. No band of high sensitivity can be seen at the northern boundary due to the small scale patches of extremely high sensitivity similar to those seen in the heat content.

Sensitivity of $\overline{K}$ to SST*

In the 0.48° model the sensitivity of $\overline{K}$ associated with convection at the northern boundary is negative, where an increase in SST* would tend to decrease the meridional temperature gradient and the available potential energy, and positive at 48°N where an increase in SST* would tend to increase the meridional temperature gradient and the available potential energy. There is no obvious pattern in the sensitivity of $\overline{K}$ to SST* in the damped or undamped 0.1° model at 278 days, due to patches of extremely high sensitivity at the northern boundary, similar to those seen in the sensitivity of $\overline{H}$ to the SST* (figure 4.5A and B).

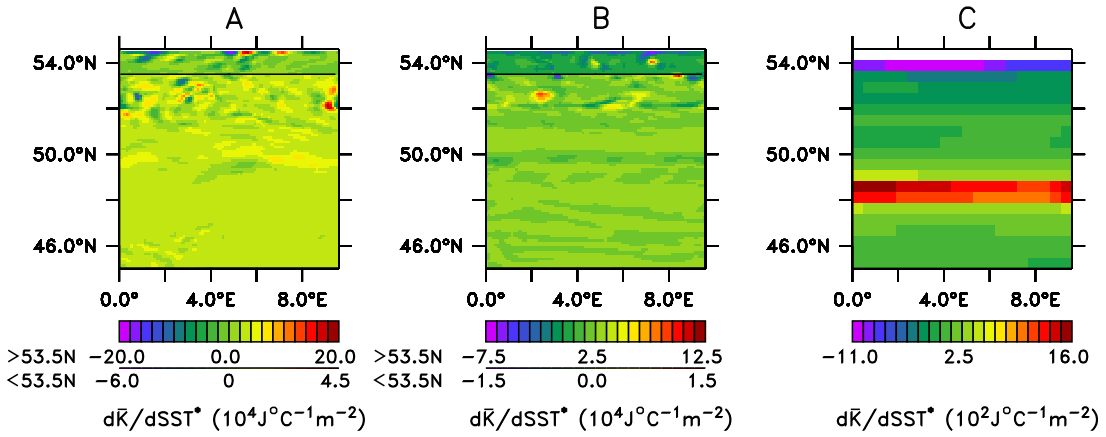

Figure 4.5: Sensitivity of $\overline{K}$ to the restoring SST* as calculated by the adjoint model at T=0 in a 278 day integration. A: undamped 0.1° model B: damped 0.1° C: 0.48° model. Due to the extremely high sensitivity to SST* along the northern boundary in both 0.1° models, a different scale is used north of 53.5°N (black line in panels A and B, and 2 sets of numbers along the colour key). Also note the difference in scale between the panels.

Sensitivity of $\overline{A}$ to SST*

The sensitivity of $\overline{A}$ to SST* in the 0.48° model is negative everywhere (figure 4.6C). Warming the surface of the channel increases the static stability and reduces the available potential energy. This dominates the sensitivity of $\overline{A}$ to SST* in the 0.48° model, so that the effect of changing the meridional temperature gradient, which is clearly visible, in the sensitivity of $\overline{K}$ to SST*, is not seen. Again the sensitivity is most negative where there is convection, as the temperature perturbations are mixed to a greater depth effecting a larger volume of water.

In the damped and undamped 0.1° models there are bands of high positive sensitivity of $\overline{A}$ to SST* at 50°N, and bands of negative sensitivity along the northern boundary (figure 4.6B). Again these correspond to bands of convection in the forward model. Although all three models have negative sensitivity along the northern boundary, the bands of sensitivity associated with convection due to upwelling at the south of the front are of opposite sign in the 0.1° models and 0.48° model. Heating here would tend to increase the meridional temperature

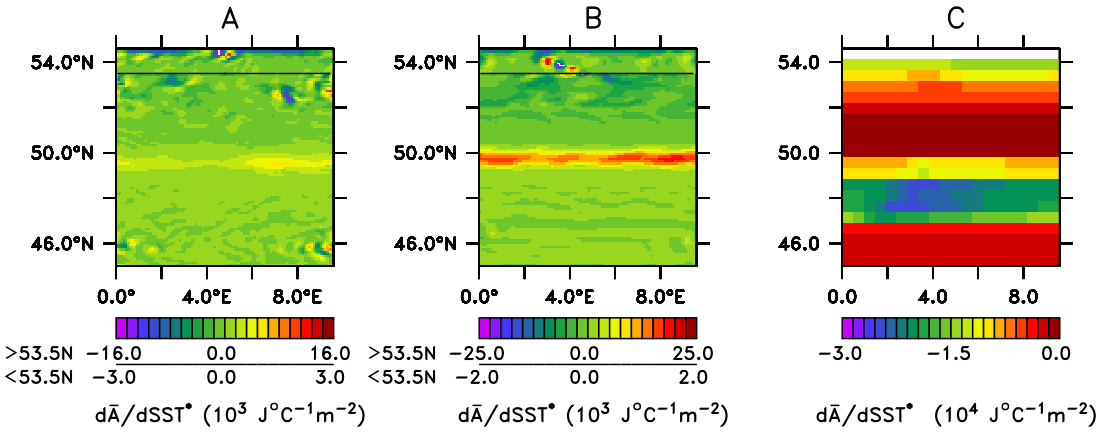

Figure 4.6: Sensitivity of $\overline{A}$ to SST* as calculated by the adjoint model after 278 days of a 690 day integration. A: the undamped 0.1° model; B: the damped 0.1° model; C: the 0.48° model. Due to the extremely high sensitivity to SST* along the northern boundary in both 0.1° models, a different scale is used north of 53.5°N (black line in panels A and B and 2 sets of numbers along the colour key). Also note the difference in scale between the panels.

gradient and hence the available potential energy. Different processes dominate the sensitivity of the available potential energy to SST* in the different resolution models.

4.2.2 Sensitivity to κ_v

The forward model uses a constant value of κ_v, but the adjoint model calculates the sensitivity to κ_v throughout the channel. The sensitivity to κ_v in the surface layer is identically zero, as the model does not use the diffusivity in the surface layer to calculate heat fluxes. The downwards diffusive heat flux in the model depends on $\kappa_v \partial\overline{T}/\partial z$, so that sensitivity of all quantities to κ_v tends to be highest where $\partial\overline{T}/\partial z$ is highest. $\partial\overline{T}/\partial z$ is highest near the surface (figure 4.7D-F) and in the south of the channel (figure 4.7A-C).

Sensitivity of $\overline{H}$ to κ_v

After 278 days sensitivity of $\overline{H}$ to the vertical diffusivity, κ_v, is highest at 50m in the 2nd model layer in the 0.48° model and both the damped and undamped 0.1° models (figure 4.8D-F). The short integration time means that the sensitivity is

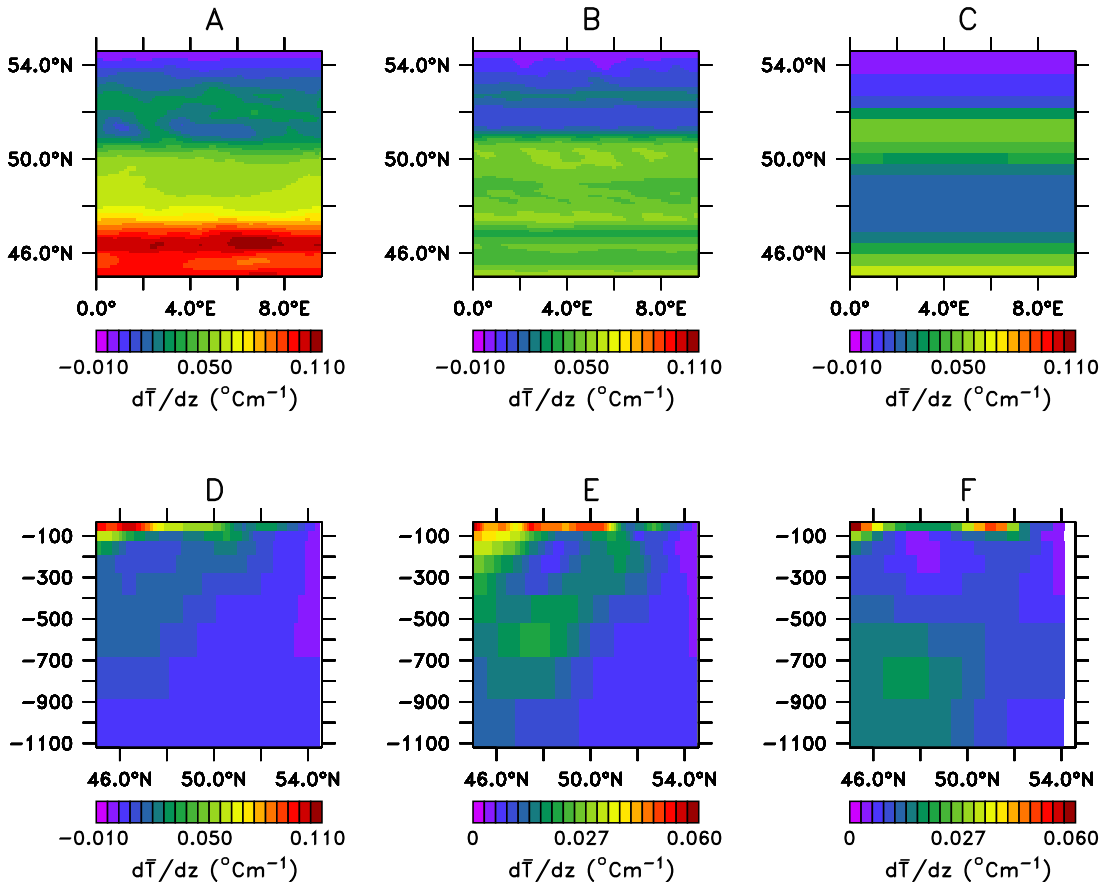

Figure 4.7: $\partial\overline{T}/\partial z$ averaged over the 690 day integration used for the adjoint model integrations. Top row (A-C),$\partial\overline{T}/\partial z$ at 50m. Bottom row (D-E), zonally averaged $\partial\overline{T}/\partial z$. Left (A,D): the undamped 0.1° model; centre (B,E): the damped 0.1° model; right (C,F): the 0.48° model. Note difference in scales between the panels.

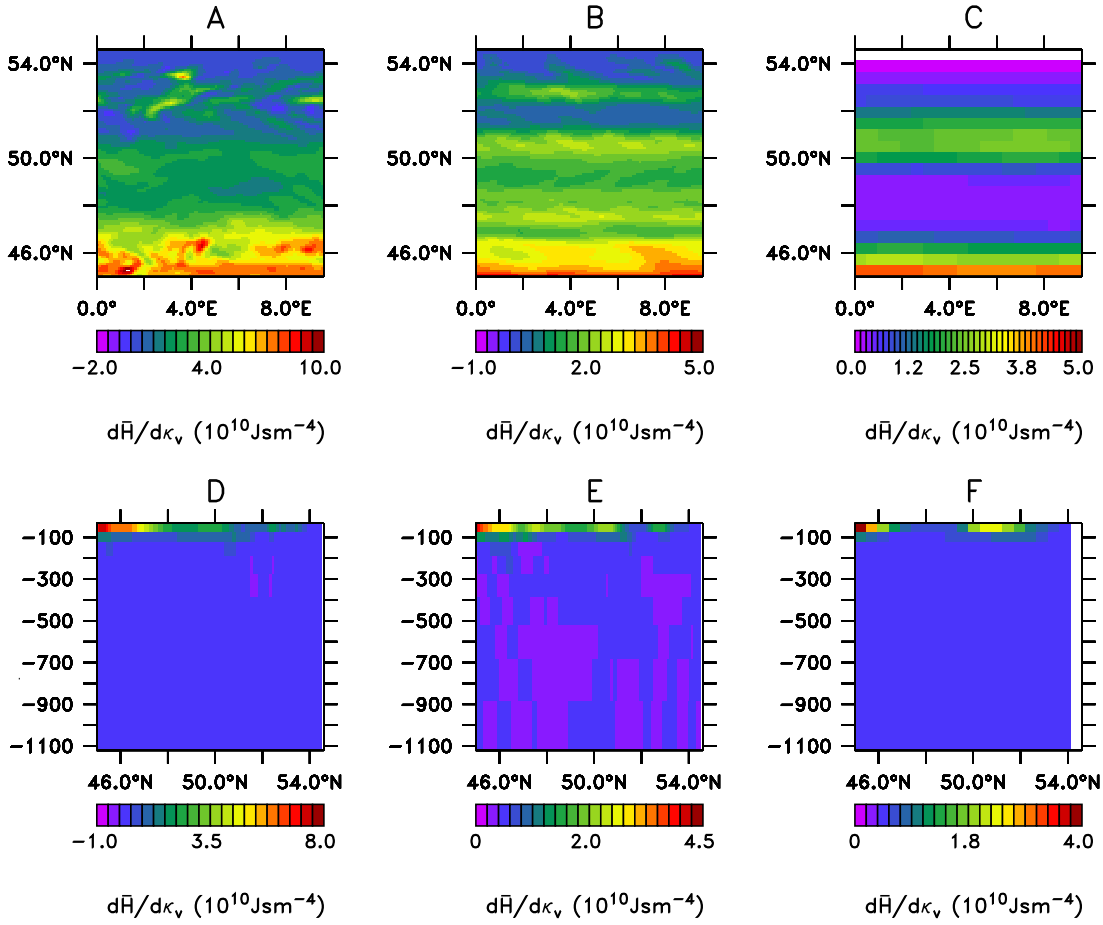

Figure 4.8: Sensitivity of $\overline{H}$ to κ_v as calculated by the adjoint model at T=0 in a 278 day integration. Top row (A-C), sensitivity of $\overline{H}$ to κ_v at 50m. Bottom row (D-E), zonally averaged sensitivity of $\overline{H}$ to κ_v Left (A,D): the undamped 0.1° model; centre (B,E): the damped 0.1° model; right (C,F): the 0.48° model. Note difference in scales between the panels.

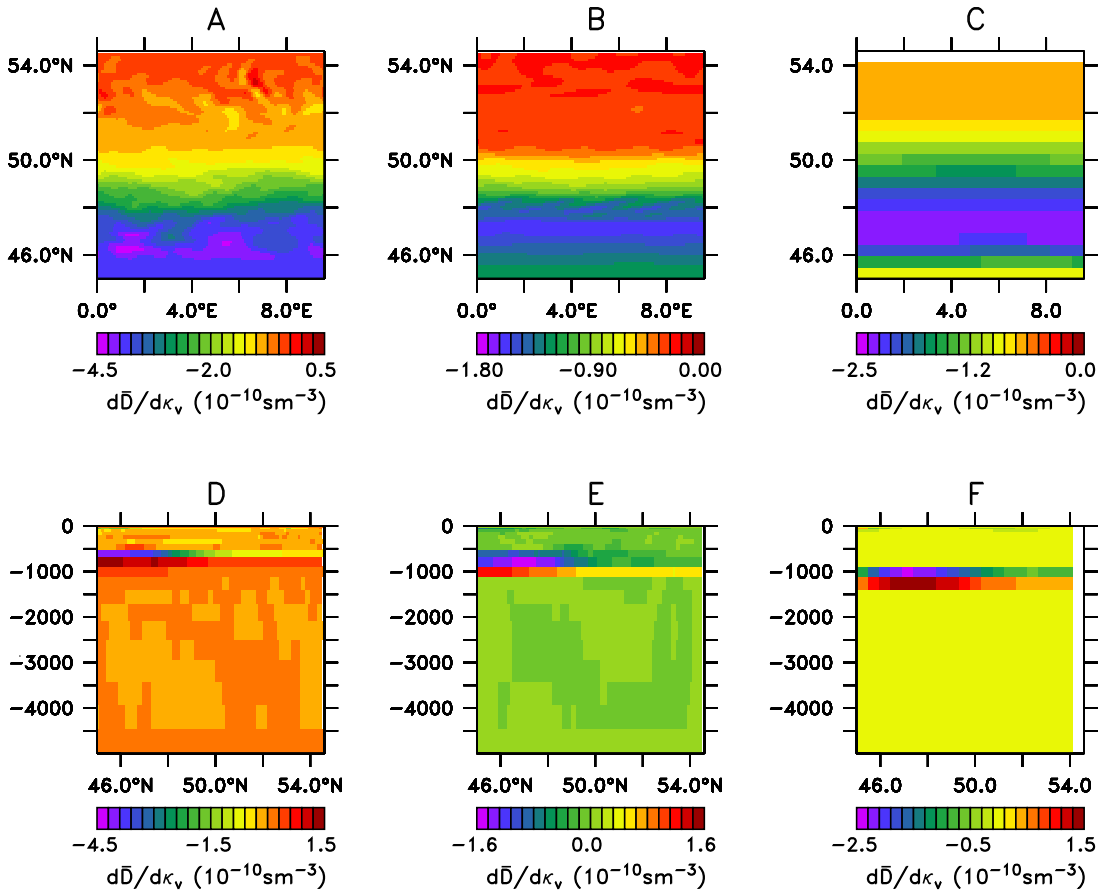

Figure 4.9: Sensitivity of $\overline{D}$ to κ_v as calculated by the adjoint model after 278 days of a 690 day integration. A: Sensitivity in the undamped 0.1° model at 593m. B: Sensitivity in the damped 0.1° model at 774m C: Sensitivity in the 0.48° model at 992m. Bottom row (D-F): zonally averaged sensitivity, in D the undamped 0.1° model, E the damped 0.1° model and F the 0.48° model. Note the difference in scale between the panels.

near zero elsewhere, as the only heat source is at the surface and the choice of κ_v ($1 \times 10^{-4}\mathrm{m}^2\mathrm{s}^{-1}$) means that heat can only diffuse by about 50m in 278 days. At 50m there are bands of high sensitivity along the southern boundary in all models and at 51°N in the 0.48° model, and at 53°N in the damped 0.1° model (figure 4.8A-C). These bands of high sensitivity all correspond to areas of high $\partial\overline{T}/\partial z$ (figure 4.7A-C).

Sensitivity of $\overline{D}$ to κ_v

The sensitivity of $\overline{D}$ to κ_v is strongly negative above the thermocline and strongly positive below it in all 3 models (figure 4.9D-F). Increasing the diffusivity in a model layer increases the heat flux into that layer so this pattern of sensitivity reduces $\partial\sigma/\partial z$. Directly above the thermocline the sensitivity is most strongly negative in the south of the channel where $\partial\overline{T}/\partial z$ is highest, and lowest in the north of the channel where $\partial\overline{T}/\partial z$ is lowest. This signal can be seen clearly in all three models. As for the sensitivity of $\overline{H}$ the sensitivity of $\overline{D}$ to κ_v is also high at the surface (figures 4.9D-F) where $\partial T/\partial z$ is highest. However, the sensitivity is much smaller than at the level of the thermocline. Elsewhere the sensitivity is near zero as the integration time is insufficient for the signal to diffuse over a greater distance.

Sensitivity of $\overline{K}$ to κ_v

After 278 days there is high sensitivity of $\overline{K}$ to κ_v, corresponding to high $\partial\overline{T}/\partial z$, between 50m and 150m in all three models (figure 4.10D-F). At 50m the sensitivity is positive south of the front and negative north of the front, in both the 0.48° model and the damped 0.1° model (figure 4.10B and C). An increase in κ_v south of the front leads to a greater meridional temperature gradient across the front, and an increase in available potential energy, while an increase in κ_v north of the front has the opposite effect. Although this can also be seen in the zonally averaged sensitivity in the undamped 0.1° model (figure 4.10D), the signal is hidden by small scale patches of extremely high sensitivity around 46°N in the plot of the sensitivity of $\overline{K}$ to κ_v at 50m (figure 4.10A). In the damped 0.1° model there is an area of high negative sensitivity of $\overline{K}$ to κ_v between 400m and 500m and 47°N to 51°N and an area of high positive sensitivity between 500m and 700m (figure 4.10E). This is at the base of the front (figure 2.5B), increasing the diffusivity below the base of the front tends to make the front deeper, increasing the extent of the zonal jet, while decreasing the diffusivity above the base of the front would tend to make it shallower. This is not seen in the 0.48° model (figure 4.10F).

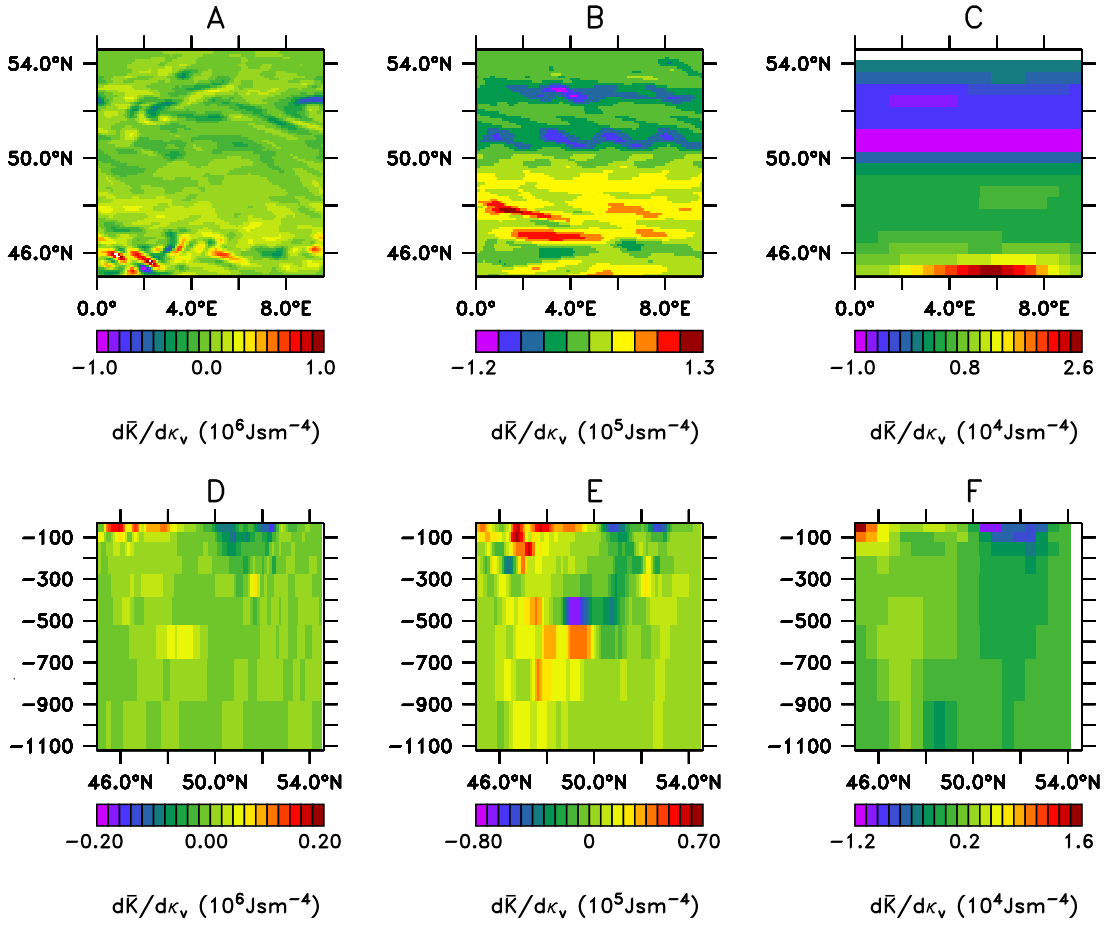

Figure 4.10: Sensitivity of $\overline{K}$ to κ_v as calculated by the adjoint model at T=0 in a 278 day integration. Top row (A-C), sensitivity of $\overline{K}$ to κ_v at 50m. Bottom row (D-E), zonally averaged sensitivity of $\overline{K}$ to κ_v Left (A,D): the undamped 0.1° model; centre (B,E): the damped 0.1° model; right (C,F): the 0.48° model. Note difference in scales between the panels.

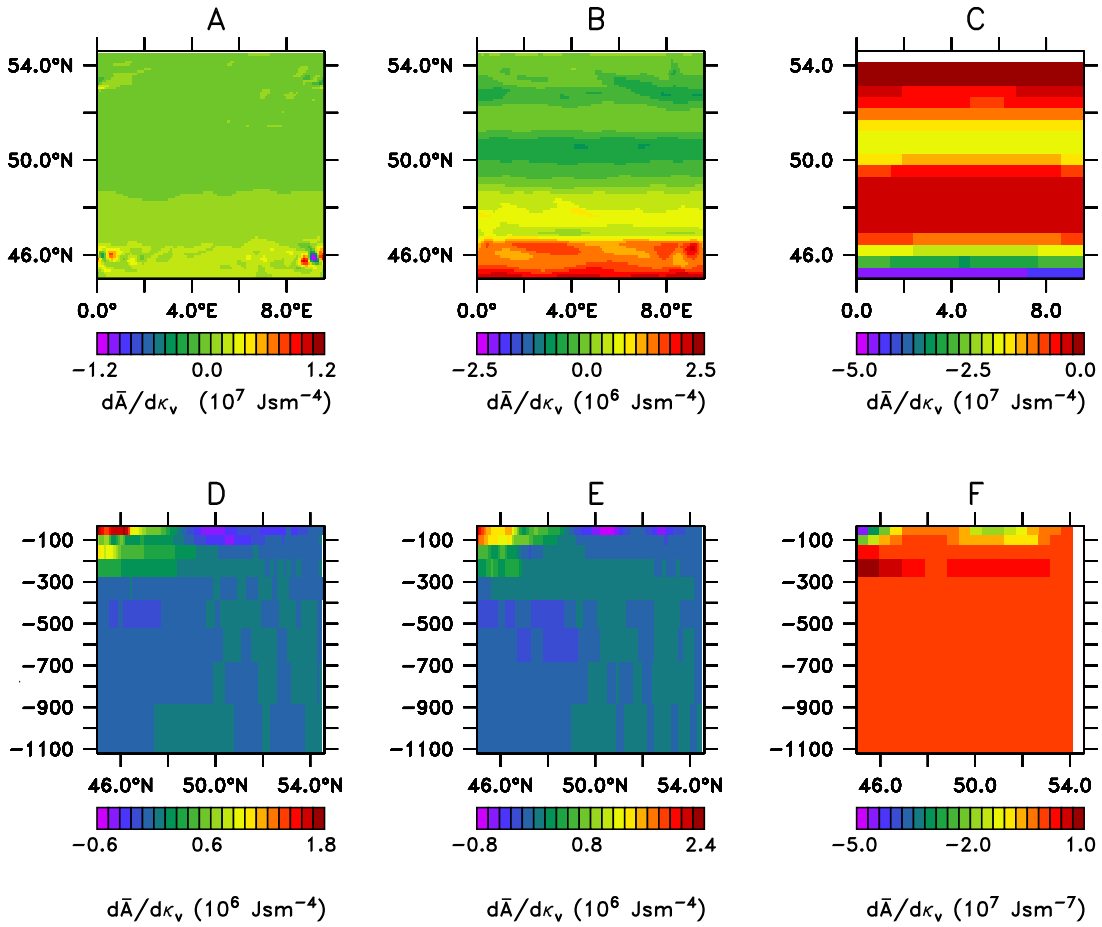

Figure 4.11: Sensitivity of $\overline{A}$ to κ_v as calculated by the adjoint model at T=0 in a 690 day integration. Top row (A-C): sensitivity at 50m. Bottom row (D-F): zonally averaged sensitivity. Left column (A,D): the undamped 0.1° model; middle column (B,E): the damped 0.1° model; right column (C,F): the 0.48° model. Note the difference in scale between the panels.

Sensitivity of $\overline{A}$ to κ_v

As for the sensitivity of $\overline{K}$ to κ_v the sensitivity of $\overline{A}$ to κ_v is positive in the south of the channel but negative in the north of the channel in the damped 0.1° model (figure 4.11B), as this would tend to increase the meridional temperature gradient. This signal can be seen in the zonally averaged sensitivity in the undamped 0.1° model (figure 4.11D) but is difficult to see in the sensitivity at 50m due to small patches of high sensitivity along the southern boundary. In the 0.48° model the sensitivity of $\overline{A}$ to κ_v is negative everywhere at 50m (figure 4.11C). The static stability dominates the sensitivity of $\overline{A}$ to κ_v in the 0.48° model in the same

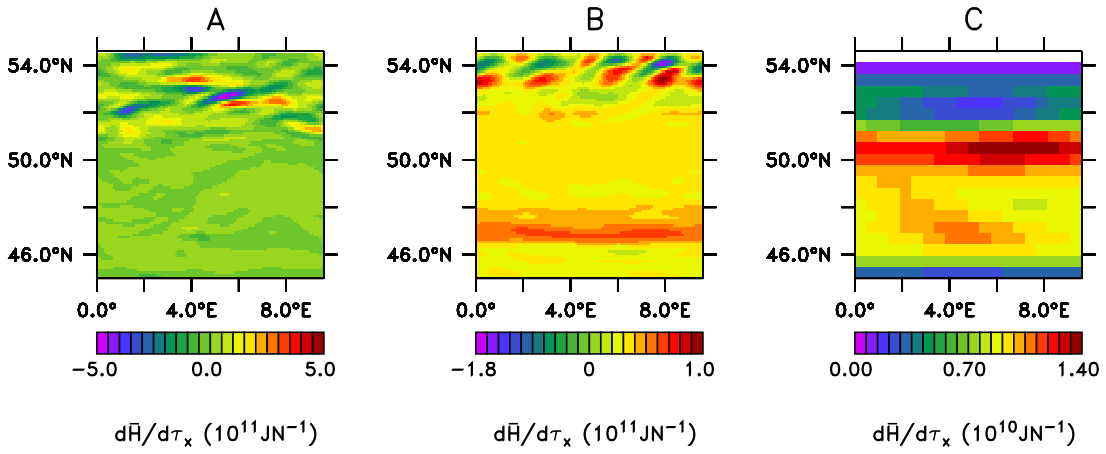

Figure 4.12: Sensitivity of $\overline{H}$ to τ_x as calculated by the adjoint model at T=0 in a 278 day integration. A: undamped 0.1° model B: damped 0.1° C: 0.48° model. Note difference in scales, between the panels.

way as it dominates the sensitivity of $\overline{A}$ to SST*. Between 200m and 400m the sensitivity to κ_v is positive in the 0.48° model (figure 4.11F). There is a minimum in $\partial\sigma/\partial z$ at 200m, and decreasing κ_v above this level, and increasing κ_v below this level further decreases $\partial\sigma/\partial z$ in this region. The sensitivity of $\overline{A}$ to κ_v is negative between 400m and 700m in the damped 0.1° model, as this again tends to increase the static stability (figure 4.11E).

4.2.3 Sensitivity to τ_x

Unlike the sensitivity to SST* and κ_v the sensitivity to τ_x is not governed by local processes, but by changes in the circulation throughout the entire channel that result from a change in the windstress, and there are no simple explanations for the patterns in the sensitivity.

Sensitivity of $\overline{H}$ to τ_x

Increased eastwards windstress leads to an increase in the southwards Ekman transport at the surface, and thus deeper down-welling of warm water south of the jet. Consequently the sensitivities of $\overline{H}$ to τ_x are positive everywhere in the 0.48° model. There is a band of high sensitivity at 51°N in the 0.48° model

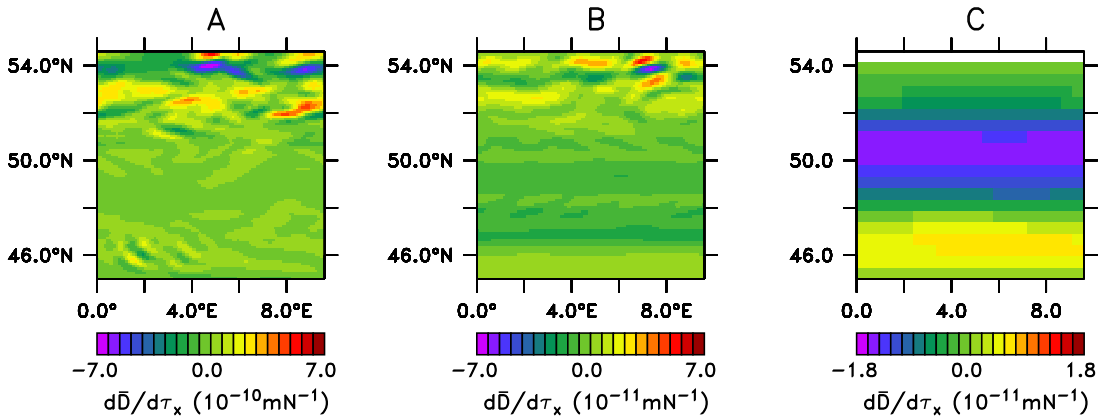

Figure 4.13: Sensitivity of $\overline{D}$ to τ_x as calculated by the adjoint model after 278 days of a 690 day integration. A: the undamped 0.1° model; B: the damped 0.1° model; C: the 0.48° model. Note the difference in scale between the panels.

(figure 4.12C) and at 47°N in the damped 0.1° model (figure 4.12B). Again the sensitivity in the undamped 0.1° model is dominated by patches of very high positive and negative sensitivity, which are unlikely to be the sensitivity of $\overline{H}$ to a non infinitessimal perturbation to τ_x in these regions.

Sensitivity of $\overline{D}$ to τ_x

In the 0.48° model the sensitivity of $\overline{D}$ to τ_x is near zero north of 52°N, negative between 48°N and 52°N, and positive south of 52°N (figure 4.13C). North of 52°N $\overline{D}$ is controlled by the depth of convection, which is constant in the adjoint model integrations. The reason for the exact pattern of sensitivity south of 52°N is unclear, but could be ascertained by making this shape of perturbation to τ_x in the forward model. There is no large scale spatial structure evident in the sensitivity of $\overline{D}$ to τ_x in the undamped 0.1° model (figure 4.13A), while in the damped 0.1° model there is a band of weak negative sensitivity at 47°N (figure 4.13B).

Sensitivity of $\overline{K}$ to τ_x

The sensitivity of $\overline{K}$ to τ_x is also positive, due to direct driving by the wind. The sensitivity of $\overline{K}$ to zonal windstress is highest in the centre of the channel, where the zonal velocity is highest, in both the 0.48° model and damped 0.1° models

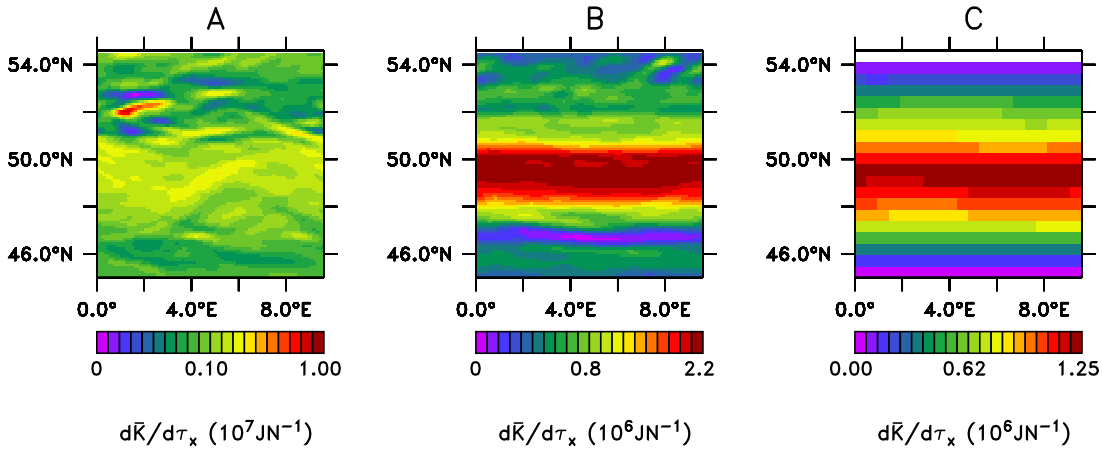

Figure 4.14: Sensitivity of $\overline{K}$ to τ_x as calculated by the adjoint model at T=0 in a 278 day integration. A: undamped 0.1° model B: damped 0.1° C: 0.48° model. Note difference in scales, between the panels.

(figure 4.14B and C). The band of high sensitivity in the centre of the channel can also be seen in the undamped 0.1° model after 278 days (figure 4.14A), but is much less distinct.

Sensitivity of $\overline{A}$ to τ_x

The sensitivity of $\overline{A}$ to τ_x is positive between 48°N and 50°N in the 0.48° model, negative between 45°N and 48°N in the 0.48° model and near zero elsewhere. In the damped 0.1° model there is a band of high positive sensitivity at 46.5°N and a less distinct band of negative sensitivity at 52°N. There is a lot of small scale structure near the northern boundary, but there may be an additional band of high sensitivity at 54°N. Again the sensitivity is near zero elsewhere. The physical explanation for the different response in the two models is not clear in this case, and further forward model experiments with perturbations to the surface windstress in this area are required to interpret the results. No structure is visible in the sensitivity in the undamped 0.1° model after 278 days.

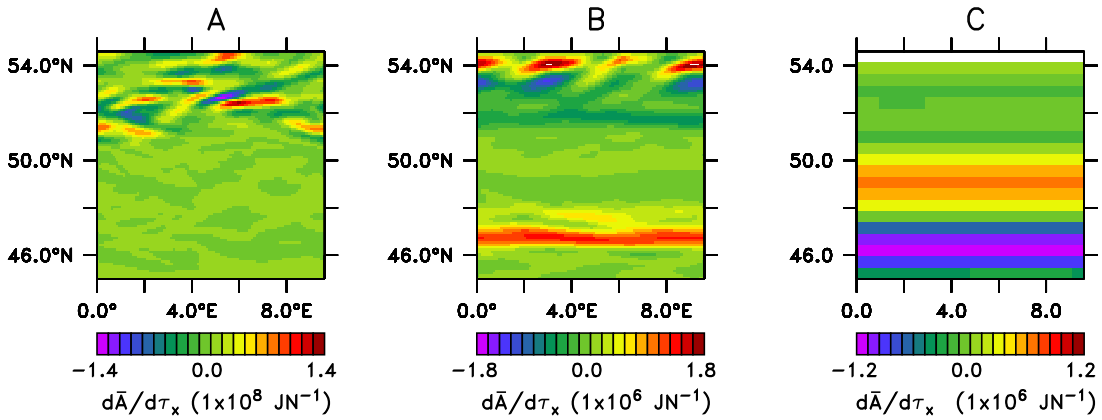

Figure 4.15: Sensitivity of $\overline{A}$ to τ_x as calculated by the adjoint model after 278 days of a 690 day integration. A: the undamped 0.1° model; B: the damped 0.1° model; C: the 0.48° model. Note the difference in scale between the panels.

4.3 Sensitivity after 690 Days

After 690 days the sensitivities calculated by the adjoint model in the 0.48° model look similar to those at 278 days for all quantities but tend to have a larger magnitude as they are sensitivities to perturbations applied over a longer time interval. There are more differences between the sensitivities calculated by the adjoint model after 278 and 690 days in the undamped and damped 0.1° models, and these are discussed in greater length in this section. As the differences between the two times are similar for most variables, here we concentrate on 3 fields, the sensitivity of $\overline{H}$ to SST*, the sensitivity of $\overline{K}$ to τ_x and the sensitivity of $\overline{D}$ to κ_v.

In both the 0.1° models there are small areas of extremely high sensitivity of $\overline{H}$ to SST^* that accumulate north of 52°N (figure 4.3A and B) after 278 days. After 690 days these areas of high sensitivity are still present, and have grown in magnitude (figure 4.16E and F). In the undamped 0.1° model the band of high sensitivity at 50°N is no longer visible after 690 days due to small scale structure in the interior of the channel (figure 4.16E). In the damped 0.1° model the band of sensitivity at 50°N is still visible after 690 days (figure 4.16F).

Unlike for the bands of high sensitivity associated with convection the small patches of sensitivity north of 52°N have no clear relation to processes in the forward model, and are unlikely to be the sensitivity of $\overline{H}$ to a non infinitessimal

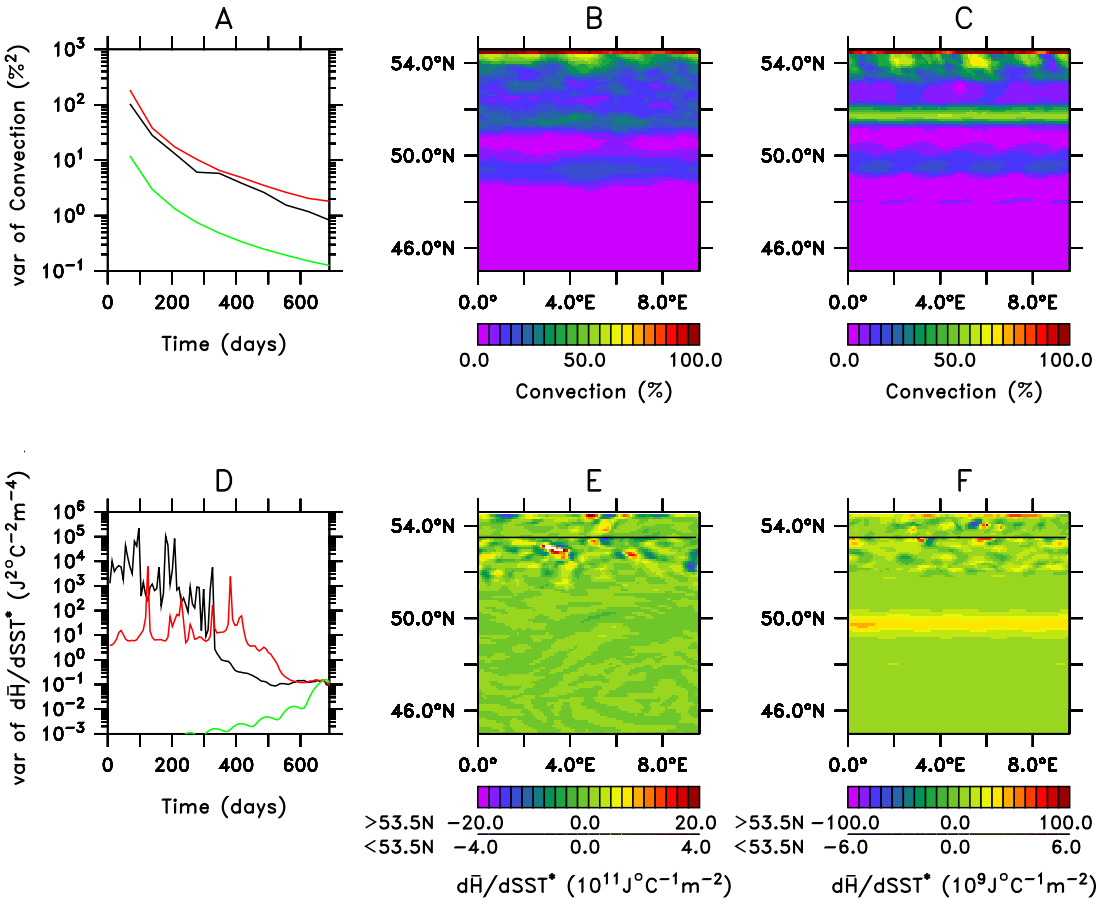

Figure 4.16: Sensitivity of $\overline{H}$ to SST* at $T = 0$ in the 690 day integration. A) Meridionally averaged zonal variance of time averaged convection as a function of integration time at 50m; black: undamped 0.1° model; red: damped 0.1° model; green: 0.48° model. B) Percentage of timesteps convection occurs between the top 2 model layers in the undamped 0.1° model, averaged over 690 days. C) Percentage of timesteps convection occurs between the top 2 model layers in the damped 0.1° model, averaged over 690 days. D) Meridionally averaged zonal variance of the sensitivity of $\overline{H}$ to SST* normalised by the mean sensitivity as a function of integration time; black: undamped 0.1° model; red: damped 0.1° model; green: 0.48° model. E) Sensitivity of $\overline{H}$ to SST* after 690 days in the undamped 0.1° model. F) Sensitivity of $\overline{H}$ to SST* after 690 days in the damped 0.1° model.

pertubation to SST*. The sensitivity of $\overline{H}$ to SST* is strongly associated with the presence of convection in the forward model. The channel and the forcing are zonally symmetric so that we expect that

$$\overline{C}'^2 = \left\langle \left(\frac{1}{T} \int_0^T C dt \right)^2 \right\rangle - \left\langle \frac{1}{T} \int_0^T C dt \right\rangle^2 \quad \begin{array}{l} \to 0 \\ \text{as } T \to \infty \end{array} \tag{4.5}$$

where $\overline{C}$, is the the percentage of timesteps where convection occurs between the top two model layers, $\langle\,\rangle$ denotes a zonal average, and T is the integration time. As the integration time is increased $\overline{C}'^2$ decreases at a similar rate in all three models (figure 4.16A) despite the differing amounts of eddy kinetic energy. Similarly it should be expected that the zonal variance of the sensitivity of $\overline{H}$ to SST* tends to zero as the integration time increases in the adjoint model and after 278 days the structure of the sensitivity of $\overline{H}$ to SST* is roughly zonally symmetric. The sensitivity of $\overline{H}$ to SST* grows rapidly backwards in time (figure 3.7A), so that the magnitude of $(\partial\overline{H}/\partial\text{SST}^*)'^2$ calculated by the adjoint model also grows rapidly backwards in time. For comparison with $\overline{C}'^2$ the sensitivity is normalised by the zonal mean before the variance is calculated. This quantity is defined as

$$\mathcal{H}_{\text{SST}^*} = \left(\frac{\partial\overline{H}/\partial\text{SST}^*}{\langle\partial\overline{H}/\partial\text{SST}^*\rangle} \right)'^2 \tag{4.6}$$

In the 0.48° model $\mathcal{H}_{\text{SST}^*}$ decreases backwards in time (figure 4.16D), while in the damped 0.1° model it increases backwards in time between 690 and 450 days and is roughly constant between 450 and 0 days. In the undamped 0.1° model $\mathcal{H}_{\text{SST}^*}$ increases backwards in time throughout the 690 days. A visual comparison of $\overline{C}$ and $\partial\overline{H}/\partial\text{SST}^*$ also shows that the small areas of high sensitivity above 52°N are not associated with any small scale structure in the convection (figure 4.16B,C and E,F). It seems likely that these areas of high sensitivity represent a loss of information in the sensitivities calculated by the adjoint model at small spatial scales.

The sensitivity of $\overline{D}$ to κ_v was very similar in all three models at 278 days, and

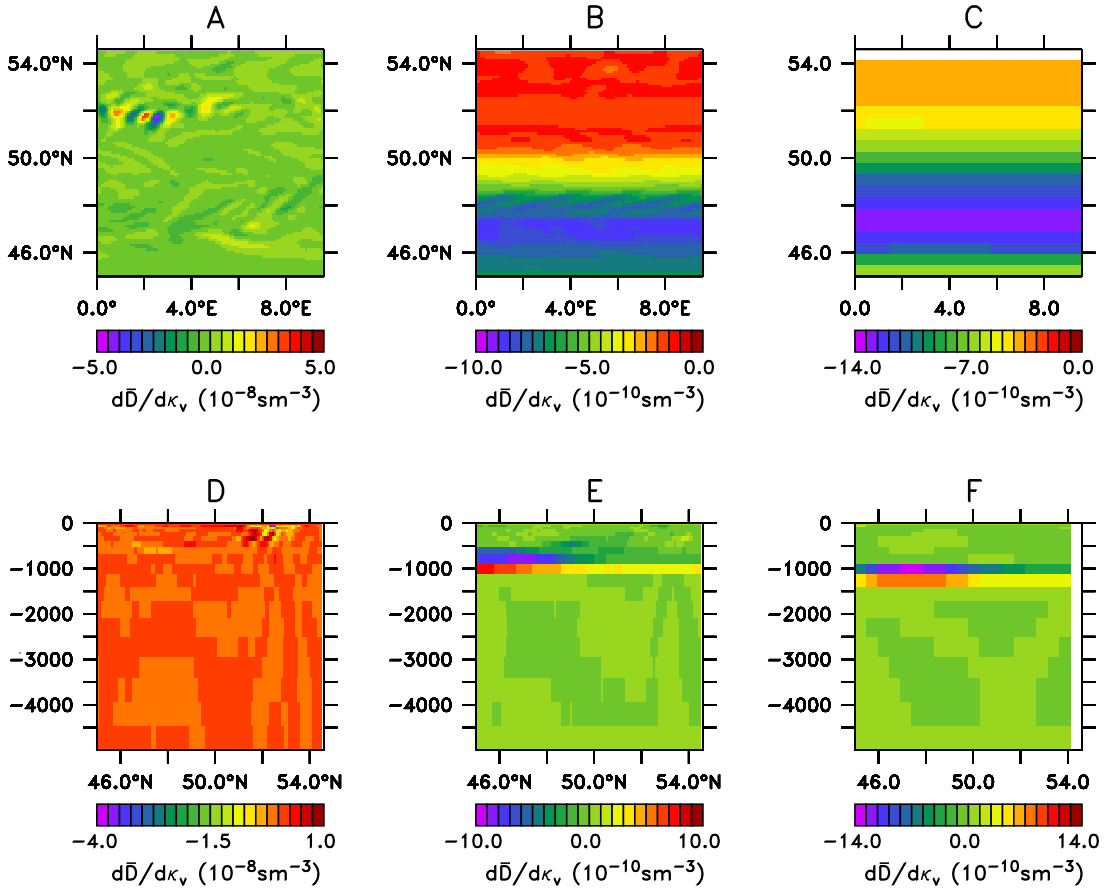

Figure 4.17: Sensitivity of $\overline{D}$ to κ_v as calculated by the adjoint model at $T = 0$ in a 690 day integration. A: Sensitivity in the undamped 0.1° model at 593m. B: Sensitivity in the damped 0.1° model at 774m C: Sensitivity in the 0.48° model at 992m. Bottom row (D-F): zonally averaged sensitivity; D: the undamped 0.1° model; E: the damped 0.1° model; F: the 0.48° model. Note the difference in scale between the panels.

was clearly related to dT/dz in the forward model. After 690 days the sensitivity of $\overline{D}$ to κ_v in the damped 0.1° model and 0.48° model has a similar pattern to that after 278 days (figure 4.17). However, after 690 days the sensitivity in the undamped 0.1° model is dominated by small scale structure (figure 4.17A,D). Again it is likely that these areas of high sensitivity represent the loss of information on small spatial scales.

The sensitivity of $\overline{K}$ to τ_x is again similar in the 0.48° model after 278 and 690 days (figures 4.18C and 4.14C). In the damped 0.1° model there is a faint band of high sensitivity at the position of the secondary zonal jet at 52°N after 278

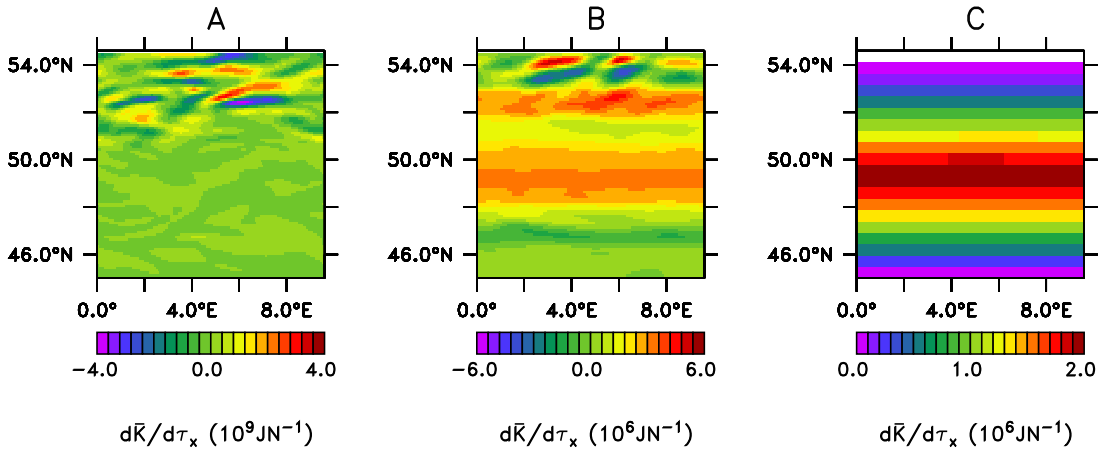

Figure 4.18: Sensitivity of $\overline{K}$ to τ_x as calculated by the adjoint model at T=0 in a 690 day integration. A: undamped 0.1° model B: damped 0.1° C: 0.48° model. Note difference in scales, between the panels.

days, but this is much more distinct after 690 days (figures 4.18B and 4.14B). In the undamped 0.1° model the sensitivity of $\overline{K}$ to τ_x is dominated by small scale spatial structures after 690 days, particularly north of 52°N (figure 4.18A).

4.4 Discussion

For most quantities, after 278 days, the spatial structures in the sensitivity of $\overline{H}$ to SST*, κ_v and τ_x and $\overline{K}$ to κ_v and τ_x as calculated by the adjoint to the chaotic undamped 0.1° model are similar to those in the less chaotic damped 0.1° model and unchaotic 0.48° models, although there is more small scale structure in the sensitivity in the undamped 0.1° model. There are clear physical explanations for most of the structures in the sensitivities and on this timescale the adjoint model is able to give useful information about the sensitivity of time averaged climate quantities at a higher resolution than could easily be achieved using perturbed forward model experiments.

After 690 days there is little structure remaining in the sensitivities calculated by the adjoint model, in the undamped 0.1° model, that can be related to processes in the forward model. This is true even for quantities such as the sensitivity of $\overline{H}$ to SST* where the adjoint model was able to give the correct sensitivity to

large scale perturbations in the forcing (chapter 3). The time over which useful information remains in the adjoint solution depends on the spatial scale of the information, with the adjoint model able to give the correct sensitivity to large scale changes in forcing for longer, although there may be little spatial structure in the sensitivities calculated by the adjoint model that obviously relates to processes in the forward model.

When small scale areas of high sensitivity, that do not relate to process in the forward model, develop in part of the model domain they do not necessarily dominate the solution in other regions, and useful information can still be gained from running the adjoint model. A good example of this are the areas of high sensitivity of $\overline{H}$ to SST* after 690 days in the damped 0.1° model, which do not effect the band of high sensitivity associated with convection in the centre of the channel at 50°N.

The location of areas of high sensitivity depends on the circulation in the forward model, so that the sensitivity fields after 278 days look different in the 3 different models. However, areas of high positive and negative sensitivity of the heat content, kinetic energy and thermocline depth are associated with similar features in the forward model circulation. Although it is possible that different processes may dominate the sensitivity in the different resolution models at longer times, this suggests that the sensitivity calculated by the adjoint to a non eddy resolving ocean model run over long timescale may give an indication of what the sensitivity will look like in an eddy resolving model over a longer timescale.

However, even on a short timescale the sensitivity of the available potential energy is very different in the 0.48° and 0.1° models. In the 0.48° model the sensitivity of the available potential energy is dominated by the sensitivity of the static stability, while in the 0.1° models it is dominated by the sensitivity to the meridional temperature gradient. As a result not only is the size of the sensitivity different in the different models, but it also has a different sign in some regions. Although the adjoint to the 0.48° model can be used to calculate sensitivities over a longer timescale than the adjoint to the undamped 0.1° model, a long integration of the adjoint to the 0.48° model could not be used to make

inferences about the sensitivity of the available potential energy in the undamped 0.1° model. The circulation in the 0.1° models is very different than in the 0.48° model; in particular there is a greater meridional temperature gradient across the front (section 2.3.4) and the thermocline is deeper (section 4.1.1). It is possible that the differences in the density structure cause the different sensitivities. Köhl and Willebrand (2002) proposed using the adjoint to a non eddy resolving model linearised about the state of an eddy resolving model averaged onto the coarser grid, for calculating sensitivities in an eddy resolving ocean model over long timescales. Although this method would still result in a lower meridional temperature gradient across the front, the circulation in the undamped 0.1° model averaged onto the 0.48° grid would be more similar to the circulation in the undamped 0.1° model than the circulation in the 0.48° model. It is therefore possible that this method could be used to calculate the correct sensitivity of the available potential energy over long time scales in the undamped 0.1° model. However, this needs to be tested in practice.

4.5 Summary

- Information remains in the spatial structure of the sensitivities calculated by the adjoint model for many times longer than the eddy timescale. In the undamped 0.1° model used here, information remains on a timescale of 9 months while the Eady growth rate in the centre of the jet is around 16 days.

- The adjoint model is able to give the correct sensitivity to large spatial scale perturbations in the forcing on a longer timescale than for smaller spatial scale perturbations. In the undamped 0.1° model used here, the adjoint model still gives the correct sensitivity to large scale perturbations for some quantities on a timescale of 2 years, although no information remains in the spatial structure of the sensitivities calculated by the adjoint model on this timescale.

- Different processes may dominate the sensitivity at different resolutions.The sensitivity of the available potential energy in the 0.48° model is dominated by the sensitivity of the static stability, while the sensitivity of the available potential energy in the 0.1° models is dominated by the sensitivity of the meridional density gradient. Conclusions drawn from the sensitivity in a non eddy resolving ocean model can not necessarily be applied to an eddy resolving ocean model.

Chapter 5

Conclusions and Outlook

In the 0.48° model the adjoint model is able to provide sensitivity information at the grid resolution on a timescale of a least 690 days. This is expected as the forward model is non chaotic. Chaos in the undamped 0.1° model means there is a time limit beyond which the adjoint model can no longer provide useful information. The non linear timescale of the 0.1° model is estimated to be around 200 days. At a slightly longer timescale of 278 days the adjoint model provides useful information in the spatial structure of the sensitivities for most quantities. After 690 days the adjoint model is able to give the sensitivity of the heat content to very large spatial scale perturbations to the sea surface temperature and the vertical diffusivity, but it is unable to provide information at a smaller spatial scale. Even at large spatial scales, after 690 days, the adjoint to the undamped 0.1° model is unable to give the correct sensitivity of the kinetic energy to any quantities and of the heat content to the surface windstress, which suggests that information remains for longer for thermodynamic quantities. Information remains in the spatial structure of the sensitivities in the damped 0.1° model on a timescale of 690 days. Although there are some signs that the sensitivities may not contain useful information on the smallest spatial scales this result is rather surprising as the non linear timescale in this model is estimated at around 280 days.

Three methods are used for assessing whether useful information remains in the results calculated by an adjoint model; these methods are: finite difference

gradient checks, forward model experiments where a larger perturbation is made to the forcing, and an assessment of whether it is believed the results are physically significant. Finite difference gradient checks may be affected by the same problems as the adjoint model and calculate the microscopic rather than the macroscopic sensitivity in a chaotic system. Very good agreement between adjoint and finite difference gradients, such as for the 0.48° model here, would indicate that the adjoint model results do contain useful information. However agreement between the sign and order of magnitude of the adjoint and finite difference gradients, as seen in the undamped 0.1° model here, does not necessarily indicate that useful information remains in the sensitivities calculated by the adjoint model.

Forward model experiments where a larger scale perturbation is made to the forcing allow us to see if a time averaged climate quantity has predictability of the second kind. However, they only show if there is useful information contained in the adjoint model at the largest spatial scales. As information loss occurs soonest at the smallest scales, these experiments probably provide an upper limit on when an adjoint model can provide useful information.

An assessment of whether the results from the adjoint model are physically significant is rather subjective. Sensitivities calculated by the adjoint to a non eddy resolving model ocean model, where very good agreement between adjoint and finite difference gradients suggests that that the sensitivities calculated by the adjoint model are likely to contain useful information, allows us to identify features in the forward model that are associated with high sensitivity, such as the presence of convection for the sensitivity to SST*. If these features are also associated with areas of high sensitivity in an eddy resolving model, it is likely that the sensitivity calculated by the adjoint model is also physically relevant in the eddy resolving model. This approach needs to be applied with caution as the processes governing the sensitivity of an output function are not necessarily the same in models of different resolution. This is clearly seen in the sensitivity of the available potential energy to SST* and κ_v in the 0.48° and 0.1° models. In the 0.48° model the sensitivity of the available potential energy is dominated by

the sensitivity of the static stability, while in the 0.1° model the sensitivity of the available potential energy is dominated by the sensitivity of the meridional density gradient.

A comparison of an eddy resolving and non eddy resolving model also fails to provide any information on what we expect the structure in the sensitivities to be at small spatial scales, as these are not resolved in the non eddy resolving model. As the channel and the forcing are zonally symmetric, and the time averaged state variables are also zonally symmetric it is expected that there will be no preferred location in the zonal direction for a perturbation to the forcing. It is thus expected that the sensitivities calculated by the adjoint model should be zonally symmetric as the integration time increases. Sensitivities calculated by the adjoint to the undamped 0.1° model become less zonally symmetric as the integration time increases, and are dominated by small scale structure, which is unlikely to be physically relevant. In the damped 0.1° model there are small areas of high sensitivity near the northern boundary for some quantities, and these are also thought not to be physically realistic.

Although explanations are given here for most of the structures in the sensitivity fields after 278 days, little explanation is offered for the sensitivity to the surface windstress. The sensitivity to the surface windstress is assumed to be physically significant, as it has a number of features we expect, i.e. zonal symmetry, and similar features between the different resolution models after 278 days. However, the results are difficult to interpret. Making perturbations to the surface windstress in the forward model that are shaped like the sensitivity patterns in the adjoint model would help in interpreting the results.

As information is present for longer over larger spatial scales, a simple comparison of the TLM solution with the difference between 2 perturbed non linear integrations as proposed by Kleist and Morgan (2005) is likely to give a conservative estimate of when the adjoint model can be used. Short spatial scales dominate the TLM solution in a chaotic system (Tanguay et al., 1995), so that the correlation between the TLM solution and the difference between 2 perturbed non linear model integrations decreases rapidly, in both 0.1° models. The TLM

model suggests that after 200 days the linearisation should no longer be valid in the undamped 0.1° model, while even at 278 days considerable climatically relevant information remains in the spatial structure of sensitivities calculated by the adjoint method to most quantities. Similarly the TLM results imply a 250 day timescale beyond which the linearisation fails in the damped 0.1° model when, even at 690 days there is useful information remaining in the spatial structure of sensitivities calculated by the adjoint method.

The useful information remaining in the spatial structures in the adjoint sensitivities out to 278 days, suggests that the limit to length of the data assimilation window caused by the presence of chaotic eddies in a 0.1° model at 45°N may be as long as 9 months. This is longer than the 0.2 years over which Lea et al. (2002) found that the sensitivities calculated by an adjoint model gave a reasonable estimate of the climate sensitivity, but far shorter than the 2 years over which Gebbie (2004) was able to successfully assimilate data into an eddy resolving model of the subtropical North Atlantic. It is still not clear why the adjoint method was so successful in Gebbie (2004)'s case.

The model used here has strong restoring to SST^* at the surface. This causes significant damping at the surface and may reduce the eddy kinetic energy in the channel and increase the timescale over which the adjoint model can provide useful information. Hogan and Hurlburt (2000) found that a resolution of 1/32° was necessary to properly represent mesoscale variability in the ocean, which is a much higher resolution than used in the present study. As smaller spatial scales tend to have shorter timescales, an increase in the resolution is likely to reduce the time over which useful information can be obtained by using an adjoint model. The effect of the use of the CD coupling scheme on reducing the exponential growth of the sensitivities calculated by the adjoint model in this study shows that this timescale is sensitive to parameterisations and computational schemes used. The predictability of a chaotic system varies along its trajectory in phase space (Palmer, 2000), and it is possible that starting from a different initial condition could change the timescale over which the adjoint model is able to give useful information about the system.

This study differs from previous work looking at the use of the adjoint method with eddy resolving ocean general circulation models as it compares sensitivities calculated using different resolution models, and used a range of output functions. This has shown that useful information remains in the sensitivities calculated by the adjoint model for longer for some quantities than for others, and that sensitivities calculated by a non eddy resolving ocean model may be very different to those in an eddy resolving model, even when the sensitivities in the eddy resolving model still contain useful information. The work also provides a guide of how to determine if sensitivities calculated by an adjoint model contain useful information.

In the present study information remains in the sensitivities calculated by the adjoint model for longer for quantities than involve thermodynamic variables only. It is not clear if this is due to these being slower components of the system, or if it is due to the strength of the forcing. The potential for the strength of the forcing to affect the time over which useful information remains in the sensitivities calculated by the adjoint model should be the subject of further study. This could be done within the current framework by varying the strength of the restoring timescale for SST^*, or the strength of the windstress, and could potentially give an insight into how the strength of the forcing affects the predictability of the system.

Climate sensitivity studies are usually performed over a much longer timescale than the timescale over which useful information was shown to remain in the sensitivity of the chaotic undamped eddy resolving model in this study. It is therefore unlikely that climate sensitivity studies would ever be possible using the adjoint to an eddy resolving GCM. This problem is likely to become more apparent in the future as greater computing power means that climate models will be run at higher resolution and will generally resolve chaotic eddies. However, the current work suggests two approaches by which the adjoint method could be used to gain information about the sensitivity of an eddy resolving OGCM at long timescales. In most cases regions of high sensitivity in the eddy resolving and non eddy ocean models coincide with similar features in the forward model. A

multi grid approach could then be used where a non eddy resolving ocean model and its adjoint are used to identify regions where the sensitivity was likely to be high in an eddy resolving ocean model. This could be tested using perturbed forward model experiments in both the eddy resolving and non eddy resolving ocean models to see if the response to the perturbation was the same. The second method is suggested by the greater timescale over which information remains in damped eddy resolving model in this study. A parameterisation scheme that reduces the growth of the sensitivities calculated by the adjoint model could be introduced into either forward and adjoint models, or into the adjoint model only in a similar method to Köhl and Willebrand (2003). This has potential benefits over the coarse resolution adjoint model approach of Köhl and Willebrand (2003) as horizontal density gradients would not be reduced by the averaging onto a coarser grid. A comparison between these methods would be a useful additional study.

Bibliography

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

Alves, M., de Verdière, A. C., 1999. Instability dynamics of a subtropical jet and applications to the Azores front current system: Eddy-driven mean flow. J. Phys. Oceanogr. 29, 837–864.

Bugnion, V., Hill, C., 2006. Equilibration mechanisms in an adjoint ocean general circulation model. Ocean Dynamics 56, 51–61.

Bugnion, V., Hill, C., Stone, P. H., 2006a. An adjoint analysis of the meridional overturning circulation in a hybrid coupled model. J. Climate 19, 3751–3767.

Bugnion, V., Hill, C., Stone, P. H., 2006b. An adjoint analysis of the meridional overturning circulation in an ocean model. J. Climate 19, 3732–3750.

Errico, R. M., 1997. What is an adjoint model? Bulletin of the American Meteorological Society 78, 2577–2591.

Errico, R. M., Vukicevic, T., 1992. Sensitivity analysis using an adjoint of the PSU-NCAR mesoscale model. Monthly Weather Review 120, 1644–1660.

Eyink, G. L., Haine, T. W. N., Lea, D. J., 2004. Ruelle's linear response formula, ensemble adjoint schemes and Lèvy flights. Nonlinearity 17, 1867–1889.

Fukumori, I., Lee, T., Cheng, B., Menemenlis, D., 2004. The origin, pathway, and destination of Niño-3 water estimated by a simulated passive tracer and its adjoint. J. Phys. Oceanogr. 34, 582–604.

Gebbie, G., 2004. Subduction in an eddy-resolving state estimate of the northeast atlantic ocean. Ph.D. thesis, MIT/WHOI.

Giering, R., 1999. Tangent linear and Adjoint Model Compiler Users Manual.

Giering, R., 2005. Transformation of Algorithms in Fortran Users Manual. http://www.FastOpt.com/taf.

Giering, R., Kaminski, T., 1998. Recipes for adjoint code construction. ACM Trans. Math. Soft. 24, 437–474.

Gill, A. E., 1982. Atmosphere-Ocean Dynamics. Academic Press.

Griewank, A., 1992. Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation. Optimization Methods and Software 1, 35–54.

Griffies, S., Böning, C., Bryan, F., Chassignet, E., Gerdes, R., Hasumi, H., Hirst, A., A-M., T., Webb, D., 2000. Developments in ocean climate modelling. Ocean Modelling 2, 123–192.

Hall, M. C. G., Cacuci, D. G., Schlesinger, M. E., 1982. Sensitivity analysis of a radiative-convective model by the adjoint method. J. Atmos. Sci., 2038–2050.

Haney, R. L., 1971. Surface thermal boundary condition for ocean circulation models. J. Phys. Oceanogr. 1, 241–248.

Hogan, P. J., Hurlburt, H. E., 2000. Impact of upper ocean-topographical coupling and isopycnal outcropping in Japan/East Sea models with 1/8° to 1/64° resolution. J. Phys. Oceanogr., 2535–2561.

Huang, B., Stone, P. H., Hill, C., 2003a. Sensitivities of deep-ocean heat uptake and heat content to surface fluxes and subgrid-scale parameters in an ocean general circulation model with idealized geometry. J. Geophys. Res 108, doi:10.1029/2001JC001218.

Huang, B., Stone, P. H., Sokolov, A. P., 2003b. Deep ocean heat uptake in transient climate change. J. Climate 16, 1352–1363.

Junge, M. M., Haine, T. W. N., 2001. Mechanisms of North Atlantic wintertime sea surface temperature anomalies. Journal of Climate, 4560–4572.

Kleist, D. T., Morgan, M. C., 2005. Interpretation of the structure and evolution of adjoint-derived forecast sensitivity gradients. Monthly Weather Review 133, 466–484.

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

Köhl, A., 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), doi:10.1029/2001JC000982.

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

Lea, D. J., Haine, T. W. N., Allen, M. R., Hansen, J. A., 2002. Sensitivity analysis of the climate of a chaotic ocean circulation model. Quart. J. Roy. Meteor. Soc. 128, 2587–2605.

Lorenz, E. N., 1955. Available potential energy and the maintenance of the general circulation. Tellus 7, 157–167.

Lorenz, E. N., 1963. Deterministic nonperiodic flow. J. Atmos. Sci. 20, 130–141.

Lorenz, E. N., 1975. Climate predictability. In: The physical basis of climate modelling WMO (GARP Publication Series vol 16). World Meteorological Organization, pp. 132–136.

Marotzke, J., 2002. Lecture notes for climate dynamics course at Southampton University.

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

Marotzke, J., Wunsch, C., 1993. Finding the steady state of a general circulation model through data assimilation: Application to the North Atlantic Ocean. J. Geophys. Res, 20149–20167.

Marshall, J., Adcroft, A., Campin, J.-M., Heimbach, P., Molod, A., Dutkiewicz, S., Hill, H., Losch, M., Fox-Kemper, B., Menemenlis, D., Ferreira, D., Hill, E., Follows, M., Hill, C., Evangelinos, C., Forget, G., 2004. MITgcm User Manual.

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

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

McLay, F., Marotzke, J., 2006. Limitations on the use of adjoint models with eddy resolving ocean general circulation models. submitted to Ocean Modelling.

Moore, A. M., 1991. Data assimilation in a quasi-geostrophic open-ocean model of the Gulf Stream region using the adjoint method. J. Phys. Oceanogr., 398–427.

Moore, A. M., Arango, H. G., Di Lorenzo, E., Cornuelle, B. D., Miller, A. J., Neilson, D. J., 2004. A comprehensive ocean prediction and analysis system based on the tangent linear and adjoint of a regional ocean model. Ocean Modelling 7, 227–258.

Morrow, R., De Mey, P., 1995. Adjoint assimilation of altimetric, surface drifter, and hydrographic data in a quasi-geostrophic model of the Azores current. J. Geophys. Res 100, 25,007–25,025.

Oort, A. H., Ascher, S. C., Levitus, S., Peixóto, 1989. New estimates of available potential energy in the world ocean. J. Geophys. Res C94, 3187–3200.

Palmer, T. N., 2000. Predicting uncertainty in forecasts of weather and climate. Rep. Prog. Phys. 63, 71–116.

Peixoto, J., Oort, A., 1992. Physics of Climate. Springer-Verlag.

Press, W H.and Teukolsky, S. A., Vetterling, W. T., Flannery, B. P., 1992. Numerical Recipes in FORTRAN 77: The Art of Scientific Computing, 2nd Edition. Cambridge University Press.

Rivière, P., Treguier, A. M., Klein, P., 2004. Effects of bottom friction on nonlinear equilibrium of an ocean baroclinic jet. J. Phys. Oceanogr. 34, 416–432.

Schiller, A., Willebrand, J., 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.

Schröter, J., Seiler, U., Wenzel, M., 1993. 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., Mintz, Y., 1977. Numerical simulation of the Gulf Stream and mid-ocean eddies. J. Phys. Oceanogr., 208–230.

Sirkes, Z., Tziperman, E., 1997. Finite difference of adjoint or adjoint of finite difference? Monthly Weather Review, 3373–3378.

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

Smith, R. D., Maltrud, M. E., Bryan, F. O., Hecht, M., 2000. Numerical simulation of the North Atlantic at $\frac{1}{10}^{\circ}$. J. Phys. Oceanogr., 1532–1561.

Stainforth, D. A., Aina, T., Christensen, C., Collins, M., Faull, N., Frame, D. J., Kettleborough, J. A., Knight, S., Martin, A., Murphy, J. M., Piani, C., Sexton, D., Smith, L. A., Spicer, R. A., Thorpe, A. J., Allen, M. R., 2005. Uncertainty in predictions of the climate response to rising levels of greenhouse gases. Nature 433, 403–406.

Stammer, D., Wunsch, C., Giering, R., Eckert, C., Heimbach, P., Marotzke, J., Adcroft, A., Hill, C. N., Marshall, J., 2002. Global ocean circulation during 1992-1997, estimated from ocean observations and a general circulation model. J. Geophys. Res 107 (C9), doi:10.1029/2001JC000888.

Stensrud, D. J., Bao, J., 1992. Behaviors of variation and nudging assimilation techniques with a chaotic low-order model. Monthly Weather Review 120, 3016–3028.

Stommel, H., 1961. Thermohaline convection with two stable regimes of flow. Tellus 2, 224–230.

Strogatz, S. H., 1994. Nonlinear Dynamics and Chaos. Westview Press.

Tanguay, M., Bartello, P., Gauthier, P., 1995. Four-dimensional data assimilation with a wide range of scales. Tellus 47A, 974–997.

Thuburn, J., 2005. Climate sensitivities via a Fokker-Planck adjoint approach. Quart. J. Roy. Meteor. Soc. 131, 73–92.

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

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

Tziperman, E., Thacker, W. C., Long, R. B., Hwang, S., Rintoul, S. R., 1992b.

Oceanic data analysis using a general circulation model. part ii: A North Atlantic model. J. Phys. Oceanogr. 22, 1458–1485.

van Oldenborgh, G. J., Burgers, G., Venzke, S., Eckert, C., Giering, R., 1999. Tracking down the ENSO delayed oscillator with an adjoint OGCM. Monthly Weather Review 127, 1477–1496.

Zhu, J., Kamachi, M., 2000. The role of time step size in numerical stability of tangent linear models. Monthly Weather Review 128, 1562–1572.